

KNOWLEDGE FROM KNOWLEDGE: AN ESSAY ON INFERENTIAL KNOWLEDGE

BY RODRIGO MARTINS BORGES

A dissertation submitted to the
Graduate School—New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Peter D. Klein

and approved by

New Brunswick, New Jersey

October, 2015

ABSTRACT OF THE DISSERTATION

Knowledge from Knowledge: An Essay on Inferential Knowledge

by **Rodrigo Martins Borges**

Dissertation Director: Peter D. Klein

Under what conditions do we have inferential knowledge? I propose and defend the following principle: S knows that p via inference only if S knows all the premises essentially involved in her inference in support of p - “KFK” for short. Even though KFK is at least tacitly endorsed by many figures in the history of philosophy, from Aristotle through Descartes, and Kant to Bertrand Russell – and, more recently, by David Armstrong – KFK has fallen into disfavor among epistemologists over the past fifty years. In response to Edmund Gettier’s legendary paper, many have proposed views according to which one’s reasoning is a source of knowledge even if one fails to know some or all premises essentially involved in one’s reasoning, while others have given up offering a theory of inferential *knowledge* and have focused on reasoning as a source of *justified belief* instead. Unfortunately, these accounts that deal with inferential knowledge are problematic; they cannot, for example, fully explain our common practice of evaluating negatively inferences with unknown premises. They also seem to rely on

an overly narrow understanding of the so-called “Gettier Problem”. This dissertation aims at updating the approach to reasoning that was popular before Gettier by building on the framework of Timothy Williamson’s “knowledge-first” epistemology. The aim is to carefully elaborate the view that reasoning yields knowledge only if all the premises essentially involved are also known.

Acknowledgements

My work has benefited immensely from my exchanges with a great number of people. First and foremost, Peter Klein has been not only the best advisor anyone could hope for (supportive, patient, challenging), but he is also the best philosopher with whom I have ever discussed philosophy. Suffice it to say that, if I were to give *full* credit to Peter's supervision, I would probably have to thank him in a footnote to virtually every paragraph in this dissertation. Needless to say, he does not agree with everything I say here and I am also fully responsible for any mistakes in the text below. I would also like to thank the other members of my committee, Ernie Sosa, Duncan Pritchard and Claudio de Almeida. They all read different versions of some or most of the material in this dissertation and they have all contributed to making it a much better work than I could have by myself. Finally, I list (in no particular order), other people who, through conversations, written comments, emails or even facebook posts helped me write this dissertation: Luis Rosa, David Black, Tiegue Vieira, Daniel Rubio, Lisa Mirachi, Tobias Wilsch, Marcello Antosh, Bob Beddor, Thomas Blanchard, Tim Campbell, Will Fleisher, Georgi Gardiner, Luvell Anderson, Matthew Benton, Simon Goldstein, Ricardo Mena, Alex Morgan, Pamela Robinson, David Rose, Mary Salvaggio, Justin Sharber, Kurt Sylvan, Guido Melchior, Carlotta Pavese, Martha Bolton, Ruth Chang, Andy Egan, Branden Fitelson, Anthony Gillies, Alvin Goldman, Doug Husak, Jeff King, Howard McGary, Brian McLaughlin, Jeff McMahan, Jonathan Schaffer, Susanna Schellenberg, Steve Stich, Larry Temkin, Michael Veber, Linda Zagzebski, Jonathan Vogel, Anthony Brueckner, Sherri Roush, Jennifer Lackey, Juan Comesaña, E.J Coffman, John Hawthorne, Fred

Adams, Timothy Williamson, Peter Baumann, Stephen Hetherington, Ram Neta and Stephen Grimm. I also need to acknowledge the institutional support I received from Rutgers, particularly, from the staff in the Philosophy Department. In the very practical sense of the word “possible”, this dissertation would not have been possible without the hard work and competence of Pauline Mitchell, Jean Urteil and, especially, the incomparable Mercedes Diaz. At home, my wife Lizette Nieves has been the most supportive and encouraging partner anyone can ask for. She has been very patient, and for that I will be forever grateful. I am indebted also to my father João T. Borges, for his unconditional support all the way from Brazil. My gratitude also to Melanie Crisfield for proofreading the dissertation. The research for this dissertation was partially funded by the CAPES/Fulbright Commission through a four-year fellowship (2007-2011). I am very grateful for their support.

Dedication

In loving memory of Norma Martins Borges (1942-2012).

Table of Contents

Abstract	ii
Acknowledgements	iv
Dedication	vi
1. Introduction	1
1.1. What Knowledge Is	2
1.2. What Is To Come	15
2. Knowledge from Knowledge	16
2.1. Clarifying KFK	24
3. KFK, E=K, and The Knowledge Norm of Inference	38
3.1. From E=K to KFK	38
3.2. Objection: No Gettier Problem	40
3.3. Objection: Dogmatism	52
3.3.1. Dogmatism: Synchronic and Diachronic	54
3.3.2. Harman's Argument for Diachronic Dogmatism	57
The Junk Knowledge Response to Diachronic Dogmatism	60
3.3.3. Kripke's Argument for Synchronic Dogmatism	66
A Failed Response to Synchronic Dogmatism	69
3.3.4. Dogmatism Depuzzled	74
3.3.5. A new argument for dogmatism?	83
3.3.6. Taking Stock	85
3.4. The Knowledge Norm of Inference	86

3.4.1. Knowledge Norms: Assertion, Action and Inference	91
4. KFK and the Gettier Problem	106
4.1. The Gettier Conjecture	107
4.2. What counts as a Gettier case?	109
4.3. Two Gettier Problems	115
4.4. Some Caveats	116
4.5. A recent ancestor of the Gettier Conjecture	118
5. Epistemic Luck	123
5.1. Anti-Luck Epistemology	123
5.2. The Appeal To The Concept Of Epistemic Luck	125
5.3. Precisifying The Anti-Luck View	129
5.4. Against Anti-Luck Epistemology	131
5.4.1. Lucky But Safe True Belief: Against Epistemic Luck	131
5.4.2. Lucky But Counterfactually Robust Events: Against Generic Luck	133
5.5. Gettier-Lucky Knowledge?	134
5.5.1. The Argument for “Gettierism”	135
5.5.2. An Alleged case of Gettier-Lucky Knowledge	136
5.5.3. What is lucky about Gettier-Lucky belief?	140

Chapter 1

Introduction

The main goal of the following essay is to defend the principle about inferential knowledge that I call “Knowledge from knowledge”:

(KFK) S knows that p via inference or reasoning only if S knows all the premises essentially involved in her inference in support of p.

I explain how this principle helps solve the Gettier Problem; I argue that it is entailed by the defeasibility account of knowledge and Williamson’s $E=K$; and I show how it can motivate a plausible norm of appropriate inference. As far as I know, this thesis amounts to the most extensive exploration of KFK in the literature today.

Since an essay is not a philosophical work without disputing other philosopher’s views, I also argue against the most popular account of epistemic and non-epistemic luck. The discussion of this issue is related to KFK to the extent that KFK provides a competing explanation of the Gettier phenomenon. I argue that, on grounds independent of KFK, anti-luck epistemology does not provide an adequate account of luck, epistemic or not.

In the remainder of this introduction, I will do two things. First, in the following section, I will establish some criteria for what knowledge is. Like knowledge-first epistemologists, I am suspicious of the project of analyzing “knowledge”, and I take knowledge that p to entail probability 1 of p. I end the introduction with an overview of the chapters that follow.

An important characteristic of this thesis (some will say “an important short-coming”) is that it assumes a certain form of externalism about evidence and justification (and, consequently, knowledge). So, for example, I assume that evidence and justification are not *luminous* states (i.e., in some circumstances e is evidence for h for S , or S is justified in believing that p , but S is not in a position to know that e is evidence for h , or that S is justified in believing that p). I do not directly argue in favor of this claim. Even though I discuss how I understand evidence and justification in chapters 3 and 4, for the most part I simply assume that those states are not luminous, and I use those notions to try to understand the Gettier Problem and the Dogmatist Paradox, amongst others. So, for the most part, I ask you, the reader, to accept (at least provisionally) this mild form of externalism, and to explore with me some of its consequences. If I am right, this exercise will yield important theoretical dividends we may want to use in the purchase of those assumptions.

1.1 What Knowledge Is

Epistemologists traditionally answer the question “What is knowledge?” by offering a definition of the *concept* “knowledge”. I, like others,¹ am suspicious of the claim that the most theoretically fruitful answer to the question “What is knowledge?” will take the form of a definition. In this section, I will offer the sketch of an argument in support of this suspicion, and later try to answer this central question in philosophy by arguing for plausible criteria something deserving of the name “knowledge” should satisfy.

I do not have anything exceptionally deep to say about why attempting to define “knowledge” might not be the best strategy to adopt when trying to give an account of what knowledge is. Giving such a “deep” argument is, as they

¹Most prominently Williamson [2000], but see also Dretske [1981a].

say, “above my pay grade”. However, I would like to stress the pedestrian point that different people giving a definition of x might not only have different goals in mind, but might also go about pursuing this in different ways.² This is a fairly simple point, I think, but one that has important consequences. I want to stress one particular consequence here. One often hears that, if one should not try to offer a definition of “knowledge”, then it is not clear what one should do when answering the question “What is knowledge?”. The remark is usually a veiled challenge to whomever wants to theorize about knowledge without offering a definition of “knowledge.” The assumption is that the epistemologist offering a definition of “knowledge” is in a better methodological position than the epistemologist who does not offer such a definition. The problem with this assumption is that it predicts that there is one “true” methodology to be followed, and that this methodology will (luck permitting) yield *the* right definition of “knowledge”. However, if there is no uniform methodology every epistemologist employs when offering a definition of “knowledge,” then it is not clear that the person offering a definition of “knowledge” has the methodological high ground. And one should not assume without argument, that she does.

Traditionally, the *purpose* or *goal* of a definition of “knowledge” is to offer a necessary truth that elucidates the nature of knowledge itself.³ The claim that, necessarily, knowledge is justified true belief is an example of a proposed necessary truth which tries to capture the nature of knowledge.

In many cases the goal of offering necessary truths about the nature of knowledge is pursued via the *method* of truth conditions whereby necessary and sufficient conditions for something to be known by someone are offered and tested against potential counterexamples. When taken together, the goal of finding a

²cf.Zagzebski [1999].

³Zagzebski [1999].

necessary truth capturing the nature of knowledge and the method of truth conditions yield proposals such as:

(JTB) Necessarily, S knows that p iff S has a true justified belief that p.

However, some reject either this traditional goal or the traditional method (or both). For example, even though Hilary Kornblith⁴ takes knowledge to be a natural kind - and, thus, sees knowledge as being what it is essentially (necessarily), not accidentally - he rejects the truth conditional methodology and argues that knowledge should be a subject of empirical investigation. As a natural kind such as gold or water, knowledge's essential (necessary) nature can be revealed by empirical investigation. The point here is to show that the enterprise of defining knowledge can be undertaken in different ways; I am not favoring in any way Kornblith's methodology.

Alvin Goldman, on the other hand, rejects the view that the nature of knowledge is captured by a necessary truth but accepts the methodology of truth conditions.⁵ According to him, "knowledge", like all concepts, is a psychological entity and, as such, contingent, not necessary.⁶ Goldman, therefore, accepts the methodology of truth conditions and counterexamples, but rejects the traditional goal of putting forward a necessary truth about the nature of knowledge.

The conclusion suggested by this brief excursion into how some epistemologists answer the question "What is knowledge?" is that there is no single goal to be pursued or single methodology to be deployed when offering a definition of "knowledge".⁷ Thus, the assumption that we know perfectly well what we should

⁴cf.Kornblith [2002].

⁵Goldman [1986, 2012].

⁶Goldman [2010].

⁷cf. Zagzebski [1999] for a detailed discussion of different epistemologists' methodologies and goals when offering a definition of "knowledge". According to her, the fact that different philosophers are trying to define "knowledge" is not evidence that they have the same goal or methodology in mind.

be doing if we were defining “knowledge” but would have “no idea” what to say about knowledge if we were not defining its concept is at best dubious and at worst false. If one thinks that defining knowledge *is* a discrete activity with clear goals and methodology, then one needs an argument to show that this is the case, for we do not get that impression if we merely look at how epistemologists actually conduct themselves while doing epistemology. The very practice of epistemology shows that one can consistently define “knowledge” in importantly different ways.

That being said, what makes the claim that those defining “knowledge” have the methodological high ground dubious or false is the assumption about definitions, not the assumption about how to understand knowledge without offering a definition. Maybe people who think they have the methodological high ground project a misleading impression of unity in their ranks, but they may be right about the enemy being in disarray. If this were the case, then the traditional epistemologist would still have the methodological high ground, if only because their view of what knowledge is would be the only viable one.

I think this conclusion is mistaken. Still, “The proof is in the pudding”, as they say, and I will *show*, rather than *tell*, that defining “knowledge” is not the only viable answer to “What is knowledge?”. I will take the considerations I am about to provide to say something about the nature of knowledge, and I do take those considerations to amount to necessary truths. So, if we are keeping score, the view of knowledge I am about to sketch could be described, I think, as accepting the *purpose* or *goal* of the activity usually associated with defining “knowledge” while *rejecting* its truth conditional methodology.

Peter Unger famously argued that “knowledge” is an absolute concept, akin to “flat” and “empty.”⁸ They are absolute because they do not admit of degrees and because their correct application presupposes the complete absence of something.

⁸Unger [1975]. See also Dretske [1981a,b]. The kind of knowledge at issue here is *propositional* knowledge, rather than, say, knowledge by acquaintance. It makes sense to say that I know the Eiffel Tower better than Sam if “know” here is shorthand for “acquainted with”.

So, if x knows that something is the case, or if x is flat or empty, then there is no y such that y knows better/worse than x , or is such that y is flatter or emptier than x .⁹ What is more, if x is empty, flat, or knowledgeable, then there is nothing inside x , there are no bumps on x , and there is no doubt whether x . This latter fact about absolute concepts explains why it is usually infelicitous to assert sentences of the following forms:

- i x is empty but there are things inside of x .
- ii x is flat but there are bumps on x .
- iii x knows that p but there are doubts as to whether p .

If “empty”, “flat” and “knows” are absolute concepts in the way described above, then asserting i-iii is infelicitous because the second conjunct of those claims is logically incompatible with the first conjunct: that x is empty entails that there is nothing inside x ; that x is flat entails that there is no bump on x ; and, that x knows that p entails that there are no doubts as to whether p .

In fact, if “knowledge” is an absolute concept, then I can explain why asserting the following also seems infelicitous:

- iv S knows that p , but there is a chance, given her evidence, that $\neg p$.
- v S knows that p , but she has not eliminated all possibilities that $\neg p$.

A direct implication of “knowledge” being an absolute concept is that, if S knows that p , then p has an epistemic probability of 1 for S .¹⁰ Consider the following

⁹Those words express, at most, relative nearness to the absolute state in question. For example, to say that x is flatter than y is to say that x is closer to being flat than y is. Strictly speaking, something is flat only if it has no bumps or irregularities, no matter how minute those bumps or irregularities are. Flatness does not admit of degrees, although a surface’s nearness to being flat does.

¹⁰See Fantl and McGrath [2012] for a view which tries to accommodate the infelicity of asserting things like iii-v while rejecting the Ungerian hypothesis that “knowledge” is an absolute concept. See Brown [2014] for a convincing case against Fantl’s and McGrath’s way of dealing with those claims.

argument:¹¹

1. $Ksp \Rightarrow \neg\Diamond\neg p$. [from the absoluteness of “knowledge”]
2. $\neg\Diamond\neg p \Leftrightarrow P(\neg p)=0$. [probability distributes over possibility]
3. $Ksp \Rightarrow P(\neg p)=0$ [from 1,2 by the transitivity of “ \Rightarrow ”]
4. $P(\phi)=0 \Leftrightarrow P(\neg\phi)=1$ [the probability of mutually exclusive events adds up to 1]
5. $Ksp \Rightarrow P(p)=1$ [from 3,4 by substitution of equivalents]

So, if S knowing that p entails that p has probability 1 for S, then what kind of probability are we talking about here? And, relatedly, what kind of possibility does “ \Diamond ” refer to in this argument?

Clearly, we should not be talking about mere *objective chance*, for perhaps with the exception of future states of the world, a state has an objective chance of 1 or 0. We should not be talking about *subjective credences* either, for even though a high enough degree of subjective credence (say, higher than .5) is necessary for knowledge (I cannot know that p if I do not believe that p), not even the highest degree of subjective credence in a true proposition is sufficient for knowledge, for one might give this true proposition credence 1 because one is mad rather than because of anything having to do with that proposition’s truth. The relevant notion of probability here should be a measure of the degree to which some body of evidence supports some proposition. Call this intuitive measure “epistemic probability.” So, if S knows that p, then S’s evidence supports p to the highest degree; i.e., 1. In symbols: $P(\phi)=P(\phi | e)$, where “e” refers to a contextually relevant body of total evidence. As usual, $P(\phi | e)=P(\phi\&e)\backslash P(e)$ and $P(e)>0$. “P” in the argument above refers to this latter measure of evidential support.

¹¹Dodd [2011] makes a similar argument.

What about “ \diamond ”? How should we understand this notion? We should understand it like this: “ $\diamond\phi$ ” is true iff $(\phi \& e)$ is the case, for some contextually relevant body of evidence e . In other words, “ \diamond ” deals with *epistemic possibility* and, as such, it has to do with what possibilities are compatible with a certain body of total evidence.

When fully spelled out, premise 1 in the argument above reads as:

1. $Ksp \Rightarrow \neg\diamond(e \& \neg p)$.¹²

If one knows that p , then it is not possible that both one’s total evidence e and $\neg p$ are true.¹³ According to the view of what knowledge is I am articulating here, “knowledge” is an absolute concept. This entails that, if one knows that p , then one’s evidence eliminates the possibility that $\neg p$ and, consequently, that p has an (epistemic) probability of 1.

But wait. There’s more!¹⁴ There is still another argument in favor of the claim that “knowledge” is an absolute concept. Given our argument above, then knowing that p entails that p has probability 1 on one’s evidence. Now, one can show that the following plausible principle requires that knowledge have probability 1:

(Conjunction Principle) If S knows that p , S knows that q , and S deduces $(p \& q)$ from this knowledge, then S knows that $(p \& q)$.

The conjunction principle is false if knowing does not require probability 1. Suppose one’s epistemology allows for one to know that ϕ even if ϕ is, say .99 probable for one. It is then possible that one knows that p and one knows that q , but one fails to know that $(p \& q)$, for, unless p and q are mutually exclusive, $P(\phi \& \psi) = P(\phi) \times P(\psi)$, and, in this case, $P(p \& q) = .99 \times .99 = .98$. Therefore, if knowledge

¹²cf.Dretske [1971].

¹³According to Dretskeian ideology, 1 says that knowing that p entails having a *conclusive reason* for p . Here, one’s “conclusive reason” is evidence. See Dretske [1971].

¹⁴cf.Dretske [1981b].

requires a probability lower than 1, then the Conjunction Principle can be counterexemplified. I believe this is a reason to favor the claim that “knowledge” is an absolute concept.

One might have two different kinds of worry about the answer to “What is knowledge?” being sketched here. One might fear that this view is both too permissive and too restrictive.

The view of knowledge which emerges from those considerations is one in which knowledge is an evidential state wherein evidence eliminates all possibilities of error. In that sense there is a fundamental difference between knowledge and true belief of any kind - even true beliefs that are highly probable on one’s evidence: knowledge that p , but not true belief that p , is incompatible with it being possible that $\neg p$.

One may think the view is too permissive because, if knowledge entails the elimination of any possibility that the opposite of what one knows is true, then this view would permit one to close one’s mind to counterevidence, for any counterevidence to what one knows is evidence for something false (namely, the negation of what one knows). Simply put, this view of knowledge breeds *dogmatism*. On the other hand, one might think that the view been sketched here is too restrictive because, with the exception of *cogito*-like propositions such as “I am me”, our evidence never raises the probability of what we know to 1. But, if one knows only if the proposition one knows has probability 1 for one, then one knows virtually nothing. According to this line of reasoning, the view of the nature of knowledge I am sketching here breeds *skepticism*.

Since I deal with the charge that this view of knowledge entails dogmatism in chapter 3, I will consider in the remainder of this section only the charge that the view I favor entails skepticism.

Unger argued that skepticism *is* a consequence of knowledge being an absolute concept. This is where I part ways with him. Unger in fact embraced the view

that all sentences of the form “x is F” where “F” denotes an absolute concept are false. According to him, nothing is literally empty, flat, or known, because, if we inspect things carefully enough, we will always find something inside the place we thought was empty, a bump on the thing we thought was flat, and an incompatible possibility we thought our evidence had ruled out. Thus Unger defended skepticism about knowledge for the same general reason he defended skepticism about all absolute concepts: nothing in the world falls under absolute concepts.

I do think there are flat things and empty things in the world. I also think we know many things. But, if those concepts are absolute and the world does not seem to be absolute, how can that be? It is true that x being flat entails that x has no bumps. It is also the case that x being empty entails that there is nothing inside x. Likewise, it is true that x being known entails that there is no non-eliminated alternative to x. The mistake that the skeptic makes is to assume that *anything* can count as a bump in x, or as an object inside x or as a relevant alternative to x. This assumption is at least far from obvious. A lighthouse with dust is empty if we want to use it as a lighthouse, but not if we want to use it as a vacuum chamber.¹⁵ There being dust particles in the lighthouse is relevant for the correct application of “empty” to the lighthouse only if we are engaged in the former type of endeavor, not if we are engaged in the latter kind of endeavor. If we are engaged in the second type of endeavor, it becomes relevant to the correct application of “empty” to the lighthouse whether there is dust in it or not. Thus, x is empty iff there is nothing inside of x, but what counts as something inside of x is relativised to some parameter or standard. If some y does not satisfy this standard or parameter, then the question of where y is located is irrelevant to the question of whether x is empty or not. In the example I just gave, dust particles do not satisfy the standard: *things that might prevent us from using*

¹⁵The example is adapted from Dretske [1981a].

this lighthouse as a lighthouse. This makes dust particles irrelevant to whether the lighthouse is empty or not. Dust particles *do*, on the other hand, meet the following standard: *things that might prevent us from using this lighthouse as a vacuum chamber.* This makes dust particles relevant to whether this lighthouse is empty. In general, whether something is relevant or irrelevant to something else being empty is determined by contextually salient standards. In this sense, “empty” is a *relationally absolute* concept. Absolute, yes - but only in relation to a particular standard. The concept “flat” displays the same relational aspect “empty” does. The road in front of my house is flat if I am looking for a surface on which to ride my bike, but the same road is not flat if I am looking for a surface I can use to sign my check. So, although x is flat iff x has no bumps, what counts as a bump is relativised to a parameter or standard. The bumps on the road fail to satisfy the standard *irregularities that might prevent me from riding my bike* and are deemed irrelevant to the question of whether the road in front of me is flat. The bumps on the road do satisfy the standard *irregularities that might prevent me from signing a check properly* and are, therefore, relevant to the question of whether the road is flat. A skeptic about flatness and emptiness would say that “no bump” or “no object” is irrelevant to whether something is flat or empty. Given this maximally general standard, it would not be a surprise that there are no flat or empty things.

If “knowledge” is a an absolute term, then we should expect it to also display this relational feature “flat” and “empty” display. To see that “knowledge” is relationally absolute, consider a version of a case discussed in Dretske [1981b, p.368-9]:

The Amateur Bird Watcher Case

An amateur bird watcher spots a duck on his favorite Wisconsin pond. He quickly notes its familiar silhouette and markings and makes a mental note to tell his friends that he saw a Gadwall, a rather unusual bird in that part of the Midwest. The Gadwall has a distinctive set of markings (black

rump, white patch on the hind edge of the wing, etc.), markings that no other North American duck exhibits. What is more, these markings were all perfectly visible to the amateur.

In this case it seems that the amateur knows that the bird in front of him is a Gadwall duck. We may suppose that the set of markings he identified (that the bird looks thus-and-so) constitute his evidence for the claim that the bird is a Gadwall duck because it eliminates all alternatives incompatible with the bird being a Gadwall duck (for example, the alternative that it is a tree stump floating in the pond).

Now suppose I change the case a little and introduce an experienced ornithologist searching for traces of Siberian Grebes in the vicinity of the amateur bird watcher. Siberian Grebes are duck-like water birds, and, when in the water, they are practically indistinguishable from Gadwall ducks. One can only tell those birds apart if one sees them flying. Siberian Grebes have a reddish underbelly while Gadwall ducks have a white underbelly. Now suppose also that the ornithologist's working hypothesis is that Siberian Grebes have migrated to that region and that he and his assistant are looking for confirmation of that hypothesis.

As Dretske noted, at this point in the story, intuitions start to diverge and it stops being obvious that the amateur bird watcher knows the bird in front of him is a Gadwall duck. For one thing, people who think that he does not know might plausibly explain that inclination by an appeal to fact that the bird watcher's evidence does not rule out the alternative that the bird in front of him is a Siberian Grebe.

The important thing for us here is that this example shows that one knows that p iff one's evidence eliminates all alternatives to p , but what counts as an alternative to p may shift from one context to another. In the context without the ornithologist, the bird watcher's evidence eliminated all the alternatives to "That is a Gadwall duck", but failed to do so in the context with the ornithologist poking

around for Siberian Grebes. Consequently, “knowledge” seems to be, like “flat” and “empty”, a relationally absolute concept, for what counts as an alternative which has to be eliminated in order for one to know shifts.

How do we know what the relevant alternatives to p are? In other words, what is the standard a possibility has to meet in order to become a relevant alternative to what one knows? Here is a plausible initial suggestion:

(PRM) x is a relevant alternative to y in C iff, x is true in C and $P(y|x) < P(y)$.

Given that probabilities distribute over possibilities, PRM is equivalent to:

(POM) x is a relevant alternative to y in C iff x is possible with respect to C and $\diamond x \Rightarrow \diamond \neg y$.

PRM and POM are supposed to say the same thing in different words. Both “possibility” and “probability” refer to epistemic possibility and epistemic probability, respectively. Moreover, if x does not satisfy those conditions, then x is an irrelevant alternative. We can also use either PRM or POM to explain what happens in the bird watcher case. In the original version of the case, there are no Siberian Grebes and there was no other truth t in that context such that $P(\text{that is a Gadwall duck}|t) < P(\text{Gadwall duck})$. The upshot is that there is no relevant alternative the bird watcher’s evidence fails to rule out in the first version of the case. Alternatively, there is no x such that x is possible with respect to this version of the case and such that $\diamond x \Rightarrow \diamond$ (It is not the case that that is a Gadwall duck). In the second version of the case, on the other hand, if we suppose that the ornithologist’s hypothesis about the Siberian Grebes is true and that those birds are in fact migrating to the Midwest, then the truth “There are Siberian Grebes around” is such that, in this version of the case, $P(\text{that is a Gadwall duck}|\text{there are Siberian Grebes around}) < P(\text{that is a Gadwall duck})$ and $\diamond(\text{there are Siberian Grebes around}) \Rightarrow \diamond(\text{it is not the case that that is a Gadwall duck})$.

As stated, PRM and POM say that it is a necessary condition on x being a relevant alternative that x is true in C (possible with respect to C). What if I change the second version of the amateur bird watcher case and make the ornithologist's hypothesis false but also true that he and the whole Midwest *believe* that his hypothesis is true? In that case, is the truth "X believes that there are Siberian Grebes around" a relevant alternative to "That is a Gadwall duck?" The short answer to this question is "no." Intuitively, it is not that "X believes that there are Siberian Grebes around" is counterevidence to "That is a Gadwall duck," but, rather, that "There are Siberian Grebes around" is. Thus, the latter, if true, and not the former would be a relevant alternative to "That is a Gadwall duck", not the former. The problem is that the latter is, by assumption, false and thus an irrelevant alternative according to PRM and POM. So, the amateur bird watcher still knows that that is a Gadwall duck even when people believe there are Siberian Grebes around, because their belief is false. Here I am agreeing with Dretske on how to understand this case.

Now, this does not mean that, upon hearing about the ornithologist and his hypothesis and people's belief that there are Siberian Grebes around, it would be reasonable for the amateur bird watcher to stick to his guns and insist he had seen a Gadwall duck. It is an objective fact that, given the bird watcher's evidence, he knows there is a Gadwall duck in front of him, but this fact does not make it reasonable for him to disregard or discount counterevidence in favor of the opposite proposition unless he knows it to be true; i.e., unless he knows that he knows that it is a Gadwall duck. The view of knowledge I am sketching here can, therefore, deal with misleading evidence in a satisfactory way.¹⁶

More would have to be said in support of this view, but this will have to suffice for now. Yet, I think my goal was achieved: I showed that much can be said about what knowledge is even if we do not offer a definition of "knowledge."

¹⁶I elaborate extensively on this account of misleading evidence in chapter 3.

Knowledge is an evidential state whose value over and above true belief comes from the special support known propositions receive from one's evidence. If one knows that p , then one's evidence eliminates all relevant alternatives to p .

1.2 What Is To Come

A brief overview of the chapters follows.

Chapter 2 introduces KFK, sets the principle into historical context, and deals head-on with alleged cases of knowledge from non-knowledge. Chapter 3 puts forward an argument according to which $E=K$ entails KFK. The chapter then scrutinizes two different arguments against $E=K$ and finds them wanting. Chapter 4 applies KFK to the Gettier Problem and argues that a plausible version of the defeasibility theory of knowledge entails KFK. Chapter 5 argues that the most popular competing account of why Gettierized subjects fail to know ultimately fails.

Chapter 2

Knowledge from Knowledge

The main thesis of this work is that reasoning (or inference) is a source of knowledge only if one knows all the premises essentially involved in that reasoning. (I will refer to this thesis as “KFK.”) This is a normative thesis about reasoning, not a descriptive one (more on this distinction in the next section). What is more, the thesis is not only immediately appealing, but it is also endorsed by many major philosophers. In the remainder of this chapter I will briefly situate this thesis in the remote and recent history of philosophy and describe what contribution this work hopes to make to the debate about inferential knowledge.

The thesis that reasoning (or inference) is a source of knowledge only if one knows all the premises it involves essentially provides a basis for the intuitive notions that reasoning generates knowledge only if it is “sound,” and that one knows the conclusion of one’s reasoning only if one’s premises provide one with a “good reason” to accept one’s conclusion. It also explains why criticisms of the form “You should not believe that – you do not know the premises in your reasoning” seem appropriate, when true. If this thesis is true and one’s reasoning is a source of knowledge only if one knows all the premises it involves essentially, then our preference for sound reasoning over unsound reasoning is explained by our preference for known premises over premises we fail to know, in conjunction with the fact that we assign great epistemic value to knowledge. What is more, this claim about inferential knowledge has been widely accepted throughout the history of philosophy. According to Aristotle, one “demonstrates” that p is true only if one has a deductive argument whose premises are items of scientific knowledge

(“primitives”) and whose conclusion is p. Given this account of “demonstration,” Aristotle says,

... if we know and are convinced of something because of the primitives, then we know and are convinced of them better, since it is because of them that we know and are convinced of the posterior items.¹

According to René Descartes

many facts which are not self-evident are known with certainty, provided they are inferred from true and known principles through a continuous and uninterrupted movement of thought in which each individual proposition is clearly intuited.²

For John Locke, in turn, “... every step in Reasoning, that produces Knowledge, has intuitive certainty.”³

Like Aristotle, Immanuel Kant also saw inferential knowledge as being the business of syllogistic demonstrations. For Kant the major premise in a syllogism is an instance of “universal *a priori* knowledge” and is called a principle. From which it follows that every syllogism is a mode of deducing knowledge from a principle.⁴ But, since, for Kant, something is a principle just in case it is an item of *a priori* knowledge, the claim that every syllogism is a mode of deducing knowledge from a principle amounts to the claim that one’s inference produces knowledge only if one is certain (and, therefore, knows) that all of the premises in one’s argument are true.⁵

According to Bertrand Russell,

... But are we to say that nothing is knowledge except what is validly deduced from true premises? Obviously we cannot say this. Such a definition

¹Aristotle [1994, 72a25-30]. Jonathan Barnes’ commentary on this passage in Aristotle [1994, p.101-2] confirms our reading of the Aristotelian passage.

²Descartes [1985, p.15].

³Locke [1975, BK IV, ch.II, sec.7].

⁴Kant [1950, p.301].

⁵Of course, it is not enough that one knows the major premise in one’s syllogism; one also has to know the minor premise of one’s syllogism, but this is assumed by all the philosophers we are discussing.

is at once too wide and too narrow. In the first place, it is too wide, because it is not enough that our premises should be true, they must also be known.⁶

This is a partial (but impressive) list of philosophers who have endorsed our thesis. Strictly speaking, with the exception of Russell, they all endorse something stronger than our thesis; something like “one knows the conclusion of one’s reasoning only if one is certain of one’s premises and one’s premises entail one’s conclusion.” We need not join these philosophers in thinking that only deduction from premises that are certain for us yield inferential knowledge. This thesis is compatible with induction from premises one knows, but of which one is not certain, yielding inferential knowledge. For the philosophers I am discussing, on the other hand, inductive arguments only occasionally generate knowledge, most of the time they generate *probable opinion*. The fact remains that the stronger claim those philosophers are making directly and obviously entails our thesis and that they are, therefore, committed to that weaker claim.⁷

⁶Russell [1912, p.132-3]. This characterization is “too narrow,” according to Russell, in that it excludes the possibility of inferential knowledge via inductive inference.

⁷Things are a little more complicated than I am suggesting here. Aristotle, Descartes, Locke and Kant seem to be talking about what is required for one to acquire *scientia* of a conclusion, while Russell and I are talking about what is required for one to acquire *cognitio* of a conclusion. That is, after reading the passages I quoted above, one could claim that the view of inferential knowledge advanced by Aristotle, Descartes, Locke and Kant does not entail the view Russell and I defend, because the view those philosophers are discussing is a view about *inferential scientia* (i.e., they are talking about how to acquire *scientia* of a conclusion) and *scientia* is different in kind from *cognitio* - for one, the evidential standard the first has to meet is a lot more demanding (certainty) than the evidential standard the latter has to meet. I do not think this is quite right. Even if *scientia* and *cognitio* are different (complex or simple) mental states, plausible bridge principles connecting the two can be easily devised. Those bridge principles would show that *scientia* does in fact entail *cognitio*. The entailment is just more indirect than I am suggesting in the body of the text. For example, if one has *scientia* that the conjunction of all the members in the set $\{q, q_1 \dots q_n\}$ is true, then one also has *cognitio* that it is true, for being certain entails both belief and an epistemic support sufficient for *cognitio*. Moreover, if $\{q, q_1 \dots q_n\}$ entails p and one acquires *scientia* that p on the basis of one’s awareness of this entailment, then one also acquires *cognitio* that p , for, again, *scientia* entails belief and the epistemic support required for *scientia* exceeds the one required for *cognitio*. If this is right, then, even though a more complicated account of how Aristotle’s view entails mine is required, it is not something we cannot achieve. I am indebted here to Peter Klein for discussion.

Even though few other normative philosophical claims enjoy such an overwhelming support, the claim that reasoning is a source of knowledge only if one knows all the premises essentially involved in that reasoning has been at least tacitly rejected by most epistemologists thinking about inferential knowledge after 1963. That year saw the publication of Edmund Gettier's groundbreaking paper "Is Justified True Belief Knowledge?"⁸ The so-called "Gettier Problem" is a problem for theories of knowledge that incorporate the idea that one may acquire inferential knowledge on the basis of unknown premises. If, like Aristotle or Russell, one thought that inferential knowledge required knowledge of the premises involved, one would likely think Gettier's cases did not present any real danger to one's account of inferential knowledge, since the protagonist in those cases failed to satisfy this necessary condition on inferential knowledge and failed to know the conclusion of his argument. To the extent that the Gettier Problem emerges only for accounts of knowledge that at least implicitly deny that inferential knowledge requires knowledge of the premises of one's inference, Gettier did *not* refute the definition of "inferential knowledge" that Aristotle, Descartes, Locke, Kant and Russell had in mind.

I will discuss how to think about the Gettier Problem within the framework of an epistemology that accepts KFK in chapter 4. I will begin by presenting some intuitive support for that principle about inferential knowledge.

Consider the following case modified from Vogel [1990]. Suppose I ask you whether you will walk home after work and that this prompts you to reason thus:

1. My car is in the parking lot.
2. If my car is in the parking lot, then I will drive home.

Thus,

⁸Gettier [1963].

3. I will not have to walk home.

On the basis of this argument you tell me that you will not be walking home. Now, consider two different versions of the case I just described. In the first version circumstances are as described and you know that your car is in the parking lot. It seems natural to say that, in this scenario, you know you will not be walking home.⁹ Consider, however, a second version of this case. In this version I tell you that many cars have been stolen lately from the same parking lot in which you parked your car. In this version of the case it seems that you do not know that your car is in the parking lot, since your evidence is not strong enough to rule out the hypothesis that your car has been stolen. Not only that, in this second version, it seems that you do not know that you will not walk home and it seems natural to say that the reason why you do not know that is because you do not know that your car is in the parking lot. This is precisely the result we should have expected if we thought that inferential knowledge required knowledge of all premises essentially involved. Notice also that in both versions we may assume that it is not only true that your car is in the parking lot, but also that this is very likely to be true given the evidence you have. Still, this true and very probable belief is not knowledge in the second version of the case, and this interferes with you knowing that you will not walk home.

Here is another case that displays the same pattern. Suppose you bought a ticket in a large and fair lottery. The odds that any particular ticket will win are as abysmal as you wish (just make the lottery as large as you want). Suppose further that the drawing of the lottery took place at noon today and that it is now 1:00pm. Without yet having looked at the lottery result, you reason, at 1:15pm, in the following way:

⁹Of course, if my car is in the parking lot without gas or it has been smashed by a truck, I know my car is in the parking lot, but I do not know I will not have to walk home. I am, therefore, presupposing “my car” refers to my *working car*.

1. My ticket lost.
2. If my ticket lost, then I will not be able to afford that Bob Dylan Stratocaster they will be auctioning at Christie's next week, for I have no other source of money to buy it.

Therefore,

3. I will not be able to afford that Bob Dylan Stratocaster they will be auctioning at Christie's next week, for I have no other source of money to buy it.

As with the second version of the previous case, you do not know the conclusion of your argument even when we suppose that it is true that you lost the lottery, and you are basing your belief on your evidence (your knowledge of the odds), making your belief highly likely to be true. Moreover, since the natural explanation for why you fail to know that you will not be able to afford Dylan's Stratocaster is that you do not know that your ticket lost, this case confirms the idea that inferential knowledge requires knowledge of all the premises essentially involved.

Further support for the Aristotelian/Russellian take on inferential knowledge can be gained if we change the example so that you come to know your ticket lost before you infer that you will not be able to buy Dylan's Stratocaster. When we modify the case in that way, then you clearly know the conclusion of your argument. This is confirmation for the idea that inferential knowledge requires knowledge of all the premises involved, because, apart from whatever explains why the subject fails to know the premise in the first version of the case, that is the only difference between the two versions of the case.

We have to be careful, however, if we do not want to overstate the importance of this intuition. I have not yet said anything about whether knowledge of all the essential premises involved is *sufficient* for inferential knowledge of the

conclusion. All I have said so far is that knowledge of the premises is *necessary* for inferential knowledge of the conclusion. The question is, is KFK also a *sufficient* condition for inferential knowledge? No, it is not. If one's reasoning instantiates a deductively invalid or inductively weak argument, then one satisfies the consequent of KFK but not its antecedent. For this reason, unless otherwise noted, I presuppose, throughout, that KFK applies only to inferences that are either deductively valid or inductively strong.¹⁰ This is one reason why KFK is a necessary condition on inferential knowledge not a sufficient one.

What is more, KFK is a necessary condition for *inferential* knowledge not for *non-inferential* or **basic** knowledge. KFK is compatible with the existence of *non-inferential* knowledge (or perhaps knowledge gained through perception is an example of non-inferential knowledge). Of course, if one acquired one's premises in a non-inferential way, then, given KFK, one knows the conclusion of one's inference only if one knows those non-inferentially acquired premises; but the point is that, for all KFK says, one may know those premises non-inferentially, i.e. in a way that does not involve KFK (e.g., perceptually).¹¹

As we will see in chapter 3, cases such as these also put pressure on the main competing account of inferential knowledge that emerged after Gettier's paper was published. If this is right, then KFK has an important advantage over the current established view of inferential knowledge.

It is my goal to argue against the current philosophical establishment; I firmly believe that the tradition from Aristotle to Russell is basically right and that reasoning is a source of knowledge only if one knows all the premises involved essentially in that reasoning. Our case in support of this thesis will draw from

¹⁰Unless otherwise stated, I also presuppose that all arguments I discuss in this essay have a true conclusion.

¹¹Even though the main focus of this project is inferential knowledge, rather than non-inferential knowledge, we will come back to the issue of non-inferential knowledge below when we discuss a particular class of Gettier cases in chapter 3.

Timothy Williamson’s “knowledge-first epistemology.”

According to this approach to knowledge, “knowledge” is explanatorily prior to “belief.” That means that, according to this approach, the best explanation of many phenomena of philosophical interest (such as rational belief and action) is to be explained in terms of what one knows (for example, one acts rationally just in case one acts on what one knows), rather than in terms of what one believes. It is, therefore, in the spirit of this view that I explain the phenomenon of knowledge-yielding reasoning in terms of the knowledge we have of the premises essentially involved in that reasoning rather than in terms of the belief we have of those premises. Knowledge-first treatments of phenomena of philosophical interest are currently burgeoning¹² and this thesis has the goal of adding the treatment of epistemically proper theoretical reasoning to that literature.

In what follows, I will offer a new knowledge-first treatment of the Gettier Problem and alleged cases of knowledge from non-knowledge. The former issue is the main motivation for the establishment’s denial of our thesis, while the latter is a potential threat to a core assumption of knowledge-first views – the assumption that all and only knowledge is evidence (aka E=K). If our treatment of those issues is successful, we will have, as a result, an extended and improved version of the knowledge-first epistemology.

In addition to making a new case for the thesis that reasoning (or inference) is a source of knowledge only if one knows all the premises it involves essentially, and in addition to offering a new contribution to the nascent philosophical literature on knowledge-first epistemology (chapter 3), I will also develop new arguments against the current philosophical establishment’s unnecessarily narrow understanding of the Gettier Problem (chapter 4).

¹²Williamson [2000], Williamson [2011], Sutton [2007], and Hawthorne and Stanley [2008] are prime examples.

2.1 Clarifying KFK

As Aristotle famously put it, “all men by nature desire to know.”¹³ Since inferring something from what we already know is a way of coming to know, we also want to be able to use our knowledge to draw inferences from what we know and thereby come to know *some more*. For example, if I know that at least two people are needed to move a couch and I learn that my friend Peter will stop by later tonight, I can reasonably infer that I will be able to move the couch tonight, for I also know that Peter is a helpful and strong fellow. If I know that there is a good chance that there are predators around, and I see what seems to be freshly-made tiger footprints by the entrance of the cave, I better infer that I should not go into the cave. To generalize: inference is a major source of knowledge, a source beings like us could not live without in environments like ours.

Many of our actions rely explicitly or implicitly on reasoning for their rationality or justification. For instance, I *went* to the bank around the corner because I reasoned from my desire to cash a check and from my knowledge that, if I wanted to cash my check, I should go to the bank, to the conclusion that I should go to the bank. If I had not *inferred* that I should go to the bank, I would not have *gone* to the bank. Likewise, if I had not drawn that inference, I would not have *told* my boss I had to go to the bank before the end of my shift (the bank is closed then).

Epistemologists try to determine the conditions under which subjects grow their body of knowledge via inference. Our view, KFK, states that the following is a necessary condition subjects must satisfy in order to expand their body of knowledge via inference:

(KFK) S knows that p via inference or reasoning only if S knows all the premises essentially involved in her inference in support of p.

¹³Aristotle [1991, 980a22].

KFK says that S knows that p as a result of an inference only if S knows all the propositions essentially involved in S's inference, but what does this mean exactly?

I will first briefly discuss what I mean by “inference” and “reasoning.” I will then clarify the notion of “essential dependence.”

For the most part, I will use “inference” and “reasoning” interchangeably. If you think this is implausible (perhaps because you think that “reasoning” refers only to *conscious* inferences), then simply swap “reasoning” for “inference” or vice versa, as the context requires and so that my use matches your preferences. I take episodes of reasoning or inferences to be particular instances of a psychological process realized by brains and that this natural process is best characterized by cognitive science (perhaps by cognitive *neuroscience*). For our purposes here, particular episodes of reasoning will be understood as tokens of a process-type which takes belief state and/or knowledge state tokens as inputs and outputs another belief state or knowledge state token. However, since there are types of cognitive processes which are also properly described in this way but which are not commonly referred to as “reasoning” (e.g., daydreaming or free association), I will reserve the terms “reasoning” and “inference” to the tokens of a sub-class of that process-type. I will reserve these terms to the tokens of cognitive processes whose inputs are usefully said to be the “premises” of a subject's reasoning and whose output is usefully said to be the “conclusion” of that subject's reasoning. Hence, particular episodes of reasoning will be understood as tokens of a process-type best describe as having roughly the logical structure of an argument, with premises and a conclusion. In a word, reasoning will be understood as form of computation.¹⁴

This is not to say, of course, that when subjects reason they explicitly take the

¹⁴In a similar spirit, Keith Holyoak and Robert Morrison characterize “thinking” as “the systematic transformation of mental representations of knowledge to characterize actual or possible states of the world, often in service of goals” in Holyoak and Morrison [2005, p.2].

content of some knowledge-state they have to be the premise of their inference and the content of other knowledge-state they also have to be their conclusion. They sometimes do that, but not always. Rather than claiming that subjects do (or should do) that, I represent reasoning as having the logical structure of arguments for the same reason psychologists do that: because it is theoretically useful. Proceeding in that way gives us insight into our main topic - the *epistemic* status of inferences and the “products” of those inferences. Traditionally, philosophers have taken particular inferences and their products to have a certain epistemic status in virtue of conforming or failing to conform to a given set of rules (e.g., the rules of deductive or inductive logic).¹⁵ These rules, in turn, are taken to determine the epistemic status of inferences and the products of those inferences primarily because inferences that conform to those rules produce some intellectual good or other (e.g., true belief, justified belief, or knowledge) often enough. This essay continues this tradition in normative epistemology.¹⁶

As we will understand them here, inferences need not be explicit processes which are monitored and/or guided by one’s conscious self; inferences can very well happen implicitly, in a way that is not directly monitored and/or guided by one’s conscious self. This characterization of inference and reasoning is common in the psychological literature on dual process theory. For example, according to this model of reasoning, *System One* inferential processes are unconscious, fast, effortless, and automatic while *System Two* inferential processes are conscious, slow, effortful, and controlled.¹⁷ This is not to say that it is always clear

¹⁵cf. Goldman [1986, 3-5].

¹⁶Of course, this is a very rough description of the way we want to understand the epistemic status of inferences and their products. One thing should be clear, however, my project is a *normative* one, not a *descriptive* one. That is, I want to know what features inferences *should* have in order to produce knowledge (hence the appeal to the normative principle KFK and later in 3 the appeal to a knowledge norm of inference - KNI); I am not engaged in the project, commonly pursued by psychologists, of trying to discover what features inferences *in fact have* (regardless of whether those features produce knowledge or not).

¹⁷See *inter alia*, Evans and Frankish [2009] and Kahneman [2011]

whether some particular episode of inference is implicit or explicit or whether an inference has occurred at all. These issues should be resolved by whatever theory turns out to be the most theoretically fruitful and empirically adequate description of human reasoning. Also, because reasoning takes time, and the subject instantiating an inference might lose the knowledge she has of the premises before she reaches the conclusion,¹⁸ unless otherwise noted, I will assume that all protagonists in the cases I discuss maintain their knowledge of the premises at least until they have formed a belief in the conclusion.

The second notion in KFK in need of clarification is the notion of essential dependence. Not all reasoning involving a proposition one does not know constitutes cases in which the subject fails to know the conclusion of that reasoning - there are cases of *harmless ignorance*.¹⁹ Here is one such case:²⁰

Harmless Ignorance Case

Smith believes that (p) someone in his office owns a Ford on the basis of the following beliefs he also has: (q) Jones owns a Ford, (r) Jones works in my office, (s) Brown owns a Ford and (t) Brown works in my office. As it turns out s is false while q, r and t are true.

There is a falsehood involved in Smith's inference, but it is *harmless* - Smith clearly knows the conclusion of his inference. This prompts the obvious question: what is the difference between cases of harmless ignorance and cases of *harmful* ignorance? The difference, of course, is that something one does not know is essentially involved in the latter, but not in the former kind of case. But how can we know when a proposition we fail to know is essentially involved in one's reasoning or not? Here is a plausible suggestion.

¹⁸Maybe because she forgets one the premises or because she learns some piece of misleading evidence against one or all the premises; or because she learns some piece of misleading evidence suggesting that the argument is inductively weak or deductively invalid.

¹⁹The label "harmless falsehoods" was introduced by Peter Klein in Klein [2008] to describe cases of harmless ignorance in which the ignorance was caused by the falsehood of the proposition believed.

²⁰Adapted from Lehrer [1965].

(Dependence) The conclusion that p depends essentially on some premise q for S in a case C iff (i) S would not have believed that p had she not believed that q in C ; and (ii) if q were not in S 's evidence set in C , p would not have been justified for/known by S in C .²¹

I will take the truth of (i) and (ii) to exclusively characterize what I mean by “ x depends essentially on y .” There might be a sense in which x depends essentially on y if either (i) or (ii) - but not both (i) and (ii) - are satisfied, but that is not the sense of “ x depends essentially on y ” that I am interested in here. Furthermore, (i) captures the sense in which one’s conclusion depends *causally* on one’s premise(s) in a crucial way;²² (ii) captures the sense in which one’s conclusion depends *evidentially* on one’s premise in a crucial way. I will thus refer to (i) and (ii) as the “causal counterfactual” and the “evidential counterfactual”, respectively. This is the sense in which the falsehood in cases of harmless ignorance cases is inessential: the relevant causal and evidential counterfactuals are false in those cases.

Some caution is in order here. One could plausibly object that in Harmless Ignorance, Smith’s true belief that someone in his office owns a Ford is causally overdetermined by his false and true beliefs about who owns a Ford. If we apply our causal counterfactual to the case, it would seem as if neither of those beliefs were the cause of Smith’s true belief, for, had either one of them not been formed, the other would have been sufficient to bring about Smith’s belief in the conclusion. What is more, once we make this problem salient, one might start to worry about whether the causal counterfactual in Dependence will infect KFK with the

²¹This account of dependence is similar to the one proposed by Klein [2008].

²²This kind of counterfactual dependence holds even if causation cannot be reduced to counterfactuals as in Lewis [1979]; and even if counterfactuals cannot be reduced to causation as in Jackson [1977]. The idea that beliefs are causes (in particular, of other beliefs) is widely accepted. But for some dissenting arguments see Klein [2012] and McLaughlin [2006]. For a recent synoptic discussion of the relationship between causation and counterfactual conditionals, see Paul and Hall [2013] and the references therein.

thorny issue of deviant causation and of how to analyze “x is the cause of y” given that issue. So, it may be argued that it is only a matter of time until someone presents cases exploiting other types of deviant causal chains (e.g., cases of early and late preemption) in order to falsify Dependence.

These worries emerge from a misunderstanding of the role Dependence plays in my story in support of KFK. Dependence is supposed to provide a heuristic, or test we may apply in order to determine if a premise is causally and/or evidentially crucial to a conclusion (in a particular case). By putting forward Dependence I am *not* trying to offer necessary *or* sufficient conditions for belief causation (or for causation in general, for that matter), or for evidential support. KFK says that inferential knowledge is incompatible with the essential dependence of one’s conclusion on something one fails to know. KFK is silent about whether or not inferential knowledge is always accompanied by something that is essential to one’s knowledge of the conclusion. That is why the KFK defender is not concerned with whether or not the true or the false belief in Harmless Ignorance is essential to the subjects’ conclusion. KFK only says that essential dependence on a premise one does not know *precludes* the possibility that the conclusion of that argument is an item of inferential knowledge. In that sense, all we need to do in order to show that some premise of which S is ignorant of does not harm S’s epistemic position vis-a-vis some conclusion is to apply Dependence. If either one of the counterfactuals is false, then the relevant proposition is inessential in the given circumstances. From this result nothing follows about *something else* being essential to S’s knowledge or not. It should, therefore, be clear that Dependence is not really infected with the deviant causation issues that plague analyses of causation. It is neither here nor there the fact that in Harmless Ignorance Dependence shows that neither the true nor the false beliefs are essential to the subject’s conclusion – I applied the test to gauge whether a belief was essential or not to his conclusion; I did not apply that test in order to uncover what was

the *true cause* of his belief in the conclusion. A similar point can be made about other kinds of cases exploiting other kinds of deviant causation.

So understood, KFK amounts to the claim that one knows that p as the result of an episode of inference only if one knows all propositions that are essential to S's inference in support of p.

Before I conclude I will address an issue which suggests that KFK is not even *necessary* for inferential knowledge.

Recently, Ted Warfield,²³ Peter Klein,²⁴ Branden Fitelson,²⁵ and Claudio de Almeida²⁶ have presented alleged cases of *knowledge from non-knowledge*, cases that allegedly show that one can gain knowledge even though one's reasoning relies essentially on at least one premise one does not know. Those cases put pressure on KFK as a universal principle on epistemically proper reasoning. I will not be able to "solve" this problem here, for, presently, I do not have the space or, honestly, a fully developed theory that resolves the issue. The hope is that what I can say here will serve as a basis for a more developed response to those cases in the future.

I will offer three different types of replies to alleged cases of knowledge from non-knowledge: a non-concessive reply, a partially concessive reply and a fully concessive reply. The thought behind each reply is that KFK maintains its status at the center of epistemological theorizing about reasoning and inference, even if I give a fully concessive reply to alleged cases of knowledge from non-knowledge.

But, before we look at these cases, I should report that I do not have the intuition that the protagonists in *all* cases discussed by Warfield, Klein, Fitelson and de Almeida know the conclusion of their reasoning. *Some* of the cases elicit

²³Warfield [2005].

²⁴Klein [2008].

²⁵Fitelson [2010].

²⁶See de Almeida [2003] and de Almeida [2004]. de Almeida [2003] was given at the 2003 Meeting of the Central APA.

that intuition in me; some *do not*.²⁷ Of course, I might be an outlier when it comes to having those intuitions.

But, even if they are robustly shared by almost everyone reading those cases, that would not necessarily mean that the protagonists in those cases do in fact know the conclusion of their inferences. That most or all intuit that p does not entail that p (just as “most or all believe that God exists” does not entail “God exists”).

Consider the following alleged case of knowledge from falsehood. This is a case discussed by Warfield.²⁸

The Handout Case

Warfield has 100 handouts for his talk; he carefully counts 53 people in attendance and concludes that his 100 handouts are enough. As it turns out he double-counted one person who changed seats during the counting of heads and there are in fact 52 people attending his talk.

This case poses a problem for KFK because Warfield seems to know that 100 handouts are enough even though this knowledge seems to depend essentially on something he fails to know. I will offer three different replies to this case and suggest that each one of those replies could be developed into a response not only to the Handout Case, but to any alleged case of knowledge from non-knowledge.

Let us look first at a non-concessive reaction to cases such as Warfield’s. This type of reply is non-concessive, because it denies that the initial impression that

²⁷After presenting his cases, Warfield admits that he cannot say anything in support of the claim that the protagonists in those cases know. The only thing he can do, Warfield says, is to “appeal to clear and widely shared intuitions about the cases” (Warfield [2005, p.408]). There are at least two problems with this claim. First, if Warfield’s cases had not been published before, how in the world could they have elicited a “widely shared intuition” that the protagonists in them know? Second, one does not have to deny that intuitions constitute evidence that confirms/disconfirms philosophical theories in order to think that the confirmation/disconfirmation it provides might be swamped by further evidence. So, even if Warfield’s intuition about knowledge from falsehood is “widely” shared, this would not prevent it from being outweighed by further considerations.

²⁸Warfield [2005, p.408].

Warfield knows that 100 handouts are enough is probative of him actually knowing that. The idea is simple enough: Warfield does not satisfy KFK and therefore does not know that 100 handouts are enough; however, this is compatible with Warfield's belief being highly justified. Intuitively, Warfield is *excused* from not satisfying KFK because he counted the number of people carefully and because he did not miss the right number by much. But, having an *excuse* to believe truly (or even for being justified in believing truly, if having an excuse to believe is sufficient for being justified in believing) is not the same as knowing. So, even though Warfield's belief is highly justified and, for all practical purposes (such as providing everyone at his talk with a handout) his belief is "as good as knowledge," he does not know that 100 handouts are enough.

I will develop this notion of excuse further in chapter 3, but here is the type of argument I have in mind here. In all alleged cases of knowledge from non-knowledge, the protagonist fails to know the conclusion of his reasoning, but it is highly probable, on the protagonist's evidence, that he knows the false premise of his argument. And, since "I know that p" is highly probable on the protagonist's evidence even though p itself is false, it is *reasonable* for him to believe that p and he has an *excuse* to believe that p. However, it being reasonable for one to believe one's premises is necessary, but obviously not sufficient to generate knowledge of one's conclusion. Thus, the protagonists of cases of alleged knowledge from non-knowledge have inferentially *reasonable* (or justified) belief, rather than knowledge of the conclusion of the argument.²⁹

So, even though it is false that Warfield knows the conclusion of his reasoning, we can accommodate the related intuition *that it is reasonable for Warfield to believe as he does*. Moreover, we are liable to mistake this fact for the non-fact that the protagonist in a case of alleged knowledge from non-knowledge knows,

²⁹This account of what it is to have an excuse is modeled on Williamson's treatment of the knowledge norm of assertion. See Williamson [2000, ch.11].

because (i) it is highly reasonable for him to believe he knows the premises of his reasoning; (ii) from the protagonist's point of view, there is virtually no difference between knowing that p and not knowing that p , since it is highly likely on his evidence that he knows that p ; and, (iii) when we assess philosophical cases we tend to put ourselves into the protagonist's shoes and let this influence our assessment.³⁰

One might object that this non-concessive treatment of alleged cases of knowledge from non-knowledge blurs the distinction between those cases and Gettier cases. In both cases the protagonist's conclusion depends essentially on a falsehood which prevents her from knowing the reasoning's conclusion. The problem with this objection is that it forgets that we can characterize Gettier cases by making essential involvement of a falsehood a *necessary condition* rather than a necessary *and* sufficient condition (or even merely sufficient condition) for a case to be a Gettier case.³¹ For all I have said, other conditions besides dependence on a falsehood are required for a case to be a Gettier case. An obvious candidate for such a further condition is that a case is a Gettier case only if it is fairly obvious that the protagonist in that case fails to know. If that is correct, then alleged cases of knowledge from non-knowledge cannot be Gettier cases, for, as I suggested above, it is not "obvious" that there is no knowledge in those cases.

I can also provide a partially concessive reply to the problem of knowledge from non-knowledge. First, note that, although non-concessive, the first reply did not dispute the idea that Warfield's true belief depended essentially on a false one.

³⁰In chapter 3 I will argue that we should accept a *knowledge norm of inference*, KNI (one must: infer only from what one knows). We can then capture the force of (i)-(iii), if we appeal to KNI. Since Warfield's reasoning relies essentially on a falsehood, he has violated KNI and should not have reasoned as he did. However, since it was reasonable for him to believe that he did satisfy KNI (i.e., since it was probable on his evidence that he knew the premise of his reasoning), it is reasonable for him to believe the conclusion of his argument. The upshot is that Warfield and all the other protagonists in alleged cases of knowledge from non-knowledge fail to know (but reasonably believe) the conclusion of their inference.

³¹I will argue for just such an account of Gettier cases in chapter 4.

The partially concessive strategy disputes this very point. The idea is that we can concede that there is knowledge in cases of knowledge from non-knowledge, but insist that, despite appearances to the contrary, this knowledge does not depend essentially on anything the subject does not know. The thought is that alleged cases of knowledge from non-knowledge are disguised cases of harmless ignorance. In Warfield's case this strategy would involve showing that

(t) There are enough handouts.

does not essentially depend on the false

(f) There are 53 people in attendance.

Remember that x depends essentially on y only if Dependence is satisfied. We can now show that it is not the case that t depends on f in both the causal and evidential way required for essential dependence. In fact, I will argue that t does not depend on f in either of those ways. Consider, first, whether t *causally* depends on f or not.

(cc) If Warfield had not believed that there are 53 people at his talk, he would not have believed that his 100 handout copies were enough.

The worlds closest to the actual world are worlds in which Warfield's belief about the number of people in attendance, n , $n < 100$, and Warfield still forms a belief about the number of attendees. Thus, those are also worlds in which Warfield believes that 100 handouts are enough. Warfield's case seems to fail the counterfactual test for evidential dependence as well.

(ec) If Warfield's evidence set did not include "There are 53 people in attendance," he would not have been justified in believing "100 handout copies are enough."

The worlds closest to the actual world are worlds in which Warfield has views about the number of people in attendance, n , and ones in which his belief about n takes the value of n to be smaller than 100. Thus, in the worlds closest to the actual world Warfield's evidence set justifies him in believing that 100 handouts are enough. In order to see that, note that we need two things to evaluate counterfactuals like ce properly: we need to make the smallest number of changes possible to the actual world so that the antecedent of the counterfactual is true; and, second, we need to consider how the world so changed would evolve and whether this evolution would make the consequent true or not. Since the actual world is one in which not only f is in Warfield's evidence set, but also one in which Warfield *counts the number of people in attendance*, a world W is a member of the set of closest possible worlds (with respect to Warfield's actual world) only if, in W , f is not in Warfield's evidence set and Warfield counts the number of people in attendance. Such a W is closer to the actual world than a world in which we change both things. But, since counting the number of people in attendance will lead Warfield to form a belief about the number of people in attendance, n , it is plausible that in the closest possible worlds some proposition about n but different from f will be in his evidence set and will be enough to justify him in believing that 100 handouts are enough.

In fact, even if we do not accept the result we arrive at via dependence, there might be a case to be made that, *in the actual world*, f is inessential. Here is why one might think that. First, note that f , by itself, is hardly a good enough reason for Warfield to believe that t . That is, Warfield's argument for t is obviously enthymematic, since it does not even explicitly mention his background (but crucial) knowledge that, for example, he brought 100 handouts to the talk. If Warfield in fact knows that 100 handouts are enough, then this knowledge cannot literally depend exclusively on the proposition that there are 53 people in attendance, for this is not a good enough reason (in and of itself) for anyone

in Warfield’s situation to believe that 100 handouts are enough. In any plausible way we might fill out the details of Warfield’s case, he will have arrived at his talk with a lot of background knowledge relevant to the question of whether he knows that 100 handouts are enough or not. For instance, in any plausible interpretation of the case he arrived at his talk knowing something like the following:

(**t***) If there are between 0 and 100 people in my talk, then 100 handout copies are sufficient.

If t^* were not part of Warfield’s background knowledge (say, because he did not even believe it), it would be hard to explain why he would think t followed from f .³² The upshot is that, whether we apply Dependence or not, the fact is that the false “there are 53 people at my talk” is inessential to Warfield’s inference, and the case is, when scrutinized, not a case of knowledge from non-knowledge, but a case of harmless ignorance. Both of these strategies can be extended to cover other cases of alleged knowledge from non-knowledge.

Finally, suppose that both the non-concessive and the partially concessive strategies fail and that we have to settle for a *fully concessive* reply to alleged cases of knowledge from non-knowledge. That is, what if we had to admit that satisfying KFK is not always necessary for inferential knowledge? How detrimental would that be for my project? Not terribly, I think. First of all, KFK is the most plausible explanation of what is wrong in lottery-like inferences we saw above. Moreover, as we will see in chapter 3, appealing to anything weaker than knowledge in one’s explanation of what is wrong in those inferences will not work. Second, I will argue in chapter 4 that KFK can be used to solve the Gettier Problem and that this solution is, all things considered, more plausible than the most popular alleged solution in the literature today. Third, KFK is

³²One might complain that t^* doesn’t partially justify t on the ground that the former does not really cause the latter. The problem with this objection is that it overlooks the fact that relevant background beliefs play an epistemic role even if not an explicitly causal one.

entailed by a reputable version of the defeasibility theory of knowledge (chapter 4) and by Williamson's influential account of evidence, E=K (chapter 3). So, rejecting KFK commits one to rejecting *both* the defeasibility theory of knowledge and E=K. If I am right about all this, then, no, KFK not being "exceptionless" is not detrimental at all, for the theoretical fruits this principle bears would far exceed the cost inflicted by cases of knowledge from non-knowledge.³³

³³There is still another possible reply to alleged cases of knowledge from non-knowledge. Some (e.g., Achinstein [2003]), have argued that there are importantly different and equally theoretically fruitful concepts of "evidence". Perhaps proponents of cases of knowledge from non-knowledge have in mind a concept of "evidence" that allows for propositions we fail to know to count as evidence while I and others use a concept of "evidence" that excludes that possibility. Following Peter Achinstein, we might dub the knowledge from non-knowledge theorist's conception of "evidence" as "subjective evidence" and the concept employed by KFK and throughout this essay "veritcal evidence". While veritcal evidence requires that evidential proposition and hypothesis are both known (if believed), subjective evidence requires instead that (i) S believes that e is evidence that h; (ii) S believes that h is true or probable; and, (iii) that S's reason for believing that h is true or probable is that e is true. (Note that Warfield satisfies conditions (i)-(iii) for subjective evidence.) Of course, this reply raises the question of whether some concept of evidence is conceptually prior to the others and whether or not they all deserve a place at the epistemological table. I will not pursue those important questions here. Not surprisingly, however, I think that veritcal evidence (or at least what Achinstein calls "potential evidence") is more fundamental than all the other concepts of evidence, but I will not argue for that here. I refer the reader to Peter Achinstein's excellent discussion of this and other issues related to the concept of "evidence" in Achinstein [2003, p.13-113] (See also Kelly [2006]). I will come back to the issue of what evidence is in chapter 3 when we discuss E=K.

Chapter 3

KFK, E=K, and The Knowledge Norm of Inference

Knowledge-first epistemology is widely known for its account of evidence (i.e., E=K) and for its knowledge-centric account of assertion, action, and belief. In this chapter I will explore those two themes in connection to KFK. First, I will show that E=K entails KFK, and then I will model a norm of inference on the knowledge norm of assertion and action.

3.1 From E=K to KFK

Williamson offers a battery of arguments in support of the thesis that one's evidence is the totality of one's knowledge. I will not revisit those arguments here. Rather, I would like to point out that one is committed to KFK, if one accepts E=K, for the latter entails the former.

Suppose, for a future *reductio ad absurdum*, that both all and only knowledge is evidence and KFK is false. Now, if KFK is false, then there is an S such that S knows that x depends essentially on y, y is evidence for x for S, but S does not know that y. But, if E=K is true, then there is no such S, for, according to E=K,¹ y is evidence for x for S, only if S knows y. This is a contradiction, and we may conclude that our hypothesis is false. This argument can be stated more precisely thus:

¹Williamson accepts the following principle about the “evidence for” relation: (EV) e is evidence for h for S if and only if S's evidence includes e and $P(h | e) > P(h)$. Given E=K, e is not evidence for h unless e is an item of knowledge for S. See Williamson [2000, p.187].

1. All and only knowledge is evidence, but for some S and case C, S knows that p inferentially in C even though p depends on e for S in C and S does not know e. [assume for *reductio ad absurdum*]
2. But, if p depends on e for S in C, then e is evidence for p for S in C. [(EV)]
3. e is evidence that p for S in C and S does not know that e in C. [1,2]
4. But, according to E=K, no x is in S's evidence set unless S knows x. [E=K]
5. e is evidence that p for S in C and e is not evidence for p in C [3,4]

It follows that E=K entails the view that reasoning yields knowledge only if all the premises involved essentially in one's reasoning are known. Step 2 is the only one in need of independent support.

1. p depends on e for S in C. [assume for a conditional proof]
2. If p depends on e for S in C, then S would not have been justified in believing (/would not have known) that p if S's total evidence did not include e in C. [from 1 and Dependence]
3. S would not have been justified in believing (/would not have known) that p if S's total evidence did not include e in C. [from 1,2]
4. If S would not have been justified in believing (known) that p had her evidence not included e in C, then S would not have been justified in believing (/would not have known) that p in part because of e. [Dependence]
5. S is justified in believing (knows) that p in part because of e. [from 3,4]
6. If S is justified in believing (knows) that x in part because of y, then $P(x | y) > P(x)$. [evidence is probability raising]
7. $P(p | e) > P(p)$. [from 5,6]

8. S's evidence includes e in C. [from 1, 2]
9. y is evidence that x for S in C iff S's total evidence includes y and $P(x | y) > P(x)$ for S in C. [(EV)]²
10. e is evidence that p for S in C. [from 7,8,9]
11. If p depends on e for S in C, then e is evidence that p for S in C. [1,2-10 by conditional proof]

This result shows that KFK is not only the view of inferential knowledge adopted by most philosophers throughout the history of philosophy, but that it is also a consequence of an influential view about what evidence is.

I will now look at two different objections to E=K. If those objections succeed, then we could not appeal to E=K as a reason for KFK. The first one says that, if E=K is true, then there is no Gettier Problem. The second problem suggests that this view of what evidence is breeds dogmatism if, as I am doing in this essay, I couple this idea with the idea that knowledge has probability 1 on one's evidence.

3.2 Objection: No Gettier Problem

A direct implication of E=K seems to be that false beliefs cannot justify other beliefs, for no false belief can be part of one's total evidence and one's total evidence is what inferentially justifies belief. The problem with this alleged implication of E=K, as Juan Comesaña and Holly Kantin³ have noted, is that it contradicts a claim Gettier cases rely on. The original Gettier cases (Case I, a.k.a "The Coin Case", and Case II a.k.a "The Ford Case") explicitly relied on two principles. One principle claims that justification is closed under known entailment. The other principle Gettier asked us to assume, the one that is seemingly incompatible with

²Williamson [2000, p.187].

³Comesaña and Kantin [2010].

$E=K$, states that sometimes one is justified in believing a falsehood. Call this principle “The Fallibility Principle” or FP, for short. In what follows, I will argue that $E=K$, contrary to what Comesaña and Kantin would want us to believe, is compatible with the agent being justified in believing a falsehood - at least in one sense of “justified” that is epistemologically important.

If $E=K$ is true, then not only is FP seemingly false, but the protagonist in the original Gettier cases (and other cases with the same general structure) do not have a justified true belief, because this belief, by the very nature of the cases, depends essentially on a falsehood for its alleged justification, and false beliefs are not part of one’s total evidence. Since the intuition that the agent in Gettier cases is justified is accepted by virtually everyone,⁴ the friend of $E=K$ is well advised to either drop the view (for it entails that the agent is not justified) or find a way to accommodate the widespread intuition about those cases. I will propose a way in which the friend of $E=K$ can accommodate this widespread intuition about Gettier cases. But first let me make the challenge to $E=K$ more precise by offering an argument in its support:⁵

Against $E=K$

1. All and only knowledge is evidence and only knowledge inferentially justifies belief. [Assume for *reductio ad absurdum*]
2. Evidence inferentially justifies belief. [from 1]
3. False beliefs are not part of one’s evidence. [from 1]
4. The protagonist in a Gettier case has a justified true belief whose justification depends essentially on a justified false belief. [Assumption]

⁴Including people such as Weatherson [2003] and Hetherington [2011] who think agents in some Gettier cases are not only justified, but also know.

⁵This is my reconstruction of Comesaña’s and Kantin’s main argument against Williamson.

5. False beliefs are sometimes part of one's evidence and do sometimes justify via inference other beliefs. [from 2 and 4]
6. Only knowledge justifies belief inferentially, and it is not the case that only knowledge inferentially justifies belief. [from 1 and 5]
7. Either it is false that all and only knowledge is evidence, or it is false that only knowledge inferentially justifies belief. [from 1, 2-5 by *reductio ad absurdum*]

I will take issue with step 4 in this argument. Since the idea that Gettiered agents rely essentially on falsehood for their justification is virtually a dogma in contemporary epistemology, Comesaña and Kantin are in good company when they appeal to this idea in their argument against E=K. However, as I will argue below, justification is a notoriously ambiguous notion and it is far from clear what notion of justification is assumed in Gettier cases. But I am getting ahead of myself. Before I turn to my preferred response to this argument, let me briefly discuss a response I find unconvincing.

Consider Gettier's Coins case.⁶ Smith comes to believe truly (e) "The person who will get the job has ten coins in his pocket" via deduction from his false belief in (d) "Jones is the man who will get the job and Jones has ten coins in his pocket." One might want to say that it is not the false proposition d that is doing the justifying of e for Smith, but the nearby truth Smith also believes (d*) "The president of the company assured me that Jones will get the job and Jones has ten coins in his pocket." Since it is plausible to think that Smith knows d*, the E=K proponent could say that d* is part of Smith's evidence and that it justifies his true belief in e (after all, the president's say-so is the only reason Smith has to accept the first conjunct of d). One could say something similar about Gettier's Ford case: even though Gettier says that Smith is justified in

⁶cf. Gettier [1963].

believing the true (g) “Either Jones owns a Ford or Brown is in Barcelona” via deduction from the false (f) “Jones owns a Ford,” one might want to say that it is Smith’s true belief in (f*) “I believe Jones owns a Ford” that is doing the justifying instead. Here too, Smith may plausibly be said to know f* and, hence, that it is also part of his total evidence (again, according to Gettier, Smith’s sole reason for believing “Jones owns a Ford” is that Smith “has strong evidence” for it. If this is right, then he has even stronger evidence for f*). One could think that E=K has been vindicated by this reply to Comesaña’s and Kantin’s argument and that the intuition that Gettiered agents arrive at a justified true belief via inference was, as a result, accommodated. I think this is not the correct argument for the friend of E=K to posit in response to this argument.

Comesaña and Kantin consider what is roughly the same strategy and also find it wanting. They think that in order for this response to work it is

not enough to find some propositions that you know and that justify you, it is necessary to argue that every proposition that justifies you is something that you know. And there is no argument that we can think of to the effect that your belief that Jones got the job plays no part whatsoever in justifying you in thinking that whoever got the job has ten coins in his pocket.⁷

The problem with Comesaña’s and Kantin’s claim that they cannot think of any argument in favor of the claim that Smith’s false belief does not do any justifying is that it simply begs the question against Williamson who has offered positive arguments for the claim that no false belief is part of one’s evidence. We can do better than that. I think we can reject this defense of E=K without begging any questions.⁸

⁷Comesaña and Kantin [2010, p.499-500].

⁸Weatherson [2011] has made a different point about Comesaña’s and Kantin’s claim that they “can’t think of any argument” in support of the claim that false beliefs are not part of one’s evidence. According to Weatherson, their claim suffers from a failure of imagination, for one can always think of an argument for p, namely God knows that p, therefore p. Comesaña and Kantin could say, I think, that this is an uncharitable reading of their claim and that they meant to say that they can’t think of any *good* argument in favor of the claim that no false belief is part of one’s evidence.

The real problem with this reply is twofold. First of all, this reply presupposes a controversial account of the basing relation. Plausibly, one is doxastically justified in believing that p on the basis of some ground g only if g at least partially causes one's belief that p . The reply we are considering assumes, however, that (d^*) "The president of the company assured me that Jones will get the job and Jones has ten coins in his pocket", justifies Smith's belief in (e) "The person who will get the job has ten coins in his pocket" even though d^* does not even partially cause Smith's belief in e . However, doxastic justification is the kind of justification required for knowledge.⁹ Hence, d^* does not justify Smith's belief in e in the way epistemologists always assumed d justified e in the original Gettier cases. Secondly, if we think that Gettier cases are such that they all instantiate the Closure Principle and the Fallibility Principle, then the reply to Comesaña and Kantin sketched above cannot help the E=K proponent fend off the charge that E=K entails the non-existence of Gettier cases, for neither d^* or f^* entail the conclusion of Smith's argument. So, if we think that Gettier cases necessarily instantiate those two principles, then the reply above, if correct, will ensure that Gettier's own cases are not "Gettier cases." As I will argue in chapter 4, we have reason to believe that Gettier cases necessarily instantiate those principles.¹⁰

Here is a better reply to Comesaña's and Kantin's argument, one that I think

⁹To see that, suppose S 's evidence set included r , r_1 , and r_2 . Suppose further that only r is in fact evidence for p and that, in spite of this, S 's belief that p is caused only by r_1 . In this scenario, S fails to know that p , even though he would have known that p if her belief that p had been at least partially caused by r instead of r_1 . The idea is that *having* the right reason is not sufficient for knowledge, knowledge requires that one "use" it in support of the target truth. Thus, one's belief that p is *doxastically justified* by a reason r only if r is partially causally responsible for one's belief that p . One is *propositionally justified* by a reason r in believing that p even if r is not partially causally responsible for one's belief that p . Doxastic justification entails propositional justification, but the converse is not true. cf. Korcz [1997] and Korcz [2010].

¹⁰Shope [1983, p.4] suggests a few plausible necessary conditions on a case C being a "Gettier case": C is a Gettier case only if S has a justified true belief that p in C , S does not know that p in C and there is some false proposition, q , S is either justified in believing is true or at least S would be justified in believing that q in C . For the purposes of my discussion here, I will accept Shope's partial characterization of a Gettier case. I will offer a slightly different characterization in the next chapter when we discuss the Gettier Problem itself.

works. As far as I know, no one has suggested it, or even considered it, even though one can easily extract it from what Williamson says about proper assertion.¹¹ According to Williamson, knowledge is the constitutive norm of assertion and this entails that one asserts that p properly only if one knows that p . When presented with an example of an assertion that seems proper but whose proposition is not known, we may appeal to a distinction between having a warrant to assert and it being reasonable for one to assert.¹² In what follows I will argue that the friend of $E=K$ should extend this distinction to cases, like the Gettier cases, in which one seems to believe properly (i.e., justifiably) propositions one does not know. This will allow her to reply to *Against E=K* in a principled way. Or so I think.

Let us consider, first how we should think about the “having a warrant/it being reasonable to believe” distinction. To that effect, consider the following case discussed by Williamson:

The Fake Snow Case

It is winter, and it looks exactly as it would if there were snow outside, but in fact that white stuff is not snow but foam put there by a film crew of whose existence I have no idea. I do not know that there is snow outside, because there is no snow outside, but it is quite reasonable for me to believe not just that there is snow outside, but that I know that there is; for me, it is to all appearances a banal case of perceptual knowledge. Surely it is then reasonable for me to assert that there is snow outside.

On the basis of this case, Williamson [2000, p.257] distinguishes between *having a warrant to assert* and *it being reasonable for one to assert* and explains how the case is consistent with the knowledge norm of assertion according to which one appropriately asserts that p if and only if one knows that p :

¹¹Williamson [2000, ch.11]

¹²DeRose [2002, p.180]; DeRose [2009, p.94-5] also make this distinction but call it “primary” and “secondary propriety”. Weiner [2006] and Benton [2011] also endorse the distinction. But see Lackey [2007] for some doubts about the distinction.

The case is quite consistent with the knowledge account [of assertion]. Indeed, if I am entitled to assume that knowledge warrants assertion, then, since it is reasonable for me to believe that I know that there is snow outside, it is reasonable for me to believe that I have warrant to assert that there is snow outside. If it is reasonable for me to believe that I have warrant to assert that there is snow outside, then other things being equal, it is reasonable for me to assert that there is snow outside. Thus, the knowledge account can explain the reasonableness of the assertion. However, granted that it is reasonable for me to believe that I have warrant to assert p , it does not follow that I do have warrant to assert p .[.]

The distinction between having a warrant to assert and it being reasonable for one to assert allows us to say that it is reasonable for Williamson to assert that there is snow outside even though he did not have a warrant to assert it. The intuitive idea is that, even though Williamson does not know that there is snow outside, he is *blameless* (has an excuse) for asserting that there is, for the proposition that he knows that there is snow outside is very probable on his evidence.¹³ We can express the idea of it being reasonable to believe that x is the case more precisely:

$$(R) \text{ Rs}\phi \leftrightarrow P(Ks\phi \mid e) > \Omega$$

According to R, it is reasonable for S to believe that ϕ if and only if the probability that S knows that ϕ conditional on her total evidence e is higher than some threshold Ω .

It should be noted that the notion of having a warrant to assert that p is an on/off notion while being reasonable to assert that p is something that comes in degrees. In the passage I quoted, Williamson says that asserting that there is snow outside is “quite reasonable” for him, rather than merely “reasonable.” But the idea is also intrinsically plausible. If p is false, then I do not have a warrant to assert that p . Full stop.¹⁴ Intuitively, however, asserting that there

¹³I will come back to this distinction later in this chapter when we look at the knowledge norm of assertion.

¹⁴One might object that it is sometimes appropriate to say that S might have *more* warrant to

is snow outside might be more reasonable for Williamson than it is for me, if we had different evidence which bears on whether there is snow outside, for the probability that *I* know that there is snow outside on *my* evidence could be different from the probability that *Williamson* knows that there is snow outside on *his* evidence. All it would take for this to be the case is that there is some proposition *m* such that *m* is a member of one of those evidence sets but not a member of the other.¹⁵

For example, suppose that I live on Williamson's, street, that I, like him, am also looking through my window and am also oblivious to the fact that there is a film crew going around spreading snow-like foam. Suppose further that I, *unlike* Williamson, heard fake news broadcasted by the film crew which said that a snowstorm was expected to dump more or less the same amount of snow I now think I see outside my window. Intuitively, the proposition that I know that there is snow outside is (even if ever so slightly) better supported by my evidence than by Williamson's. For one thing, my knowledge that it was broadcasted that there would be a snowstorm is a reason I have for believing that there is snow outside that is independent from it appearing to me that this is a "banal case of

assert that *p* than *S** does. For example, when *S* has two conclusive arguments for *p* while *S** has only one. However, since having a warrant to assert that *p*, on a knowledge-first kind of picture, requires that *p* have probability 1 on one's evidence, strictly speaking *S* cannot have "more" warrant than *S**, for 1 is the highest degree of probability according to the probability calculus. We can, however, accommodate the intuition that it is sometimes adequate for someone to say "S has more warrant to assert that *p* than *S** does." Sentences like "S has more warrant to assert that *p* than *S** does" are easily understood as expressing something about the comparative degree of psychological certainty of *S* and *S**. Those sentences express, in appropriate contexts, that both *S* and *S** have a warrant for asserting that *p*, but *S* is more (psychologically) certain of *p* than *S** is. Thanks to Peter Klein for discussion here.

¹⁵That Williamson takes the reasonability of an assertion that *p* to vary with how probable "I know that *p*" is on one's evidence is explicit in the following passage: "One may reasonably assert *p*, even though one does not know *p*, because it is very probable on one's evidence that one knows *p*" (2000: 256). The locution "very probable on one's evidence" should be understood as saying that it is reasonable for one to assert that *p* only if the evidential probability of "I know that *p*" on one's evidence falls within some arbitrarily high range of values short of 1, say, .95 to .99. If the probability of "I know that *p*" on one's evidence falls within that range, then one is in a position to assert that *p* *reasonably*. See Williamson [2000, p.256 fn.9] for a formalization of this account.

perceptual knowledge.” The consequence is that, given R, even though asserting that there is snow outside is reasonable for both Williamson and I, asserting that proposition is more reasonable for me than it is for him, for “I know that there is snow outside” has a higher probability of being true conditional on my evidence than on Williamson’s.

Of course, since neither Williamson nor I know that we know that there is snow outside, neither one of us have a warrant to assert “I know that there is snow outside.” Note also that from the fact that it is reasonable for both Williamson and I to assert that there is snow outside it does not automatically follow that it is reasonable for either one of us to assert that we know that. In order for it to be reasonable for us to assert “I know that there is snow outside” we would have to satisfy R, in which case “I know that I know that there is snow outside” would have to be sufficiently likely to be true on our evidence. Maybe that proposition, conditional on either of our evidence sets, is sufficiently likely to be true. Maybe it is not. For one thing, it is not clear that the evidence Williamson and I have makes “I know that I know that there is snow outside” as likely to be true as it makes “There is snow outside”. I raise this issue only to put it to the side, however, since it is not very relevant to my concerns here. What *is* relevant, however, is that we are clear about the fact that it being reasonable for me to assert that ϕ does not automatically entail that it is reasonable for me to assert that I know that ϕ .

Also, whether I have a warrant to assert that p and whether it is reasonable for me to assert that p may not overlap. No matter how reasonable it is for someone to assert that p, it does not follow that S has a warrant to assert that p. Neither does it follow from one having a warrant to assert that p that asserting that p is reasonable for one, for one might know that p (and thereby have the warrant to assert that p) even though “I know that p” has an evidential probability on one’s evidence arbitrarily close to 0.

Now, belief and assertion are clearly related. If one asserts that p sincerely and in one's "own voice," rather than, say, as the spokesperson for an institution, then, other things being equal, one believes that p . Likewise, if one believes that p , then, one is disposed to assert that p (under the right circumstances). It is natural to think that, at least in the usual cases, assertion is the external counterpart of belief or, conversely, that belief is the internal counterpart of assertion. Williamson himself accepts something that comes really close to this view. He says that assertion is "the verbal counterpart of judgment and judgment ... the occurrent form of belief"¹⁶ and that asserting p represents oneself as at least believing that p .¹⁷ The tight connection between assertion and belief suggests that the distinction between having a warrant to assert and it being reasonable for one to assert may be extended to the case of belief.¹⁸ Here are a couple of noteworthy similarities between assertion and belief.¹⁹

One might have the warrant to assert that p , but withdraw from actually asserting that p for many different reasons (e.g., one is tired and asserting that p would start a new conversation and prevent one from going to bed; or one might be afraid one will be harmed if one asserts that p in the presence of p -deniers). Likewise, we should think that some agent having a warrant to believe that p does not entail that the agent knows that p , for he might have the warrant to believe that p and not believe that p (e.g., because one fails to put "two and two together" or because one is epistemically timid).

Secondly, in the case of assertion, one might have a warrant to assert that

¹⁶Williamson [2000, p.10].

¹⁷Williamson [2000, p.252 fn.6].

¹⁸As I will suggest below when we discuss the knowledge norm of inference, a similar strategy can be used to address cases where, even though the agent inferred something from something else she does not know, it seems appropriate to say that she knows the inferred proposition.

¹⁹See Adler [2002] for an apt defense of the claim that belief is the external analog of assertion. See, specially Adler [2002, p.274-7] for an extensive list of similarities between belief and assertion.

p, assert that p, but do so not in virtue of having a warrant, but in virtue of something else altogether. For example, one might know that p and therefore have a warrant to assert that p, but assert that p because one wants to hurt one's audience's feelings. In cases like this, even though one has a warrant to assert, one's assertion is not as appropriate as it would have been if one asserted *in virtue of* having that warrant.²⁰ Similarly, one might have a warrant to believe that p, believe that p but do so not in virtue of having such a warrant, but in virtue of something else. For instance, suppose I form the belief, at t, that Mary is having an extra-marital affair in virtue of being ill disposed towards members of Mary's ethnicity. Suppose further that, at t1, I receive an email with pictures showing Mary and her lover kissing during a business trip they took to Europe. In this case I have a warrant to believe that Mary is having an extra-marital affair, but I believe that she is having an affair not in virtue of having that warrant, but in virtue of the prejudice I have. In cases like this, even though one has a warrant to believe that p, one's belief that p is not as appropriate as it could have been if one believed that p *in virtue of* having that warrant.

The same two points apply, *mutatis mutandis*, to it being reasonable for one to assert/believe that p. It might be reasonable for me to assert/believe that p, but I might either (i) fail to assert/believe that p; or (ii) assert/believe that p in virtue of something different than what makes asserting/believing that p reasonable for me.

One other important point before we move on. Even though there are all those similarities between the propriety of asserting and believing, there is also one noteworthy difference between them. If one has a warrant to assert that p, then one knows that p and, given that knowledge is factive, p is true. The consequence is that, in the case of assertion, what provides the warrant is always a *known proposition*. The same is not true in the case of belief, since one's belief

²⁰Turri [2011] calls this view the "Express Knowledge Account of Assertion".

that p may be warranted by something *different* from a known proposition (e.g., by perceptual experience).²¹ This is important because, even though one can have a warrant to believe some falsehood f , f itself cannot serve as a warrant for anything else, for only known propositions warrant belief.

So, keeping those points in mind, I am now ready to characterize two distinct notions of justification based on this brief discussion of the distinction between having a warrant and it being reasonable to believe:

(Jw) S is w -justified in believing that p only if S has a warrant to believe that p and S believes that p in virtue of having that warrant²²

(Jr) S is r -justified in believing that p only if it is reasonable for S to believe she has a warrant to believe that p .

We may now apply these two senses of “justified” to the snow case in the following way. Given Jw , Williamson is not w -justified in believing that there is snow outside. Given Jr , however, Williamson is r -justified in believing the same proposition.

The protagonist of Gettier cases is in a similar situation. Given Jw , it is true that Smith is not w -justified in believing either that Jones will get the job or the truth he infers from this falsehood (i.e., that the man who will get the job has ten coins in his pocket). On the other hand, given Jr , Smith is r -justified in believing both propositions.

The distinction between w -justification and r -justification also helps bring to light a central feature of Gettier cases. Even though the target true belief in those

²¹cf. Williamson [2000, p.201-2].

²²This account may bear some relationship to Alvin Plantinga’s definition of “warrant” in Plantinga [1993] in the sense that my account, like his, entails that all things being equal, believing truly in virtue of having a warrant w is sufficient to give you knowledge of that truth. I say “may” because I am not sure I fully understand Plantinga’s account. I am sure of one thing, though: I do not want my account of “warrant” to be committed to anything like a theory of proper functioning. For criticism of Plantinga’s account of warrant and proper function, see the contributions to Kvanvig [1996], specially Klein [1996].

cases is clearly not an item of knowledge, it does enjoy some positive epistemic status. J_w helps explain why the target true belief is not known - it is not w -justified - while J_r helps explain what is epistemically good about it - it is r -justified. Moreover, the reason why Smith is r -justified but not w -justified can be traced back to the fact that his true belief relies essentially on a falsehood. Since falsehoods cannot be w -justified, the fact that Smith deduces a true proposition from the false one cannot yield w -justification for his belief in the true one, for his belief in the false one has no w -justification to transmit via deduction to begin with. Because Smith deduces the true proposition from the false one and Smith's belief in the false one is r -justified, Smith is r -justified in believing the true one as well. Unfortunately for Smith, r -justification is not sufficient for knowledge.

So, Comesaña's and Kantin's argument against $E=K$ fails. The argument mistakenly presupposes that "justification" expresses only one epistemically relevant concept and that this unequivocal concept is present in Gettier cases. With the help of the distinction between warranted and reasonable assertions we can sort out different senses of "justified" and thereby accommodate the intuition that Gettiered subjects are justified. This completes the task of discharging the objection that $E=K$ eliminates the Gettier Problem.

3.3 Objection: Dogmatism

One might have a different worry about taking $E=K$ on-board: dogmatism. In this section I will look at two arguments that, if sound, would strongly suggest that $E=K$ breeds dogmatism. This would be an extremely bad result since dogmatism is the vicious intellectual trait of close-mindedness.

Dogmatism or close-mindedness arises because the arguments I will discuss below suggest that, if one knows that p , then one can rationally ignore or discount counterevidence against p . If those arguments succeeded, that would give us a

reason to reject $E=K$, and, given that $E=K$ entails KFK , we would lose support for that latter thesis. Both arguments exploit the fact that it seems to be a consequence of the kind of probability 1 epistemology I favor in this essay (i.e., the view which says that one knows that p iff the probability of p on one's evidence is 1) entitles knowers to their known propositions *no matter what counterevidence against what they know they encounter*. The idea is that, if p has probability 1 on S 's evidence, then S is entitled to discount any evidence against p , since new evidence either will be redundant (p already is maximally probable for S) or misleading (it could lower the probability of p below 1 and thus destroy S 's knowledge). As we will see, both arguments fail to establish this conclusion.

The arguments I will consider originated from a lecture Saul Kripke gave at Cambridge University to the Moral Science Club in 1972. The first version of the argument was presented by Gilbert Harman in his book "Thought."²³ The second version was recently published by Kripke himself in the first volume of his collection of papers "Philosophical Troubles."²⁴ Both arguments are supposed to show that, if knowledge eliminates the possibility of error, then it breeds dogmatism or close mindedness. I will present my own solution to the dogmatism puzzle raised by those arguments. But, before I do that, I will have to distinguish two forms of dogmatism - synchronic and diachronic dogmatism - and try to convince you that the most prominent alleged solutions do not work. Once I have done that, I will propose that the puzzle arises only if one has the wrong view of what makes disregarding counterevidence rational. The last section sums up our discussion of dogmatism and draws lessons from it.

²³Harman [1973].

²⁴Kripke [2011].

3.3.1 Dogmatism: Synchronic and Diachronic

Saul Kripke's discussion of dogmatism²⁵ is an effort to "prove" a position suggested by Norman Malcolm in the following passage:²⁶

It could happen that in the next moment the ink-bottle will suddenly vanish from sight; or that I should find myself under a tree in the garden with no ink-bottle about; or that one and more persons should enter this room and declare with apparent sincerity that they see no ink-bottle on this desk... Having admitted that these things *could* happen, am I compelled to admit that if they did happen then it would be proved that there is no ink-bottle here now? Not at all! I could say that when my hand seemed to pass through the ink-bottle I *should* then be suffering from hallucination; that if the ink-bottle suddenly vanished it would have miraculously ceased to exist...

...No future experience or investigation could prove to me that I am mistaken. Therefore, if I were to say 'I know that there is an ink-bottle here', I should be using 'know' in the strong sense...

In saying that I should regard nothing as evidence that there is no ink-bottle here now, I am not *predicting* what I should do if various astonishing things happened. If other members of my family entered this room and, while looking at the top of this desk, declared with apparent sincerity that they see no ink-bottle, I might fall into a swoon or become mad. I *might* even come to believe that there is not and has not been an ink-bottle here. I cannot foretell with certainty how I should react. But if it is *not* a prediction what is the meaning of my assertion that I should regard nothing as evidence that there is no ink-bottle here?

That assertion describes my *present* attitude towards the statement that here is an ink-bottle. It does not prophesy what my attitude *would* be if various things happened. My present attitude towards that statement is radically different from my present attitude toward those other statements (e.g., that I have a heart). I do *now* admit that certain future occurrences would disprove the latter. Whereas no imaginable future occurrence would be considered by me *now* as proving that there is not an ink-bottle here. These remarks are not meant to be autobiographical. They are meant to throw light on the common concepts of evidence, proof, and disproof.

Malcolm thinks his knowledge that there is an ink-bottle in front of him exemplifies a strong sense of "know," a sense that denotes the kind of knowledge the

²⁵Kripke [2011, p.43].

²⁶Malcolm [1952, 185-6]. The emphases are all Malcolm's.

knower might properly hold on to when faced with counterevidence. I have nothing to say about the linguistic claim that there are different senses of “know” and that they denote knowledge relations of different strength.²⁷ I simply note, in passing, that nothing in my discussion of dogmatism hangs on this linguistic claim. There are two points in Malcolm’s passage that I would like to emphasize. First, the fact that he distinguishes a synchronic version of dogmatism from a diachronic one. Second, that he takes his remarks to shed light on the concept of “evidence.” In relation to this second point, I will argue that dogmatism presents us with a puzzle only if we fail to appreciate the fact that one is entitled to treat counterevidence to *p* as misleading evidence only if one knows that one knows that *p*. The brute fact that *e* is misleading evidence against *p* is not enough to rationalize one’s discounting of *e*. The fact that *e* is misleading evidence has to be *my reason* for this discounting and this is achieved only if I know that I know that *p*. This reply will be developed further below. I will first analyze the arguments for synchronic and diachronic dogmatism and then show how we can block the conclusion of those arguments.

Let me flag an important ambiguity in the arguments for dogmatism I will be discussing.²⁸ There are two distinct formulations of the dogmatist thesis, and those two formulations are not always kept separate. There are *descriptive* and *normative* formulations of dogmatism. A descriptive formulation of dogmatism is one where the author merely describes the attitudes of a dogmatist, but says nothing about whether those attitudes are the ones that agents, when in similar circumstances, *ought* to adopt. Normative or prescriptive formulations of dogmatism might also describe the attitudes of the dogmatist, but they also take the normative step of prescribing those attitudes to agents in similar circumstances.

²⁷But see Hetherington [2001], and Goldman [1999], Goldman and Olsson [2009] for similar linguistic claims about “knows.”

²⁸Thanks to Ernie Sosa for making me appreciate this issue.

In this essay I will be mainly concerned with the normative or prescriptive formulation of dogmatism. The passage I quoted above from Malcolm can be seen as presenting both a descriptive and a normative formulation of dogmatism.

According to Malcolm, when he says that he should not take anything as evidence that there is no ink-bottle in front of him, he is not “*predicting* what [he] should do if various astonishing things happened” and they all seemed to suggest to him that there was no ink-bottle in front of him. He *could* be making a prediction, but he does not want to be taken as having made one. But there is a *normative* view being endorsed by Malcolm here as well, and that is the view according to which, if I know that p, then I should not let any future (apparent) evidence undermine my knowledge that p. Call this view *diachronic dogmatism*. This is the relevant normative claim I will take the diachronic dogmatist to be making.

In this passage, Malcolm also discusses what we might call *synchronic dogmatism*. He says he is not making a prediction, but that he is concerned with his “present attitude” towards the proposition that there is an ink-bottle in front of him. He thinks that “no imaginable future occurrence would be considered by [him] now as proving that there is not an ink-bottle” in front of him. This passage is also ambiguous between a descriptive and a normative form of (synchronic) dogmatism. On the one hand, Malcolm can be taken to be *describing* what he would do now in case someone challenged his knowledge. On the other hand, he can also plausibly be taken to be *prescribing* what anyone in his situation should do at that moment.

My discussion will focus on the normative version of both synchronic and diachronic dogmatism. Even though both versions of dogmatism are clearly absurd, they seem to be at least consistent with the view that knowledge has probability 1 on one’s evidence.

3.3.2 Harman's Argument for Diachronic Dogmatism

Gilbert Harman presented the following argument for diachronic dogmatism and went on to reject its conclusion.²⁹

If I know that h is true, I know that any evidence against h is evidence against something that is true; so I know that such evidence is misleading. So once I know that h is true, I am in a position to disregard any future evidence that seems to tell against h .

Harman's claim that once he knows that h is true he is in a position to disregard any *future evidence* that counts against h tells us that he is concerned with the diachronic form of dogmatism. The dogmatist Harman is describing will use the following argument to justify her dogmatism with respect to any future counterevidence she may find:

1. If S knows that h at t_0 , then S knows at t_0 that any counterevidence e^* against h he may find at t_1 is evidence for something false (i.e., evidence for $\neg h$) and S is in a position at t_1 to disregard e^* . [Assumption]
2. S knows that h at t_0 . [Assumption]
3. S knows at t_0 that any counterevidence e^* to h he may find at t_1 is evidence for something false and S is in a position at t_1 to disregard e^* . [1, 2 by *modus ponens*]
4. S knows at t_1 that she has been offered counterevidence e^* against h . [Assumption]
5. Thus, S knows at t_1 that e^* is evidence for something false and S is in a position at t_1 to disregard e^* . [3,4]

²⁹Harman [1973, p.148].

Since (1)-(5) are true (if true at all) of any subject, proposition, and body of evidence, generalized diachronic dogmatism appears to be justified. For instance, according to this view, since I now know that my bike is in my backyard, chained down to a chain-link fence, I can rationally disregard any future evidence that seems to indicate that my bike is not in my backyard. So, if my neighbor calls me and says he drove by and did not see my bike chained to the fence in my backyard, I am in a position to disregard the counterevidence he offers me. Similarly, I am in a position to do the same to any counterevidence to that proposition. If my wife calls me and says she cannot see my bike through the window in the back of our house, I am still in a position to disregard her report. I may also disregard as a fabrication or a prank the phone call I get from the police saying that they found my bike (with my phone number and address on it) at the outskirts of town. This is clearly a ridiculous result. The question is: what is the problem, exactly, with the dogmatist's argument for diachronic dogmatism?

Maybe we should start by saying something about what the dogmatist might mean by "disregard." According to the *Oxford English Dictionary*, "to disregard x" means, roughly, "to pay no attention to x." That cannot be exactly what the dogmatist means, however, for as the bike example shows, one is not always in a position to avoid paying attention to counterevidence - we answer phone calls, turn on TVs and listen to what other people say. Our attention is grabbed by our information-rich environment; we are usually not capable of turning off our attention at will. Thus, the diachronic dogmatist must mean something else by the claim that we may disregard future counterevidence against what we know. One must "disregard" counterevidence even if it grabs one's attention. It follows that, according to the diachronic dogmatist, it is rational for me to disregard future counterevidence to my knowledge that my bike is now in my backyard, even when I attend to it.

So, according to premise 3, agents may rationally disregard future counterevidence for what they know, because counterevidence e for S 's knowledge that p is "evidence for something false;" that is, S may rationally ignore counterevidence e against p , because S knows that e is evidence for something false, i.e., $\neg p$. But, knowing that e is evidence for something false is a good enough reason not to believe that e . Thus, in the context of diachronic dogmatism, "to (rationally) disregard x " entails "to (rationally) forbear believing x ."

Against diachronic dogmatism, Harman argues that as soon as an agent acquires evidence against what she knows, she loses her knowledge that evidence against what she knows is misleading. That is, the agent may know at t_0 that any future evidence against what she knows is misleading, but as soon as she actually acquires the evidence against what she knows, she loses her knowledge of the "misleadingness" of counterevidence. Harman makes this point in the following passage:

The argument for the paradox overlooks the way actually having evidence can make a difference. Since I now know that [h], I now know that any evidence that appears to indicate something else is misleading. That does not warrant me in disregarding any further evidence, since getting that further evidence can change what I know. In particular, after I get such further evidence I may no longer know that it is misleading. For having the new evidence can make it true that I no longer know that new evidence is misleading.³⁰

Harman's solution emphasizes the defeasible character of knowledge. However, his reply does not go far enough. It seems to me that there is still something wrong with his suggestion that someone could know *now* that any future evidence against her knowledge that p is misleading just by knowing that p . Intuitively, people cannot know that any future evidence against p is misleading by merely knowing that p . So, to say that the agent loses knowledge of the misleadingness of future counterevidence once she acquires that counterevidence fails to get at the

³⁰Harman [1973, p.149].

heart of the dogmatist problem. The result is that Harman's "solution" leaves the synchronic version of the problem untouched. Furthermore, this proposed solution seems to be dialectically inefficient. The dogmatist is arguing that knowing that p , for any p , is an overriding reason for one to disregard counterevidence against p . So, since the dogmatist knows (q) "e is misleading evidence against p ," she knows that any counterevidence against q is misleading and she has an overriding reason to disregard this counterevidence. By insisting that the dogmatist loses her knowledge that q when she finds counterevidence for p , Harman fails to find common ground between him and the dogmatist by merely assuming that q does not give the dogmatist an overriding reason to disregard counterevidence.³¹ But why does Harman treat the problem in this way? Here is a plausible explanation: Harman took the dogmatist to be *describing what she would do* were she to encounter counterevidence to what she knows, rather than *prescribing what she should do* in those circumstances. By conflating the normative and the descriptive versions of the problem, Harman failed to address the normative problem raised by Malcolm. The idea is that Harman's solution is a lot more plausible if we take him to be *describing* what usually goes on when one is presented with counterevidence.

The Junk Knowledge Response to Diachronic Dogmatism

One might also wonder: even if Harman's approach does not solve the synchronic version of the problem, how good is it as a solution to the diachronic problem? In order to answer this question appropriately I will look at a recent and more elaborate version of Harman's account. I will assess a few objections to this view and argue that it has enough resources to answer them. This will only take the view so far, however, for, as we will see in 3.3.4, the account ultimately employs

³¹One could point out that, just like epistemologists are not trying to convince the skeptic that he should give up skepticism, Harman is not trying to convince the dogmatist that he should give up dogmatism. If my reply to the dogmatist below is on the right track, then we can do better than Harman, and the analogy between dogmatism and skepticism breaks down.

a mistaken notion of rational disregarding of counterevidence.

Roy Sorensen³² proposed a solution to the diachronic dogmatism in the spirit of Harman's solution. His view, like Harman's, leaves synchronic dogmatism unsolved. Nonetheless, the view has put Harman's solution in a different (and perhaps better) perspective. Sorensen considers the following version of the argument for rational dogmatism:³³

1. My car is in the parking lot.
2. If my car is in the parking lot and Doug provides evidence that my car is not in the parking lot, then Doug's evidence is misleading.
3. If Doug reports he saw a car just like mine towed from the parking lot, then his report is misleading evidence.
4. Doug has reported that a car just like mine was towed from the parking lot.
5. Doug's report is misleading evidence.

According to Sorensen, Harman's "talk of further evidence changing what I know [is] a point about the co-justifiability of (3) and (4)."³⁴ The thought is that I cannot justifiably believe the conjunction of (3) and (4) even though I know (3) and (4) individually. (3) and (4) are not independent, in the sense that the more confident I am that Doug has reported that a car just like mine was towed from the parking lot, the less confident I am that his report constitutes misleading evidence. Hence, when my confidence in the proposition that Doug has reported that a car just like mine was towed from the parking lot is high enough for me to know it, my confidence in the proposition that his report constitutes misleading

³²cf.Sorensen [1988, 2003, 2012].

³³The example comes from Sorensen [2012, section 6.2].

³⁴Sorensen [1988, p.438].

evidence decreases below the level required for knowledge.³⁵ Of course, this line of reasoning is correct only if Doug and I are in the same location and we both know that. If I am sitting in what I know to be my car in what I know to be Detroit, speaking on the phone with Doug whom I know to be in Chicago, then, if he says he saw a car like mine towed from the parking lot adjacent to the office building in Chicago where we both work, I can be quite certain that he is reporting that a car just like mine was just towed from the parking lot in Chicago and also know that his report is misleading evidence that my car has just been towed from that parking lot.³⁶

That our confidence in a conditional decreases as our confidence in its antecedent increases is an odd feature, for, usually, we can come to know the consequent of a conditional we know whenever we learn its antecedent.³⁷ This feature of Harman's solution might bother us. One might worry that Harman's solution to the paradox is *ad hoc*, in that it suggests that a very particular type of indicative conditional (i.e., conditionals about the misleadingness of counterevidence to what one knows) cannot be used to extend one's knowledge through *modus ponens*. Harman's response to the paradox would be less attractive if it turned out that only conditionals about a particular subject matter (i.e., misleading evidence) were the object of defeasible knowledge in the way he suggests. Fortunately, there are other examples of indicative conditionals that are such that, even if agents know them, the more confident agents are in their antecedents, the

³⁵Sorensen [1988, p.438].

³⁶Thanks to Ernie Sosa for discussion here.

³⁷Sorensen [1988, p.438-9] thinks that this feature of the paradox helps explain why it is a mistake for us to read (3) as the subjunctive conditional: (3*) I know that if Doug were to report that he saw a car just like mine towed from the parking lot, then his report would be misleading evidence. Taking (3) to be (3*) would allegedly hide the tension between accepting (3) and accepting (4). Also, the inference from (1) and (2) to (3*) is invalid, for the usual Lewis/Stalnaker semantics for " $\Box \rightarrow$ " takes $p \rightarrow (x \rightarrow p)$ to be necessarily true while $p \rightarrow (x \Box \rightarrow p)$ is only contingently true.

less confident they become in their consequents.³⁸ Consider:

(J1) If Dick Cheney is an honest guy, then I am a monkey's uncle.

(J2) If she has spoken falsely, I shall always be deceived.

I can know (J1) and (J2) but I cannot (rationally) perform *modus ponens* using either one of them if I learn their antecedents. In fact, the point of asserting (J1) is to invite hearers to perform *modus tollens*. (J2) has an unknowable consequent and I would not perform *modus ponens* if I learned the antecedent - I would stop believing the conditional. So, the phenomenon explored by Harman in response to the dogmatism paradox is actually widespread and is not restricted to any particular conditional about misleading evidence. Moreover, this phenomenon is not restricted to conditionals. Some disjunctions also can be known, but not reasonably used in disjunctive syllogisms. John Hawthorne discusses the following example concerning disjunctions:

Suppose there are two newspapers, The Times and The Guardian, which I trust equally well for the purposes of obtaining soccer information. With good reason: both are extremely reliable in their reporting of soccer results. I look in The Times and find a Manchester United victory reported. I trust the report. The report is in fact correct. Under such circumstances people are inclined to say I know both that The Times said that Manchester United won and also that Manchester United won. Let us suppose I also know that The Guardian will have reported a result for the Manchester United game. I deduce that either The Times and The Guardian correctly reported a Manchester United victory or else The Guardian made a mistake about the Manchester United result. Suppose, in fact that, unbeknownst to me, The Guardian did make such a mistake.³⁹

In this example Hawthorne knows the disjunction

³⁸See Sorensen [1988] for a detailed discussion of the examples I am about to synoptically describe. These examples are modified from Ginet [1980].

³⁹Hawthorne [2004, p.71-4].

- (a) Either The Times and The Guardian correctly reported a Manchester United victory or else The Guardian made a mistake about the Manchester United result.

Hawthorne then notes that people might want to deny that he knows (a) because they notice that he could not reasonably infer that the Guardian made a mistake if he were to learn that it reports a Manchester United loss, and that they think that this is a sign of a belief that is not knowledge. This is a mistake, says Hawthorne:⁴⁰

Following Roy Sorensen, let us use the term ‘junk disjunctive knowledge’ to refer to cases where knowledge of a disjunction is such that, for one of the disjuncts, if one were to come to believe that it is false, that would destroy knowledge of the disjunction... We are apt to confuse useful disjunctive knowledge with junk disjunctive knowledge. Having recognized that disjunctive knowledge is not useful, we are prone to think that it is not knowledge at all.

In a nutshell, some conditionals and disjunctions are “useless” or “junky,” in the sense that one’s knowledge of them cannot withstand *modus ponens* or disjunctive syllogism. In contrast, knowledge that is “useful” in this way is dubbed “robust.” S’s knowledge that p is robust with respect to evidence e iff, $P(p | e) > \Omega$, where “ Ω ” refers to some probability threshold required for knowledge. Thus, S’s knowledge of a conditional $(p \rightarrow q)$ is robust with respect to p iff, $P(p \rightarrow q | p)$ is high enough to make the threshold required for S to know $(p \rightarrow q)$, whatever that threshold may be. So, for example, “If Dick Cheney is an honest guy, then I am a monkey’s uncle” is not robust with respect to “Dick Cheney is an honest guy,” because $P(\text{If Dick Cheney is an honest guy, then I am a monkey’s uncle} | \text{Dick Cheney is an honest guy})$ is as close to 0 as one would like. This conditional is intuitively robust with respect to “I am not a monkey’s uncle,” however, for $P(\text{If Dick Cheney is an honest guy, then I am a monkey’s uncle} | \text{I am not a$

⁴⁰Hawthorne [2004, p.72-3].

monkey's uncle) is as close to 1 as one would like. The thought here is that, the conditional is robust with respect to my learning its antecedent or the negation of its consequent just in case detaching (and thereby coming to believe) the consequent or the negation of the antecedent is more justified for me than believing the negation of the conditional.

According to Assaf Sharon and Levi Spectre⁴¹ (henceforth, “Sharon&Spectre”), though, there is a problem with the junk knowledge story. The problem is that the Harman/Sorensen solution argues that “knowledge isn’t vulnerable until it is.” Sharon&Spectre call this the “Hypothetical Problem” raised by the argument for dogmatism. This is basically what I have been calling (the descriptive version of) the problem of “diachronic dogmatism”. Thus, the Hypothetical Problem has to do only with what the agent would do in the future if she comes across counterevidence. This, they think, does not address the problem raised by the knowledge the agent allegedly *now* has of the misleadingness of counterevidence in general. This is essentially the point I raised above against Harman’s proposed solution to diachronic dogmatism. So, we can agree with Sharon&Spectre that Sorensen, like Harman, does not deal with the real issue involved in the dogmatism puzzle, namely the synchronic problem of saying that the agent can know *now* that counterevidence against *p* is misleading by the mere fact that she knows that *p*. According to them,

[Sorensen and Harman] claim that, knowing that *p*, *S* can know that any counter-evidence is misleading. Independently of what is to happen at t_1 , it does not seem plausible to say that at t_0 , having only the knowledge that *p* is true, *S* can infer from the general truism that if something is true, evidence against it is misleading, that counter-evidence to *p* is misleading. In other words, given the epistemic means at his disposal, the knowledge ascribed to *S* [to the agent in the relevant contexts] is dubious.⁴²

Sharon&Spectre call the problem raised by synchronic dogmatism the “Epistemic

⁴¹Sharon and Spectre [2010, p.310-1].

⁴²Sharon and Spectre [2010, p.310].

Problem.” This is the problem of “providing a reasonable account of the state of the agent regarding knowledge” of the indicative conditional about the misleadingness of counterevidence in Harman’s and Sorensen’s discussions.⁴³ As my discussion of Harman’s argument for diachronic dogmatism above suggested, I think Sharon&Spectre are right about this. They are also right about it not being clear how one would extend Sorensen’s account to cover cases, like Hawthorne’s, in which the subject is not aware of the counterevidence.⁴⁴ For all Sorensen has said, one can know that p even if one’s “environment” is saturated with counterevidence against p. All that it is required is that one not be aware of this counterevidence.⁴⁵ One problem with Sharon&Spectre’s views is that they seem to be saying that the dogmatist has in mind the wrong *description* of how people react to counterevidence against what they know. They might be right about this. Still, the *normative* version of the puzzle remains untouched and the dogmatist can insist that he is concerned with what agents should do, not with what they actually do.

3.3.3 Kripke’s Argument for Synchronic Dogmatism

I now turn to synchronic dogmatism. I have argued that synchronic dogmatism is at the heart of dogmatism. It is now time to unpack this problem and solve it. In his discussion of dogmatism Kripke makes it clear that he is interested in the synchronic argument for dogmatism. He says the following about Harman’s diachronic version of the argument:

Well, one need not disagree with what Harman says about the acquisition of the new evidence (at least for typical cases). But remember that I was

⁴³Sharon and Spectre [2010, p.311].

⁴⁴Sharon and Spectre [2010, p.312-3].

⁴⁵In a sense, this aspect of Sorensen’s account should not come as a complete surprise to us. Sorensen takes himself to be following Harman’s approach to diachronic dogmatism and Harman explicitly denies that counterevidence one is unaware of can defeat one’s knowledge (cf. Harman [1973, p.146-9]). I will come back to this issue below.

talking about a resolution to be made in advance. Just because the subject wishes to avoid a loss of knowledge such as Harman describes, so for that reason she or he makes the resolution. ... Harman is right that if [contact with counterevidence] nevertheless occurs, one may well lose the knowledge that p , and hence no longer know that the counterevidence is misleading. But just this is why the subject resolves not to get into such a situation!

Kripke seems to get the order of Harman's explanation backwards here. As we saw above, Harman thinks we lose knowledge of p because we lose our knowledge that future counterevidence to p is misleading, not the other way around. Kripke [2011, p.43] wants to "prove" the following principle of synchronic dogmatism:

(d) If I know something now, I should, as a rational agent, adopt a resolution [now] not to allow any future evidence to overthrow it.

(d) might seem initially plausible. For example, one might think that we can appeal to something like (d) in order to rationally disregard counterevidence produced by car salesmen, politicians, and charlatans in general. Since I know that equal work should be paid equal, it seems that I should, as a rational agent, adopt a resolution *now* not to allow any future evidence to overthrow this knowledge of mine. So, one might think that (d) does sometime rationalize the discounting of evidence. I think even this more cautious endorsement of (d) is wrong and I will explain why below. For example, (d) yields the implausible result that I may rationally discount, in the bike example above, any evidence suggesting that my bike is not in the backyard.

A fair reconstruction of Kripke's "proof" of d can be stated thus:

1. S knows that p . [assumption]
2. If p is true, then any evidence against p is misleading (where "misleading" is to mean "leads to a false conclusion"). [assumption]
3. S knows that 2 is the case. [assumption]

4. S knows that any evidence against p is misleading. [from 1, 2, 3, by knowledge closure]
5. If S knows that taking an action T leads to consequence C, and S wishes above all else to avoid C (i.e., that is the only relevant issue), then S should resolve now not to take action of type T. [assumption]
6. If S knows that accepting any evidence against p has, as a consequence, a false belief and S wishes, above all else, to avoid acquiring a false belief, then S should resolve now not to accept any evidence against p. [from 5 by instantiation]
7. S knows that accepting any evidence against p has, as a consequence, a false belief. [from 4]

Thus,

8. S should resolve now not to accept any evidence against p. [from 6,7 by *modus ponens*]

Clearly, 2 and 5 are the crucial premises in this argument. Kripke assumes premise 2 is analytic and offers no support for it. He also seems to assume that an intuitive, non-explicitly articulated sense of “rational” is all the support premise 5 needs. In section 3.3.4 I will argue that, all things considered, premises 2 and 5 make problematic assumptions about counterevidence and should not be accepted. Before I turn to my reply to this argument, I will discuss, in the next section, why the only existing detailed solution to synchronic dogmatism I know fails. If I am right and the proposed diagnosis is mistaken, then the account I offer in section 2 is the only plausible solution.

A Failed Response to Synchronic Dogmatism

Sharon&Spectre have not only offered a criticism of Harman’s solution to diachronic dogmatism but also a solution to synchronic dogmatism. In this section I will discuss their proposed solution and why it fails.

According to Sharon&Spectre, the dogmatist puzzle is “instructive regarding formal features of knowledge, in particular the closure of knowledge under entailment.”⁴⁶ Let us look at why they posit that and why I think this is the wrong thing for us to say about this puzzle.

Sharon&Spectre argue that the dogmatist cannot rationally infer premise 4 from 1, 2 and 3 in Kripke’s argument, because of a knowledge closure failure.⁴⁷ If they are right, synchronic dogmatism would be avoided. Their “solution” is mistaken, however. The argument for synchronic dogmatism does not involve knowledge closure failure. Rather, it involves the failure of a much less plausible, but different principle about the transmissibility of positive epistemic status. Consider, first, the following popular formulation of the closure principle for knowledge Sharon&Spectre endorse:⁴⁸

(SPC) $\Box(\forall p)(\forall q)(\text{If } S \text{ knows that } p, \text{ competently deduces } q \text{ from } p, \text{ and thereby comes to believe } q, \text{ while retaining knowledge of } p \text{ throughout, then } S \text{ knows } q)$ ⁴⁹

The literature on the validity of SPC is large and I will not revisit it here.⁵⁰ Suffice to say that most epistemologists think that SPC enjoys strong intuitive support.⁵¹

⁴⁶Sharon and Spectre [2010, p.320].

⁴⁷Sharon and Spectre [2010, p.320].

⁴⁸Sharon and Spectre [2010, p.309].

⁴⁹This version of SPC is John Hawthorne’s. cf.Hawthorne [2004, 34-6] and Hawthorne [2005, p.29].

⁵⁰But see, among others, Dretske [1970, 2005], Klein [1995], Nozick [1981], Hawthorne [2004, 2005], and de Almeida [2011].

⁵¹See Hawthorne [2005] for some of the costs of denying SPC. Also, even if knowledge-closure

Now, consider a closely related principle that deals with the transmissibility of evidential warrant. Here is an intuitive statement of this principle:

(TES) $\Box(\forall p)(\forall q)(\text{If } e \text{ is a source of positive epistemic status for } S\text{'s belief that } p, \text{ and } p \text{ entails } q, \text{ then } e \text{ is a source of positive epistemic status for } S\text{'s belief that } q)$ ⁵²

TES is stronger than SPC; that is, the former entails the later, but the converse is not true. TES is stronger than SPC; that is, the former entails the latter, but the converse is not true. In particular, while TES excludes the possibility that the entailing reason is itself a source of positive epistemic status for the entailed proposition, SPC is compatible with this possibility. So, if, in a case C, e is a good enough source of positive epistemic status for the belief that p, S deduces q from p, and e is not a good enough source of positive epistemic status for the belief that q, then TES is false in C. SPC is not necessarily false in C, however. In order to show that SPC is false in C, it also has to be true, in C, that p is not a good enough source of positive epistemic status for the belief that q.

I will now show that Sharon&Spectre ignore the distinction between TES and SPC and that this leads them to believe that the dogmatist argument features a failure of SPC. I will claim, on behalf of the dogmatist, that those arguments feature, at most, a failure of TES. But, since we already knew TES was false⁵³ and the dogmatist can use SPC instead, Sharon&Spectre's case against dogmatism fails.

is not valid (i.e., even if it has some false instances), it does not follow that all instantiations of the principle are false. Hence, knowledge-closure deniers owe us a story about what makes some instantiations of the principle true and others false. Nozick [1981] tells us that the true instantiations preserve tracking while the false ones don't. For a recent, alternative story see de Almeida [2011].

⁵²Klein [1995], for reasons that will become clear below, calls TES "mistaken target". de Almeida [2011] calls it "evidential closure."

⁵³Cf. Klein [1995] and Wright [2002].

The mistake of taking counterexamples to TES to be counterexamples to SPC is most prominent in discussions of skepticism. Consider the following example adapted from Dretske [2005, p.14-6]. I know there is a cookie jar in front of me on the basis of normal visual perception. That there is a cookie jar full of cookies in front of me entails that those things inside the jar are not cookie-like papier-mâché fakes. But, the source of my knowledge that there is a cookie jar full of cookies in front of me (namely, by seeing the jar of cookies in front of me) is not itself a potential source of knowledge that those things inside the jar are not cookie-like papier-mâché fakes. Thus, I can know that there is a cookie jar in front of me by seeing it, know that there being a cookie jar full of cookies in front of me entails that those cookie-like things inside the jar are not cookie-like papier-mâché fakes, but fail to know the later by vision alone.⁵⁴ Therefore, TES is false.

Now, Dretske and others took examples such as this to show that SPC is false. The idea is that the failure of SPC blocks the skeptic from arguing by *modus tollens* that, since I do not know that the things inside the jar are not cookie-like papier-mâché fakes, I do not know that there is a cookie jar in front of me. This is a mistake, however, and the skeptic should resist Dretske's move.

The example, and others like it, show (at most) that TES is false, not that SPC is. Those cases are compatible with my knowledge of the entailing reason *itself* being a good enough source of positive epistemic status to ground knowledge of the entailed proposition. TES is the mistaken target of Dretske and others who try to undermine SPC. Sharon&Spectre make structurally the same mistake with respect to dogmatism Dretske makes with respect to skepticism. They look at two instances of the general argument for dogmatism I gave above and claim that one can resist the dogmatic conclusion if one rejects SPC.

⁵⁴Of course, this is Dretske's more recent version of the classic zebra-in-the-zoo example discussed originally in Dretske [1970, p.1015-6].

Sharon&Spectre claim that they have identified a certain pattern in Hawthorne's and Sorensen's discussion of junk knowledge and that this pattern is evidence that the real challenge posed by the dogmatism paradox has to do with SPC. Remember Sorensen's example:

1. My car is in the parking lot.
2. If Doug reports he saw a car just like mine towed from the parking lot, then his report is misleading evidence (i.e. evidence for something false).
3. Doug has reported that a car just like mine was towed from the parking lot.
4. Doug's report is misleading evidence.

According to Sharon&Spectre,⁵⁵ the dogmatist cannot rationally infer 4 and thereby come to know it because of the following closure failure: my evidence (e) "I just parked the car in the school lot" warrants my belief in (1) "my car is in the parking lot," 1 entails (2) "If Doug reports he saw a car just like mine towed from the parking lot, then his report is misleading evidence (i.e. evidence for something false)," but e fails to warrant 2.

The dogmatist can reply to Sharon&Spectre in the following way: "You are right that I could not rationally infer 4 from 1-3 *if* e were the only reason to accept 2. The problem is that I know 1 and I can deduce 2 from 1. So, I can come to know 2 by deducing it from something else I know, 1." The point of the dogmatist reply is that there is more than one evidential path to 2. One evidential path, the one Sharon and Spectre think does not rationalize the dogmatist's belief in 2, requires that e be a good enough source of positive epistemic status to accept *both* 1 and 2. The other evidential path, the one Sharon and Spectre neglect, requires e to be a good enough source of positive epistemic status for 1 and 1 itself to be a

⁵⁵Sharon and Spectre [2010, p.320].

good enough source of positive epistemic status for 2. So, it is not the case that we can reject the dogmatist conclusion for the reason Sharon&Spectre say we should. As we will see below, even if I am wrong about all of this, Sharon&Spectre reply to dogmatism does not generalize to include cases of dogmatism which do not explicitly rely on TES or SPC.

This reply notwithstanding, Sharon&Spectre make a structurally identical point about Hawthorne's Manchester United case:

reading in The Times that Manchester United won the match seems like proper justification for the belief that Manchester United in fact won. It is, however, utterly inappropriate evidence for knowing that if The Guardian says otherwise, it is mistaken.

Again, the claim is that Hawthorne's evidence (e) "the Times says that the Manchester United won the match" is a good enough source of positive epistemic status for his belief in (p) "the Manchester United won the match," p entails (q) "if The Guardian says that $\neg p$, it is mistaken," but e is not a good enough source of positive epistemic status for Hawthorne's belief in q. The dogmatist should reply to this in the same way she replied above: "This case would pose a problem to SPC only if the only evidential path available to Hawthorne required e to be a good enough source of positive epistemic status for both p and q; however, all that SPC requires is that p itself be a good enough source of positive epistemic status for q, and that is satisfied in this case."

Thus, I conclude, that Sharon&Spectre's proposed solution to the dogmatism paradox ultimately fails. The paradox poses, at most, a challenge to TES, a principle we already had independent reason to reject.⁵⁶ The dogmatism puzzle is not generated by SPC.

⁵⁶Taking a counterexample to TES to be a counterexample to SPC is not an unfamiliar mistake. As Klein [1995] shows, Dretske [1970] also took himself to be refuting SPC with his famous zebra example when in fact the example only clearly refutes the stronger TES. For that reason Klein called TES Dretske's "mistaken target."

3.3.4 Dogmatism Depuzzled

Both Kripke's argument for synchronic dogmatism and Harman's argument for diachronic dogmatism are fallacious. And they are fallacious for the same general reason: they both employ mistaken accounts of counterevidence and evidential defeat. We can see that most clearly if we scrutinize Kripke's argument for synchronic dogmatism. Focusing on this argument allows us to deal with the most virulent version of the paradox. If I am not justified in being dogmatic about my knowledge that p *now*, when I have no knowledge of counterevidence against p , how could it be justified for me to be dogmatic about my knowledge that p *later*, when I am presented with counterevidence for p ? Hence, undermining the argument for synchronic dogmatism will give us a good enough reason to think the argument for diachronic dogmatism is also unsound. In particular, I will now try to show that two premises in Kripke's argument are false and that his characterization of dogmatism is inherently implausible. First, remember premise 2 of his argument:

- (2) If p is true, then any evidence against p is misleading (where "misleading" is to mean "leads to a false conclusion").

Kripke thinks 2 is analytic - i.e., true in virtue of meaning. This premise needs some tidying up before we can evaluate its plausibility, however. (Since something may be analytic but not obvious, this is not *per se*, an objection to 2 itself.) First of all, 2 assumes that p being true entails that any evidence against p is evidence for $\neg p$. This idea is hard to reconcile with clear intuitions about how counterevidence undermines knowledge, though. Suppose I know there is a tree in front of me in virtue of my reliable vision. Suppose also that a friend I know to be generally reliable tells me that I took a drug that generates tree illusions 95% of the time. I did not take such a drug, however. In this case, (e) that my friend told me that I took the drug, is evidence against (p) "there is a tree in front of

me,” without it being necessarily evidence in favor of (not-p) “it is not the case that there is a tree in front of me.” The lesson is that having a reason to doubt the reliability of one’s knowledge is not always a reason to believe the negation of what one believes on the basis of that source. Sometimes evidence that undercuts one’s warrant for p is not evidence for $\neg p$, but evidence that favors suspension of judgment as to whether p.

John Pollock⁵⁷ called this type of counterevidence “undercutting evidence” and contrasted it with what he called “rebutting evidence” - i.e., counterevidence to p that *is* evidence for $\neg p$. For example, (e*) “the tree-like thing in front of me is made of styrofoam” would count as rebutting evidence against “there is a tree in front of me.” *Undercutting evidence* defeats knowledge in virtue of undermining the evidential connection holding between one’s evidence and one’s knowledge. *Rebutting evidence*, on the other hand, defeats knowledge in virtue of it being evidence for the negation of what one knows. So, contrary to what premise 2 assumes, counterevidence for p is not always evidence for $\neg p$. In light of this, let’s rewrite premise 2 thus:

(2*) If p is true, then any *rebutting evidence* against p is misleading (where “misleading” is to mean “leads to a false conclusion”).

There is still a problem with 2*: it presupposes that, once accepted by the agent, rebutting evidence always destroys knowledge. However, this assumption is clearly false. In certain cases rebutting evidence e against p does not destroy the agent’s knowledge that p, even though e is a reason for believing $\neg p$. This seems to be the case whenever the agent has all the resources needed to neutralize the undermining effect of rebutting evidence. For example, suppose that after having seen thousands of black ravens, S, a novice ornithologist, has a strong inductive argument for (p) “All ravens are black.” Suppose further that one day

⁵⁷Pollock and Cruz [1999, 196-7].

a trustworthy friend tells S that (e) he has seen a bird that looks like a green raven. (This friend, we may suppose, has no reason to lie to S, but he cannot be characterized as someone who recognizes birds reliably. S knows that, and S does bring this knowledge to bear on his assessment of the situation.) Now, if S knows (q) that there are many Common Crackles around and (r) that this fact explains how both p and e can be true at the same time, it appears that S can use q and r (and his knowledge about his friend's reliability) to neutralize the defeating effect of e against p while, at the same time, coming to know e. Hence, contrary to what premise 2* says, accepting counterevidence does not always reduce the plausibility of something known so that it is no longer known (i.e., it does not always lead to belief in the negation of what one knows). This also falsifies step 4 in Kripke's argument: one does not know *a priori* that a particular bit of counterevidence increase the plausibility of something false.

This shows that premise 2 in Kripke's argument cannot be *literally* true. Maybe there is a more charitable interpretation of this premise, one that enables the dogmatist to rationally retain it. As far as I know no one has offered such a premise while at the same time securing the normatively strong dogmatist conclusion.

Remember premise 5 in Kripke's argument:

- (5) If S knows that taking an action T leads to consequence C, and S wishes above all else to avoid C (i.e., that is the only relevant issue), then S should resolve now not to take action of type T.

What does it mean for an agent to "resolve" not to take a certain type of action? Perhaps the most natural interpretation is that, for S to resolve to proceed in some way w, is for S should *form the intention to proceed in way w*. So, according to 5, rational agents should form the intention not to proceed in a way w if they know that proceeding in this way will lead to a consequence they want, above

all else, to avoid. If I want, above all else, not to be harmed in my sleep, and I know that nothing that might come through my cabin door during the night has my best interest in mind, I should *now* form the intention to bolt the door before I go to sleep.⁵⁸ Similarly, suppose that, after running many experiments to find out whether *p*, Liz comes to know that *p*. Since Liz wants to avoid, above all else, forming a false belief about whether *p*, premise 5 would then entail that Liz should form the intention *now* not to read any journal articles, conduct any more experiments, or otherwise collect any evidence that might bear on whether *p*, for proceeding in those ways would only lead her away from the truth about whether *p*.⁵⁹

A first objection to premise 5 is that even if one forms an intention to proceed in a certain way, one might fail to carry out this intention either because of weakness of the will or through a change of mind. This is particularly troubling in the epistemic case, for, even if we suppose that agents do not change their minds about the way they should proceed, we do not seem to have direct control over whether we form, retain or lose beliefs. As a result, forming the intention to avoid counterevidence to what one knows will prevent one from losing one's knowledge only if we can successfully "hide" from counterevidence. If one does run into counterevidence, not even the strongest intention not to lose one's knowledge can guarantee that that will not happen.

Maybe we can waive this objection by noting that the fact that we *cannot* avoid losing our knowledge on the face of counterevidence does not change the

⁵⁸I owe this example to Ernie Sosa.

⁵⁹Although this is the most charitable reading of premise 5, from a purely exegetical point of view it is most likely not the reading Kripke himself had in mind. Kripke [2011, p.44] explicitly denies that one could maintain the resolution to avoid places, people, and things potentially containing misleading evidence. According to him, it should be possible for us to "ignore" misleading evidence, regardless of whether we wanted to be confronted with it or not. This suggestion also betrays what Kripke takes to be the motivating thought behind his argument for dogmatism: the thought, defended in Malcolm [1952, 185-6] and Hintikka [1962, p.20-1], that there is a sense in which knowledge is conclusive (cf. Kripke [2011, p.39]).

fact that we *should* at least try to avoid doing so. The dogmatist will remind us that his argument is about how we *should behave* and not about *how we do in fact behave* in the epistemic domain. Here the dogmatist can employ a strategy familiar to any one who has had to teach skepticism about the external world to undergrads. Invariably, some student will say in reply to a skeptical argument that “The argument must be mistaken, because we do form beliefs about the external world and we do rely on them when we act.” It is part of the instructor’s job to remind those students that the skeptical challenge is about whether we have the *right* to hold those beliefs, and not about whether we do in fact hold them.

There are other problems with premise 5, however.⁶⁰

First, the mere fact that I know that p cannot justify me in taking counterevidence against p to be misleading evidence against p. That I know that p must be *my reason* for disregarding counterevidence against p, if such disregarding is to be rational. A plausible suggestion is that r is my reason to ϕ only if *I know that r*.⁶¹ If I know that p, then I disregard evidence against p rationally only if I know that I know that p. To think that one *should* discount or ignore counterevidence to what one knows as misleading evidence even if one does not know that one knows is analogous to thinking that we should ignore what a corrupt police officer tells us to do, even though we do not know we are legally obligated to ignore it. The fact that S knows that p, and the fact that there is a law which states that S may not ϕ make it the case, respectively, that counterevidence to p is misleading and that orders to ϕ are inappropriate calls to action. But this does not entail that S can *rationally* disregard counterevidence against p or disobey orders to ϕ ; we ignore counterevidence and disobey orders rationally only when

⁶⁰Thanks to Duncan Pritchard for discussion here.

⁶¹This is in line with the knowledge-first approach to practical reasons in Hawthorne and Stanley [2008]. If Williamson’s E=K is right (Williamson [2000]), then r is a reason for believing only if r is known.

the reason for doing either is *our reason* for doing those things. It is not enough that *there is* such a reason “out there,” as it were; I need to be aware of it in order for my ignoring or disobeying to be rational. Hence, S ignores or disregards counterevidence against p *rationally* only if she at least knows that she knows p . Otherwise she cannot rationally take counterevidence against p to be *misleading* evidence against p . In other words, knowing that one knows is at least necessary for rational discounting of counterevidence as misleading evidence. The upshot is that the dogmatist’s normative claim that one should form the intention to disregard counterevidence is false. At most, one should form the intention to disregard counterevidence to p *only if one knows that one knows that p* . But is second order knowledge sufficient for rational disregarding of first-order counterevidence?

If S knows that she knows that p , then she is clearly in a stronger epistemic position than if she merely knew that p . If S knows that p on the basis of evidence e and she knows that she knows that p , then she knows that e is good evidence for p - how else could she know that she knows that p ?⁶² That e is good enough evidence for believing that p has to be part of her reasons - it seems - for believing she knows that p , if she is to know that she knows that p . When is evidence for believing that p good enough to base knowledge that p , and how does S get to know that? The following seems plausible:⁶³ given that one knows that p on the basis of evidence e iff it is not (epistemically) possible that $(e \& \neg p)$, evidence is good enough to base knowledge when p has epistemic probability 1 on one’s total evidence (since epistemic probabilities distribute over epistemic possibilities, if the epistemic probability were less than 1, then $\neg p$ would be epistemically possible). So, in order for S to know that she knows that p , (1) “S knows that p ” must be

⁶²Even if God had told her that she knew that p , it still seems true that she could easily infer from the fact that God told her that she knows that p that her evidence for p is good enough to give her knowledge that p .

⁶³Remember that we are supposing, with Kripke, that knowledge requires the elimination of all possibilities of error.

true; (2) S must believe “S knows that p;” and, (3) “S knows that p” must have probability 1 on S’s total evidence. Notice that, S is not required to believe (or to know) that she satisfies conditions (1)-(3) in order to know that she knows that p, those conditions just need to be satisfied in order for her to know that she knows that p. Hence, knowledge of knowledge is clearly possible. Now, I ask again: is knowing that I know that p a good enough reason to make disregarding counterevidence to p rational for me? Not always. At a minimum, there are cases in which it is far from clear that having second-order knowledge about one’s knowledge constitutes a good enough reason to disregard misleading evidence rationally. To see that, consider, first, a case in which it seems plausible to think that second-order knowledge gives the subject a good enough reason to ignore misleading evidence.⁶⁴

Liz is watching the news on TV. The news anchor comes on. He has breaking news: Air Force One has crashed and no one knows yet if the President survived the crash or not. The anchor says they have double-checked the information with a reliable but anonymous source from the Pentagon. Liz believes the report. Moreover, what the report says is true and, we may suppose, it being reported by that trustworthy TV network raises the probability of what is reported to 1 for Liz. Now, suppose Liz thinks about whether she can trust the report and, after deliberating about how likely what was reported is on her evidence, she believes (truly) that she knows that Air Force One has crashed. Since Liz’s second-order belief satisfies conditions (1)-(3) above, she knows that she knows that Air Force One has crashed. Now, suppose Liz runs into her thoughtful and reliable neighbor Rob, who has not heard the news about the crash, but argues, on general grounds that, since plane crashes are statistically improbable events and Air Force One is very well maintained, Liz should not believe the President’s plane crashed. It

⁶⁴The case I am about to present is a modified version of the Assassination Case discussed in Harman [1973, p.143-4].

is plausible to think that, if Liz is able to marshal her second-order reasons for believing she knows the plane crashed in response to Rob's challenge, she can rationally discount or neutralize the counterevidence Rob offered her and keep her knowledge that Air Force One crashed. I take this to be a clear case in which second order knowledge secures first order knowledge in the face of counterevidence by rationalizing the treatment of such evidence as misleading evidence. It is important to note that the appeal to second order knowledge also dissipates the charge of dogmatism on Liz's part. Her reaction to Rob's challenge is clearly non-dogmatic. That seems true even if she had anticipated Rob's challenge and formed the intention to ignore the counter evidence as misleading on the basis of her second order knowledge. No dogmatism here. So far so good.

Liz's second-order knowledge does not shield her from *all* counterevidence, however. If we modify this case a bit, it becomes a lot less clear that Liz's second-order knowledge is sufficient to rationalize her discounting counterevidence. Suppose that, sometime after Liz hears the report about the crash the network reports they have made a mistake and that the President's plane has not crashed. (The plane *did* crash and this is just part of a misinformation campaign to prevent political and economical turmoil that is certain to ensue after the President's sudden death.) In this modified case one can argue (plausibly) that Liz's second order knowledge that she knows the plane crashed does not constitute a good enough reason for her to rationally disregard the counterevidence she has now been presented with. In fact, since she believes she knows the plane crashed in part *because* she believes the network is a reliable informant, it seems that she should reconsider her views on the crash, given that this reliable informant is now taking back something it reported earlier. I take it that cases such as this strongly suggest that knowledge of knowledge is not always a sufficient reason for an agent to disregard counterevidence rationally.

At this point one might raise the following objection:⁶⁵ I don't need to know that I know that p in order to run the dogmatist argument; I just need to know *that* p . Since I know p I can properly form the intention to disregard counterevidence against p .

This is a mistake. The dogmatist argument has no chance of engaging me if I am unsure or in doubt as to whether I do in fact know that p or not. If I am unsure about whether I know that p or not, then how can I even form the intention to protect my *knowledge* that p ? What is more, even if I form this intention, this is an appropriate (rational) intention for me to form only if the reason I have for forming it is (epistemically) good enough. But, if this is right, then I properly (rationally) form the intention to disregard counterevidence against p only if I know that I know that p .

Contrary to what Kripke's (d) suggests, the fact that S knows that p is not a good enough reason for S to ignore counterevidence against p . "S knows that p " is a good enough reason for S to ignore counterevidence against p only if "S knows that p " is S 's reason for thinking counterevidence against p is *misleading* evidence (how else can S know that counterevidence to p is misleading evidence against p ?), and this requires that S knows "S knows that p ." And even then, there are no guarantees that our first-order knowledge can rationally withstand the effect of counterevidence.

This brings us to the second problem with premise 5. Even if I know that I know that p , it might not be rational for me to discount counterevidence against p , because counterevidence against p might well include truths I value more than I value p . Here, it matters whether counterevidence for what we know is constituted by false or true propositions, because dogmatism tells us we should "ignore" or "disregard" counterevidence for what we know. If ignoring/disregarding p entails believing $\neg p$, then, whenever counterevidence to what the dogmatist knows is

⁶⁵Thanks to Ernie Sosa for pressing this point in conversation.

constituted by true propositions, disregarding that evidence will entail that the dogmatist believes something false. If ignoring/disregarding p does not entail believing not- p , but only not forming the belief that p , then the dogmatist does not believe something false when he disregards false counterevidence, but he still misses out on truths that are potentially more important, by his own lights, than what he knows.

In sum, Kripke's argument for synchronic dogmatism fails - premises 2 and 5 are very likely false. Dogmatism has been depuzzled.⁶⁶

3.3.5 A new argument for dogmatism?

John Hawthorne has presented the following case for dogmatism.⁶⁷ This example provides an argument for dogmatism neither the junk knowledge strategy nor Sharon&Spectre can account for. Since the considerations we offered in the previous section can account for this case, the example can be seen as evidence that our strategy fares better than those other accounts.⁶⁸

Suppose a scientist makes me the following offer: "You know that p . However, there is a lot of misleading evidence out there and it is very likely that you will lose your knowledge that p if you run into it. Luckily for you, I have developed a pill that immunizes your knowledge against misleading evidence. If you take this pill, your knowledge that p will be immune to misleading evidence, i.e., you will not lose your knowledge that p ."

Should I take the pill?

Thinking that I should take the pill seems to vindicate some form or other of synchronic or diachronic dogmatism: one could think that it would be rational

⁶⁶In writing this section I have greatly benefited from discussion with Claudio de Almeida. Claudio also graciously suggested I used the title "Dogmatism Depuzzled."

⁶⁷cf.Hawthorne [2004, p.181]. I have modified Hawthorne's case to fit both synchronic and diachronic dogmatism, instead of diachronic dogmatism only.

⁶⁸Thanks to John Hawthorne and Peter Klein for discussion on the issues in this section.

for me *now* to take the pill (synchronic) or that it would be rational for me *later*, after I take the pill, to disregard counterevidence against *p* (diachronic).

Note that, at least *prima facie*, neither Sharon&Spectre, nor Harman have much to say about this case. Since it does not involve knowledge closure, Sharon&Spectre's reply is just not relevant here. On the face of it, it seems also out of place for us to say that taking the pill is not the rational thing for me to do, because, as Harman would have it, once I encounter counterevidence for what I know I stop knowing that this evidence is misleading (by assumption, the pill will not allow that to happen). Those are not the right kind of answer to this case.

On the other hand, the considerations I marshaled above in reply to Kripke's argument tell against taking the pill and avoid dogmatism in this case. First of all, how trustworthy is the "scientist" offering me the pill? If he is not trustworthy at all, I am not in a position to know what he says. In particular, I am not in a position to know that I know that *p*, or if he is untrustworthy, that his pill will in fact protect my knowledge against counterevidence. So, for the sake of argument, let us make the scientist perfectly trustworthy (i.e., deserving of the highest confidence) and the pill he offers flawless (i.e., it never fails to work). Am I then justified in taking the pill? Not quite. I have made the scientist and his pill perfectly reliable, but that does not mean that I know those things about the scientist and his pill. If I do not know anything about the scientist's trustworthiness, I would be gullible to take him for his word and believe either that I know that *p* or that his pill actually works, and I cannot be justified in taking the pill on the basis of my gullibility. But suppose we concede all that; that is, suppose we assume that I *know* the scientist and his pill are 100% reliable. Am I *then* justified in taking the pill? No. Even if I come to know that I will not lose my knowledge that *p*, if I take the pill, I cannot be sure I will not be missing out on truths that I myself consider more important than *p*. Suppose "p" stands for the (actually true) proposition "Newark New Jersey and Porto Alegre Brazil are

sister-cities” and suppose that the counterevidence against p I will find if I do not take the pill includes truths about global markets and financial regulation, truths that, as it turns out, I am lot more interested in learning than I am in learning the truth about Newark and Porto Alegre. In fact, if you asked me, I would say I would trade losing my knowledge about the cities for the more important (in my opinion) knowledge about world economy. But, if I cannot be sure the pill will not prevent me from learning truths I care a lot more than I care about the truth that p , I have no reason to take the pill.⁶⁹

3.3.6 Taking Stock

This brief discussion of Harman’s and Kripke’s arguments, as well as of a case offered by Hawthorne allow us to draw some general lessons about dogmatism.

Knowing that counterevidence is misleading. Kripke and Harman assumed that counterevidence is always evidence for something false. This ignores undercutting evidence, which undermines knowledge without making the negation of what we believe probable. They also assumed that one cannot know a piece of counterevidence and maintain one’s knowledge. This, they thought, happens either because counterevidence is constituted by false propositions or because it is not rational for an agent to believe the target proposition and the counterevidence at the same time. Both assumptions are mistaken, however. Some counterevidence to what we know is true and, potentially, more valuable than the item of knowledge for which it is counterevidence. Also, the undermining effect of rebutting evidence may be neutralized even if we incorporate the rebutting evidence

⁶⁹What if the scientist designed a pill capable of filtering out only potentially undermining evidence I care *less* than I care for p ? Would I then be justified in taking the pill? Maybe, but then the pill would not protect my knowledge that p *against potentially undermining evidence I care more about than I care for knowing that p* . The result would be a partial form of dogmatism, not the fully general form we are considering here. Something similar could be said, I think, if there was nothing I cared more than I cared for p . In that case, the resulting view is skepticism, not the generalized form of dogmatism people usually have in mind. Thanks to Carlotta Pavese and Daniel Rubio for discussion here.

into our total body of evidence.

Ignoring misleading evidence. Whether counterevidence for what we know is constituted by false or true propositions matters because dogmatism tells us we should “ignore” or “disregard” counterevidence for what we know. If ignoring/disregarding p entails believing $\neg p$, then, whenever counterevidence to what the dogmatist knows is constituted by true propositions, disregarding that evidence will entail that the dogmatist believes something false - this contradicts our epistemic goal of believing all and only truths. If ignoring/disregarding p does not entail believing $\neg p$, but only not forming the belief that p , then the dogmatist does not believe something false when he disregards true counterevidence, but he still misses out on truths that are potentially more important, by her own lights, than what she knows.

Higher-order knowledge. Knowing that p is not, strictly speaking, a reason for us to ignore/disregard counterevidence to p as misleading. One needs to at least know that one knows that p in order to rationally disregard counterevidence against p as misleading. This lesson applies both to Harman’s and Kripke’s arguments. Another related lesson is the following: if I do not know that I know that p , then I do not know that my evidence makes p highly likely to be true and I cannot, thus, rationally believe that evidence against p is evidence against something true. For all I know, evidence against p is evidence *for* something true, namely $\neg p$!

3.4 The Knowledge Norm of Inference

Knowledge-first epistemologists have proposed a variety of knowledge norms: the knowledge norm of belief (believe only what you know);⁷⁰ the knowledge norm

⁷⁰cf. Williamson [2011, 2000].

of assertion (assert only what you know);⁷¹ and the knowledge norm of action (act only on what you know).⁷² According to these philosophers, these norms constitute the normative benchmark for our attributions of appropriateness and inappropriateness of belief, assertion and action. In the same spirit, in this section I propose and explore a knowledge norm of inference.

If KFK is true, then one would expect our practice of appraising reasoning, from an epistemic point of view, to show sensitivity to it - at least when the reasoning in question is non-suppositional. One would expect reasoners to be sensitive to a norm which would constitute the normative benchmark against which particular inferences would be assessed. In that spirit, consider the following norm concerning non-suppositional reasoning:

(KNI) One must: infer only from what one knows.

KNI seems to suggest itself given that, according to KFK, one has inferential knowledge only if one knows all the premises essentially involved in one's inference, and given that we value knowledge. For instance, suppose I reason that I will not have to walk home because my car is in the parking lot and if my car is in the parking lot I can drive home. However, suppose I reason that way right after you told me that at least one car is stolen every day from that parking lot. In this scenario, it seems natural to say that, unless I have a good reason to dismiss what you said, I should refrain from inferring that I can drive home. I should refrain from inferring that I can drive home because I do not know my car is in the parking lot and that is an essential premise in my reasoning. To the extent that this negative assessment of my inference is appropriate, it is appropriate because one is sensitive (if only tacitly) to the norm that one must infer only from what one knows.⁷³

⁷¹cf. Williamson [2000, ch 11].

⁷²cf. Hawthorne and Stanley [2008].

⁷³Notice that this assessment of the case (and KNI itself) are independent of any particular

KNI explains also what is wrong with the following inference. Suppose I have one ticket in a large and fair lottery paying ten million dollars to its sole winner. Now, suppose the drawing was this morning and, without having checked the result yet, I reason thus:

1. My ticket is a loser.
2. If my ticket is a loser, I do not have the means necessary to pay my mortgage in full.

Thus,

3. I do not have the means necessary to pay my mortgage in full.

Clearly, I should not believe that I do not have the means necessary to pay my mortgage in full. The most obvious explanation for that fact is that I do not know my ticket is a loser. Knowledge requires probability 1 on one's evidence and mere knowledge of the odds of winning is not enough evidence for me to meet that threshold. KNI captures this idea and accounts for the propriety of the criticism.

Consider also the situation in which the protagonist of Gettier cases finds himself. Suppose Smith infers that the man who will get the job has *coins* in his pocket from his belief that the man who will get the job has *ten coins* in his pocket. Clearly, Smith does not know the inferred proposition and it seems natural to say that he should not have inferred it, for he inferred that proposition from something he does not know. Again, to the extent that this negative assessment of Smith's inference is appropriate, it is appropriate because one is sensitive (if only tacitly) to the norm that one must infer only from what one knows.

Epistemic norms that are weaker than KNI cannot account for the propriety of this criticism in a way that is as straightforward as KNI. For instance, if the

explanation of why I fail to know that my car is in the parking lot. Plug in your favorite explanation for *why* I fail to know that premise; whatever that explanation is, that is independent from the fact that it seems appropriate to criticize my inference. Thanks to Lisa Mirachi for discussion here.

normative standard of proper inference was given by NFG (i.e., “One must: infer only from what one truly and justifiably believes”), then I should expect at least my criticism of Smith to “sound bad,” for he has a justified true belief in all the premises of his reasoning (as do I in the car and lottery cases). A “true belief norm” fares even worse: if such a norm were correct, it would not only predict that our criticism of Smith was mistaken, but it would also entail that one could acquire inferential knowledge even if the sole reason one believed the premises was because of a coin toss. Hence, KNI has an advantage over weaker epistemic norms.

If KNI is true, and one infers something on the basis of premises one does not know, then one is at fault for breaking that epistemic rule. One might have an *excuse* for breaking KNI, however, if it is reasonable for one to believe that one knows the premises involved in one’s reasoning. In other words, one has an excuse for breaking KNI if it is reasonable for one to believe that one conforms to KNI. If that is the case, then one has an excuse to believe the conclusion of one’s inference, but that belief does not amount to knowledge, because KNI was broken. The thought is that a belief needs to be more than excusable in order to be a case of knowledge - just like an action needs to be more than excusable in order to be virtuous. One’s belief in the conclusion of an inference is more than excusable only if one conforms to KNI.

KNI is the obvious counterpart, in the theoretical domain, of the knowledge norm of action in the practical domain.⁷⁴ If one thinks that one should not act unless one knows all the propositions one uses as reasons for one’s actions, then one thinks that rational action requires good reasons and that reasons are good only if they are known. This is, *mutatis mutandis*, the idea behind KNI: rational

⁷⁴John Hawthorne and Jason Stanley make a similar point, in passing, in Hawthorne and Stanley [2008, p.577].

belief requires good reasons and reasons are good only if they are known.⁷⁵

What is more, it seems that someone who accepts the knowledge norm of belief is committed to KNI: if one's belief that *p* depends essentially on one's belief that *q* and one fails to know *q*, then one has not only violated KNI, but one has also violated the knowledge norm of belief, for one fails to know something one believes, namely, *q*.

One familiar complaint about other knowledge norms is that, even if we are not skeptics, they are too demanding. It might be pointed out that we are not always in a position to know whether or not we know the premises in our reasoning and that we should settle for a weaker epistemic standard such as justified belief, for we are relatively good at finding out whether or not we are justified in believing the premises of reasoning.

The usual reply to this kind of objection to knowledge norms can also be used in defense of KNI. First, KNI does not require one to know that one knows the premises to one's reasoning - first order knowledge (i.e., *knowing* the premises) is enough. Second, the alleged fact that KNI is broken frequently does not show KNI is not the epistemic rule governing epistemically proper inference; just like a high number of people breaking the rule "one must: stop at stop signs" does not show that this rule is not what governs proper driving behavior at intersections with stop signs.

KNI would be too demanding only if we knew too little, for, if that were the case, then we would have very little to work with in our inferences. However, since we clearly sometimes reason properly, we might plausibly take that to confirm KNI. Skeptics about knowledge will, of course demur, but we have as much reason to think the skeptic is right about this as we have to think that there is no knowledge at all. While skepticism is an important philosophical question, it poses no special problem for KNI.

⁷⁵cf. Unger [1975], Williamson [2000], Sutton [2007].

Since no one I know has written about a knowledge norm of inference, I will use the knowledge norms of assertion and action as a model and draw comparisons between those more familiar norms and the knowledge norm of inference. Hopefully, that will be a good first step towards making this novel idea clearer.

3.4.1 Knowledge Norms: Assertion, Action and Inference

In Williamson [2000, ch.11], it is argued that knowledge and assertion have a “special relationship” not exhausted by the fact that when we assert something we usually express or communicate what we know. This special relationship is a *normative* one and it is analogous to the relationship between knowledge and belief.⁷⁶

Since knowledge and assertion are related normatively, it makes sense to attribute to assertions the property of being “known.” But, since assertions are also impolite, sincere, relevant, and so on, it follows that assertion bears a normative relation with all those other properties, not only with knowledge. The same happens with any other act. A jump can be criticized or praised because it is long or short, cowardly or brave. Williamson argues, however, that “some norms are more intimately connected to the nature of asserting than any norm is to the nature of jumping.”⁷⁷ Someone who knowingly asserts something false breaks a rule in the same sense in which someone who knowingly breaks a rule in a game - she cheats. According to Williamson, the speech act of asserting, like a game and unlike jumping, is constituted by rules. It follows that not all possible norms of assertion are on a par. Williamson’s goal, then, is to characterize the constitutive rule of assertion using the rules of games as an analogy. But what is the difference between the constitutive and the non-constitutive rules of an act?

For Williamson, constitutive rules are essential to an act, and necessarily, the

⁷⁶Williamson [2000, p.238].

⁷⁷Williamson [2000, p.238].

rule governs every performance of the act. This implies that the analogy between games and assertion is somewhat idealized, for games change their rules over time without changing their identity. What is more, the rules of games and languages are arbitrary conventions. So, in this sense, the constitutive rule of assertion is not a convention, because conventions are not stable, they change over time.

It is also worth noting that one might have an answer to the question “What are the rules of G?” without having an answer to the more ambitious question “What are non-circular necessary and sufficient conditions for a population to play a game with rule R?” Compare: the umpire who knows that the constitutive rule of game G is R* might have no idea what the non-circular necessary and sufficient conditions for someone to play G with R* are. So, we can know what is the rule R that constitutes assertion, without knowing what are the necessary and sufficient conditions that need to be satisfied in order for us to be said to assert something. However, we need to have *some* grasp of how to answer the more ambitious question about assertion in order to answer the question about what rule constitutes assertion, otherwise we would have no idea whatsoever if we were asserting something or not. Having this minimal grasp of how to answer the more ambitious question should not require “a full philosophy of games”, however, according to Williamson.⁷⁸

Even if one castles *after* one has moved one’s king (thereby breaking the rule governing castling in chess), that does not mean that one has stopped playing chess. Likewise, even if one lies (thereby breaking the rule of assertion), that does not mean that one has stopped asserting. Actually, the opposite seems to be the case: one is subject to criticism *because* one has flaunted the constitutive rule of assertion.⁷⁹

What kind of normativity is involved in the rule of assertion? The normativity

⁷⁸Williamson [2000, p.240].

⁷⁹Williamson [2000, p.240].

of a constitutive rule is neither moral nor teleological, according to Williamson. Breaking a constitutive rule is as moral as breaking a rule of a game or language. Lying and cheating in a game can have moral implications, but, according to Williamson, this implication is “made possible only by the non-moral rules which constitute the game.” The normative bite of a constitutive rule does not come from the internal or external aims involved. One may break the rules *and* achieve the aims of the game, or one may fail to achieve the aims of the game *while* obeying all its rules.

Williamson puts forward the hypothesis that there is only one constitutive rule of assertion. It should take the following form:

(The C rule) One must: assert p only if p has C.

The idea is that assertion is the *unique* speech act A whose unique rule is the C rule. Mastering the speech act of assertion involves implicitly grasping and conforming to the C rule in much the same sense that mastering a language or a game involves implicitly grasping and conforming to the rules governing that language or that game. This seems particularly true when we reflect on the activity of learning a second language and the activity of learning a new game. If one were to start either of those activities by first memorizing all rules involved in the language or in the game, one would never start speaking the language or playing the game. Immersion in the culture and in the *modus vivendi* of the native speakers using the language or in the ways of experienced players is the most effective way of learning a new language or a new game; memorizing all the rules of that language or game is not necessary (and perhaps not even sufficient) for learning how to speak a new language or how to play a new game. In both cases, one is more or less aware of the most important rules, and by immersing oneself in the language or game and imitating the “moves” native speakers or experienced players make, one gradually improves one’s understanding of the

language and of the game.

Williamson's analogy between games and assertion is somewhat idealized, of course. What game do we know which is such that it is constituted by one single rule? Clearly, not many.⁸⁰ The idea of a game constituted by one single rule is possible, however.

Thus, I am inclined to accept the less strict reading of Williamson's analogy between games and assertion, one that keeps in mind that the analogy involves, to a great extent, a high degree of idealization. So, if the reader feels uncomfortable with the talk of a single norm constituting the act of inferring, I would be happy to fallback on the more modest (but perfectly fitting for my purposes) claim that KNI is the single constitutive norm of *epistemically appropriate inference*. One could say a similar thing about assertion.

Two obvious candidates for property C in the C rule are truth and warrant. Respectively,

(The Truth Rule) One must: assert p only if p is true.

(The Warrant Rule) One must: assert p only if one has warrant to assert p.

If having a warrant to assert p amounts to knowing p, then the warrant rule becomes:

(The knowledge rule) One must: assert p only if one knows p.

Since knowledge entails truth, one satisfies the truth norm whenever one satisfies the knowledge norm. Since one knows that p only if one has *good evidence* for *believing* the *truth* that p, the knowledge norm also explains what is right about Grice's two maxims of quality: "Do not say what you believe to be false" and "Do

⁸⁰Perhaps, the game children play called "Simon says" is an example of a game constituted by one single rule - one must: do as Simon says. See Maitra and Weatherston [2010] for a criticism of Williamson's claim that the unique rule that uniquely constitutes assertion is the knowledge rule.

not say that for which you lack adequate evidence,” as well as the supermaxim “Try to make your contribution one that is true.”⁸¹

The idea is that only knowledge *warrants* assertion. “Warrant” here is understood as a term of art “or that evidential property (if any) which plays the role of property C in the correct simple account of assertion.”⁸² This is not the same sense of “warrant” as in everyday English. Sometimes we say that an assertion is warranted when we mean to be saying that the assertion is “reasonable.”

(As we suggested above, this same model can be applied with great explanatory success to the case of action and inference. The unifying methodological idea behind all those knowledge norms is that knowledge (or lack thereof) is the simplest explanation for the appropriate/inappropriateness of the relevant state or act.)

One potential problem for the truth rule is that assertion is not the only act to aim at truth. This is a potential problem for the truth account, for it claims that assertion is the unique speech act A whose unique rule is “Perform A with the content p only if p is true.” The speech acts of *conjecturing* and of *swearing* also aim at truth. One may conjecture that p without asserting that p - even though conjecturing what is true is better than conjecturing what is false. Asserting does not entail swearing, but the converse is likely to be true. Maybe swearing is a solemn way of asserting.

But Williamson’s main argument against the truth rule is a *reductio ad absurdum* of the view.

1. One must: assert p only if p is true. [assume for *reductio ad absurdum*]
2. One ought to have evidence for one’s assertions *because* they ought to be true (i.e., the evidential norm is derivative from the truth rule) [assumption]

⁸¹cf. Grice [1991].

⁸²Williamson [2000, p.243].

According to Williamson, the philosopher proposing the truth rule appeals to the following sub-argument in favor of 2:

3. One must assert something only if it is true. [assumption]
4. If one must (ϕ only if p is true), then one should (ϕ only if one has evidence that p is true). [assumption]
5. Thus, one should (assert only if one has evidence that p is true) [from 3,4]

Williamson proposes that a contradiction can be derived from a case involving a lottery.⁸³ Suppose that S bought a ticket in a large and fair lottery. Only one ticket wins, the draw has already been held, but neither S nor I know the result. On the basis of the mere probabilistic grounds that S's ticket is the only ticket among many, I assert flat-out, to S, "Your ticket did not win," without giving her my grounds for asserting that. Even if it is true that S did not win the lottery, S will still be entitled to feel some resentment when she discovers that I based my assertion on merely probabilistic grounds. I represented myself as having the authority to make that flat-out assertion. However, I lacked that authority. What I did was akin to cheating.

The infelicity in my assertion "Your ticket did not win" comes from the fact that I have exceeded my evidential authority, but the truth norm cannot explain that fact since "Your ticket did not win" is true, but asserting it infelicitous. The upshot is that one ought to have evidence for what one asserts independently of whether what one asserts is true or not. That is, if the lottery example is sound, we can infer the negation of premise 2 above:

6. It is not the case that one ought to have evidence for one's assertions *because* they ought to be true.

⁸³Williamson [2000, p.246].

6 and 2 form a contradiction, however, and we are entitled to deny 1. The philosopher proposing the truth rule might reply in at least two different ways.

First, she might appeal to Grice's maxim of quantity ("Do not make your contribution more informative than is required") and say that asserting "Your ticket is a loser" violated this maxim and generates the false implicature that I have privileged information about the lottery result (maybe because I saw the draw or because I know the lottery was fixed). The problem with this reply is that it ignores the fact that conversational implicatures are cancelable, and I do not seem to have more evidential authority to assert "Your ticket did not win, but I do not mean to imply that I have inside information" than I have to assert the plain "Your ticket did not win."

Secondly, the proponent of the truth rule might suggest that the argument appealing to the lottery case cannot be generalized. She might point out that for each ticket in the lottery, I have a similar basis for asserting that it did not win. And that if I make all those assertions, I will have asserted something false. The problem with this objection is that it shows only that the lottery example, if generalized, commits one to asserting something *false*; the objection does not show that one has *adequate evidence* to assert the conjunction of all claims of the form "Ticket x is a loser".

Thirdly, one might argue in defense of the truth rule that, if I lacked warrant to assert "Your ticket did not win," then I would lack warrant to assert mostly any other proposition, since few of them are certain for us. This objection assumes that there is no way to discriminate between asserting "Your ticket did not win," and most ordinary propositions except by an appeal to the truth rule. That assumption is not clearly true.

Again, the general structure of this argument for the knowledge norm of assertion can be used to model a knowledge norm of action and inference. If I have a ticket in a large and fair lottery that pays ten million dollars to the person

holding the winning ticket, then it is wrong (in the practical sense of the word) to act on the basis of the following argument:

1. I will lose the lottery.
2. If I keep the ticket I get nothing.
3. If sell the ticket, I will get a cent.

Thus,

4. I ought to sell the ticket.

The reason it is wrong to act in this way is because I do not know I will lose the lottery.⁸⁴

Similarly, if I infer from 1-3 that I can at most make one cent out of my ticket, I will have done something wrong from the *epistemic point of view*, for I do not know I lost the lottery.

Williamson generalizes the lottery argument and argues that one may lack the authority to assert that p , for any p , if p is highly probable on one's evidence but less than 1. Let us consider his example.

Consider a fair lottery with one million tickets. Let p be a proposition whose truth value is known to an expert but about which S has no evidence. S has one ticket in this lottery. The expert does not announce the number of the winning ticket; instead, she hands each participant a piece of paper. If one's ticket is the winner, the true member of the pair $\{p, \neg p\}$ is written on one's piece of paper. If one's ticket is lost, then the *false* member of the pair is written on one's paper. Everyone knows the way the lottery works and no one is in a position to talk to anyone else about the results.

⁸⁴Hawthorne and Stanley [2008, p.572].

Suppose S gets a piece of paper with “ $\neg p$ ” written on it. Given S’s evidence, there is a probability of $1/1,000.000$ that her ticket won and $\neg p$ is true, and a probability of $999,999/1,000.000$ that S’s ticket lost and $\neg p$ is false. If S asserts p , the probability on S’s evidence that her assertion is true is $999,999/1,000.000$. Williamson argues that this is not enough for S to have a warrant to assert p outright, however.⁸⁵ As in the original lottery case, one is not entitled to assert that one’s ticket did not win in the generalized case because one’s sole evidence is that one’s ticket is one in a million. The fact that “ $\neg p$ ” is written on one’s piece of paper tells one nothing about the truth of that claim, for one has no independent evidence for or against it. The result seems to be that one is not entitled to assert that p , even though there is only a one-in-a-million chance that this assertion is false.

Moreover, since the example applies to any proposition we happen to have little or no evidence for, high probability does not entail assertibility. Some propositions, like “I exist”, are not covered by this since we have independent non-probabilistic grounds for them. Thus, one seems never to have a warrant to assert anything on the basis of its probability alone, if this probability is not 1.

This argument can also be applied to the action case and to the inference case. Assume the same set up of Williamson’s generalized lottery. Further suppose that one ought to ϕ only if $\neg p$ is true and that q is true only if $\neg p$ is also true. Now, given the knowledge norm of action and the knowledge norm of inference, if one is holding the piece of paper with $\neg p$ written on it, then one should not ϕ and one should not infer that q . This would suggest that one should act and infer only on the basis of what has evidential probability 1 for one.

At this point, one might wonder: is the requirement that we assert only what has probability 1 on our evidence too demanding? All our assertions about empirical matters seem to have probability less than 1. That depends on what

⁸⁵Williamson [2000, p.250].

we take “probability” to mean here.

If what I mean by “probability” is “objective probability,” then this worry would apply only to our knowledge about the future, since propositions about the present and the past have either objective probability 0 or 1. This is not the sense of “probability” Williamson has in mind, though. According to him, this sense of “probability” is too objective to warrant assertion: “of two past tense assertions whose objective probability is 1, I may have excellent evidence for one and none for the other.”⁸⁶

Williamson also does not mean to use “probability” in the sense of “degree of belief” or “subjective probability”. He says this is “too subjective” to warrant assertion. Intuitively, a proposition one has no reason to believe does not become assertible for one because one is subjectively certain that it is true (for example, “I am Napoleon” does not become assertible for me, because I am deranged and, thereby, psychologically certain that it is true).

The notion of probability that Williamson thinks is relevant to assertion (and to belief) is the notion of probability *on one’s evidence*, or the notion of *epistemic probability*. So, the requirement that the proposition one asserts has to have probability 1 amounts to the requirement that the proposition one asserts has to have probability 1 on one’s evidence. Thus, since, as we saw above, one’s evidence is all and only what one knows, according to this view of assertion, one has a warrant to assert that p only if p has probability 1 on one’s evidence, and p has probability 1 on one’s evidence if and only if one knows that p.

Besides the lottery examples, Williamson argues that conversational patterns support the knowledge norm of assertion.⁸⁷

⁸⁶Williamson [2000, p.251].

⁸⁷Williamson’s discussion here matches, to a good extent, Unger’s discussion in Unger [1975, ch.7].

First, it is noted that “How do you know?” is usually an appropriate response to assertions and that this appropriateness is explained by the fact that the speaker knows what she asserts. What is more, the absence of an answer to this question suggests that the speaker does not have a warrant to assert. This should not be the case if something else besides knowledge conferred warrant on one’s assertion.

Again, as we saw with the cases involving lotteries and with the Gettier case, if knowledge were not the norm of inference, it would be hard to explain the propriety of our criticism when the reasoner fails to know at least one of the premises on which her conclusion relies essentially.

Second, the knowledge norm can explain why it is infelicitous for someone to assert Moorean sentences of the form “p and I do not know that p” and “p and I do not believe that p.”⁸⁸ One cannot properly assert the propositions expressed by those sentences because one cannot know those propositions and only knowledge warrants assertion. It is easy to show that one cannot know the proposition those sentences express. Consider “p and I do not know that p” first.

1. $K(p \& \neg Kp)$ [assume for a *reductio ad absurdum*]
2. $Kp \& K\neg Kp$ (1, by distribution of knowledge over conjunction)⁸⁹
3. $\neg Kp$ (2, by conjunction elimination and K-factivity)
4. Kp [2, by conjunction elimination]
5. $\neg K(p \& \neg Kp)$ [1, 2-4, by *reductio ad absurdum*]

A similar argument shows that one cannot know “p and I do not believe that p”.

⁸⁸Moore [1993].

⁸⁹Most everyone accepts the validity of $K(p \& q) \vdash Kp \& Kq$. But see Nozick [1981] for a dissenting view on the matter.

1. $K(p \& \neg Bp)$ [assume for a *reductio ad absurdum*]
2. $Kp \& K\neg Bp$ [1, by distribution of knowledge over conjunction]
3. $\neg Bp$ [2, by conjunction elimination and the factivity of knowledge]
4. Bp [2, by conjunction elimination and the fact that knowing entails believing]
5. $\neg K(p \& \neg Kp)$ [1, 2-4, by *reductio ad absurdum*]

One might worry that this proves too much. What about “p and I cannot be (am not) certain that p”? Does the fact that it is always inappropriate to assert this suggest that the norm of assertion is certainty and not knowledge? Williamson suggests that if we allow the contextually set standards for knowledge and certainty to diverge, the worry disappears since we let assertibility be determined by knowledge and not by certainty.⁹⁰

Perhaps there is a more pressing worry here. It has been suggested that besides explaining why it is inappropriate to assert the *omissive* form of paradoxical Moorean sentences, one’s theory also has to explain why it seems inappropriate to assert the *commissive* form of these sentences such as “p and I know that $\neg p$ ” and “p and I believe that $\neg p$ ”.⁹¹ It is clear how the knowledge norm should deal with “p and I know that $\neg p$ ”:

1. $K(p \& K\neg p)$ (assume for a *reductio ad absurdum*)
2. $Kp \& KK\neg p$ (1, by distribution of knowledge over conjunction)
3. $\neg p$ (2, by conjunction elimination and the factivity of knowledge)

⁹⁰Notice that the Truth norm of assertion cannot explain this data by itself.

⁹¹As far as I know Roy Sorensen was the first to press this point in Sorensen [1998, ch.7]. For a more recent adoption of this point see the collection of papers on the Moorean absurdity in Williams and Green [2007], specially Claudio de Almeida’s contribution to that volume, de Almeida [2007].

4. p (2, by conjunction elimination and the factivity of knowledge)
5. $\neg K(p \& K\neg p)$ (1, 2-4, by *reductio ad absurdum*)

What is not as clear is how the knowledge rule should deal with “ p and I believe that $\neg p$.” From the fact that I know that I believe that $\neg p$ it does not follow that $\neg p$. This prevents us from deriving a contradiction in the same straightforward way we have been able to derive one from the other sentences. One might try to derive a contradiction from 3 and 4 with the help of the suggestion that no single agent can believe both p and $\neg p$:

1. $K(p \& B\neg p)$ (assume for a *reductio ad absurdum*)
2. $Kp \& KB\neg p$ (1, by distribution of knowledge over conjunction)
3. $B\neg p$ (2, by conjunction elimination and the factivity of knowledge)
4. Bp (2, by conjunction elimination and the factivity of knowledge)

If that were the case, then 3 and 4 describe an impossible state of affairs and we should deny 1.⁹² The problem with this suggestion is that it is far from clear that 3 and 4 do describe something that is impossible; that is, it is far from clear that agents cannot believe p and $\neg p$ simultaneously.

In contrast with philosophers who think one cannot believe both p and $\neg p$, *gullibilists*⁹³ defend that, given the right circumstances, anyone would believe anything. In fact one could hold either a strong or a weak version of gullibilism.⁹⁴ According to the strong version of the view,

(SG) $(\forall p)(\forall x) \diamond Bxp$.

⁹²Jaakko Hintikka defended a similar claim in this context in Hintikka [1962].

⁹³Fain and Griffiths [1972].

⁹⁴cf.Sorensen [1998].

Strong Gullibilism says that, no matter what proposition or what agent we care to consider, it is possible that this agent believes that proposition. According to the weak form of gullibilism, on the other hand,

(WG) $(\forall p)(\exists x) \diamond Bxp$

Weak Gullibilism says that no matter what proposition we consider, there is an agent such that it is possible for this agent to believe that proposition.

I take it to be obvious that *some* version of the gullibilist thesis is true. Maybe weak gullibilism as stated by WG is just that version. However, the plausibility of the gullibilist thesis is, to a great extent, a function of it being a *psychological* claim, rather than a *normative* claim. That is, WG's plausibility stems from the fact that there seems to be no necessary psychological fact that would exclude the possibility that a proposition is always a potential the object of belief, no matter how absurd. After all, this world has a huge collection of fools, irrational and crazy agents (not to mention all possible worlds!). The same is not so obviously true with respect to the normativity of belief. Quite the opposite, there seems to be many propositions that are such that no one could possibly believe them rationally: for example, no one could rationally believe "I do not have any beliefs whatsoever" or "I do not exist." Thus, there seem to be some normative fact or facts that prevent any agent from believing rationally *some* propositions. If this is right and irrationality is incompatible with knowledge, then commissive cases of Moorean sentences are unknowable and, thus, only infelicitously asserted.

Rather than insisting that the commissive cases of Moore's paradox should be taken to be only irrationally believed, I would like to make the case that, first, it is not so clear that the commissive case is as paradoxical as the omissive case and, second, even if the commissive case is just as contradictory sounding as the omissive case surely is, other accounts of assertion do not fare better than the knowledge account does in accommodating the paradox-sounding nature of

commissive Moorean sentences.

As some evidence that the commissive case is not as paradoxical as the omissive one consider the following example modified from de Almeida [2001]:

The Moorean Couple Case

John and Mary are very good friends who know each other for over thirty years. They enjoy the same things and, in particular, they enjoy discussing their personal lives with each other. One day, after telling John about her brother-in-law and the weekend they spent together with their respective partners in a cabin in the woods, John starts suspecting that Mary is in love with her brother-in-law, but that she doesn't believe she is. John thinks Mary is in love with her brother-in-law because she is really excited about everything he does and it is always very happy when she talks about the time they spend together. Plus, John knows Mary well and he has seen the same signs of love when she was getting to know her current husband. John thinks Mary doesn't believe she loves her brother-in-law because she sincerely tells him she loves only her husband. A few weeks after their conversation about Mary's weekend John decides to confront Mary with her own feelings for her brother-in-law. John explains to her how she seems excited every time she talks about her brother-in-law, while insisting that she loves only her husband.

Now, suppose that, after listening attentively to John's heartfelt and honest description, Mary is surprised with herself. After thinking for a few seconds, she turns to John and says "Yes, I love my brother-in-law, but I believe I don't." What Mary said seems plausible to me. There does not seem to be anything wrong with what she said - at least not to my ears. But Mary has just asserted the commissive form of Moore's paradox and, if Sorensen and others are right, what Mary said should have sounded absurd. What is more, compare the seemingly appropriateness of Mary's assertion with the clear inappropriateness of the related omissive form of what Mary asserted: "Yes, I love my brother-in-law, but I don't believe I do." I believe that cases such as this should give us pause when considering Sorensen's claim that commissive and omissive cases are equally absurd.

Chapter 4

KFK and the Gettier Problem

At the heart of contemporary philosophy's understanding of what it means for someone to know something via inference are the so-called "Gettier cases."¹ Consider one such case, Edmund Gettier's Case I.²

The Coin Case

Smith has a justified true belief that the man who will get the job has ten coins in his pocket. He infers this proposition from his justified but false belief that Jones will get the job and Jones has ten coins in his pocket.

Smith does not know that the man who will get the job has ten coins in his pocket. That he does not know that proposition is something virtually all epistemologists accept. The agreement ends here, however. Since Gettier published his paper a whole industry of theories that claim to "solve" the "Gettier Problem" has dominated most of the epistemological literature. In this chapter I defend the view that the principle according to which inferential knowledge requires knowledge of all the premises essentially involved in one's reasoning (i.e., KFK) solves the Gettier Problem.

¹The view that Gettier cases are pivotal to inferential knowledge is advocated, most prominently, by Gilbert Harman in Harman [1973] and Harman [1980].

²Gettier [1963].

4.1 The Gettier Conjecture

Gettier uncovered an important epistemological phenomenon: sometimes deductively valid arguments will not give us inferential knowledge of their true and justified conclusion, even though we seem to be justified in believing all the premises in those arguments. The lesson many drew from Gettier's paper was one about conceptual analysis, however - namely, that "knowledge" was not coextensive with "justified true belief." This reaction to Gettier's paper seemed appropriate in part because Gettier himself was ostensibly using his cases as counterexamples to the claim equating both concepts. In what follows my focus will be on understanding the phenomenon of inferential knowledge uncovered by Gettier rather than on the conceptual problem raised by his cases. In particular, I want to argue for one main claim. I will argue that the explanation as to why protagonists in Gettier cases fail to know is the fact that their reasoning depends essentially on something they do not know. I call this the "Gettier Conjecture". This explanation follows from the principle KFK I introduced in chapter 2. Therefore, it gives us a straightforward explanation of an important epistemological phenomenon. The fact that this conjecture is entailed by KFK helps to put the conjecture in a wider epistemological context.

In order to get a better grasp at the Gettier Conjecture, consider Gettier's Case I. Smith has a justified true belief that the man who will get the job has ten coins in his pocket. He infers this proposition from his justified but false belief that Jones will get the job and Jones has ten coins in his pocket. Smith does not know that the man who will get the job has ten coins in his pocket. But why? If the Gettier Conjecture is true, then a simple explanation is readily available: Smith fails to know the conclusion of his inference because his belief in the conclusion depends essentially on a false (and, thus, unknown) belief in one of the premises. A similar explanation applies to Gettier's Case II.

The Ford Case

Smith infers that either Jones owns a Ford or Brown is in Barcelona, which is true, from his justified but false belief that Jones owns a Ford. Nonetheless, he does not know the conclusion of his inference.

Again, if our conjecture is true, then the reason why Smith does not know is because he does not know the premise on which his belief in the conclusion depends essentially.

Before we go any further, let me state our conjecture more explicitly:

(GC) Necessarily, the protagonist of a Gettier case fails to know that p because her inference in support of p depends essentially on at least one premise she does not know.

The first thing to note is that GC is what we get if we apply KFK to Gettier cases. Notice, too, that the “because” clause in GC denotes the relation of explanation, rather than the relation of causation. Causation and explanation might be intimately related, but GC is intended to be silent on this relationship.

Second, notice that whether GC succeeds in explaining the ignorance that is present in *all* Gettier cases depends in part on what counts as a Gettier case. GC is ostensibly an explanation of the ignorance we find in Gettier cases where the protagonist acquires a justified true belief via reasoning. Hence, GC covers the original Gettier cases, since those are cases where the target belief is an inferential one. But, does it cover all the *other* Gettier cases? I will try to answer this question in the next section.

Before I finish this section let me point out that at least one other philosopher seems to have held something similar to GC. In his 1973 book “Belief, Truth and Knowledge” David Armstrong says the following about the Gettier cases:

Gettier produces counterexamples to the thesis that justified true belief is knowledge by producing true beliefs based on justifiably believed grounds,

... but where these grounds are in fact false. But because possession of such grounds could not constitute possession of knowledge, I should have thought it obvious that they are too weak to serve as suitable grounds.³

Armstrong not only seems to accept GC in this passage, but the reason why he accepts it is because he accepts something like the general principle about inferential knowledge I call KFK. His acceptance of KFK also echoes Russell's endorsement of this principle in the beginning of the twentieth century.⁴

4.2 What counts as a Gettier case?

This is a vexing question and epistemologists usually disagree on how to answer it.⁵ Even though I do not expect to close this question definitively, I do hope the argument in this section will make progress on the issue.

First of all, it might be thought that not all Gettier cases involve inference. For example, Alvin Goldman's Barn-façade case⁶ is sometimes taken to be a Gettier case in which the protagonist acquires a justified true belief via perception. Russell's stopped-clock case⁷ and Roderick Chisholm's sheep-in-the-field case⁸ are also sometimes taken to be Gettier cases whose protagonists acquire a justified true belief via perception. If these are Gettier cases in which there is no reasoning involved in the acquisition of the protagonist's justified true belief, then the appeal

³Armstrong [1973, p.152].

⁴See ch.2.

⁵cf. Lycan [2006].

⁶cf. Goldman [1976, p.772-3].

⁷Russell originally used the example to show that there can be true belief without knowledge in Russell [1948/2009, p.140]. Israel Scheffler, in Scheffler [1965, p.112], was the first to suggest that the case could be turned into a counterexample to the JTB analysis of "knowledge" if we assume that the subject looking at the clock "has good grounds to suppose the clock *is* going." This formulation of Russell's case turns it into a case in which the protagonist's justified true belief that it is 10:00AM depends essentially on the falsehood "the clock *is* going." I come back to issue of whether or not this is a non-inferential case below.

⁸cf. Chisholm [1977].

of GC might be taken to be somewhat diminished. We do not have to necessarily accept this result, however. There are a few things we can say about these cases that are in line with GC explaining ignorance in *all* Gettier cases.

For example, one could hold the view that only the two original cases presented by Gettier deserve to be called “Gettier cases.” One might think this because one thinks that the point of a Gettier case is to show that “knowledge” is not the same as “justified true belief,” and that this point can be made with one single Gettier case. The problem with this view of what counts as a Gettier case is that it is far from clear that we cannot multiply Gettier cases by discerning the essential features of the original cases. For instance, we might note that Gettier accepted two principles (fallibility and justification closure) and that, *mutatis mutandis*, any case in which the protagonist instantiates these principles and, as a result, arrives at a justified true belief that is not knowledge should also be deemed “Gettier case.” Here is how Gettier presents those principles about justification:

I shall begin by noting two points. First, in that sense of ‘justified’ in which S’s being justified in believing P is a necessary condition of S’s knowing that P, it is possible for a person to be justified in believing a proposition that is in fact false. Secondly, for any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q.⁹

If we go down that route - that is, if we accept Gettier’s suggestion that he is interested only in cases in which both the fallibility principle and the closure principle are instantiated - then it follows that all Gettier cases are cases in which the protagonist acquires a justified true belief via deductive reasoning which relies essentially on a false premise. At this point someone might complain that even though instantiating fallibility and justification closure is a *sufficient* condition for a case to be a Gettier case, it is not a *necessary* one. The idea would be that cases that do not instantiate either one of those principles may also be deemed

⁹Gettier [1963, p.121].

Gettier cases. There are three mutually inconsistent and exhaustive options here, all of which take the original Gettier cases to be Gettier cases, of course, but they also suggest that we expand the class of those cases.

The first alternative is to say that instantiating fallibility and justification closure is *sufficient but not necessary* for a case G to be a Gettier case. The proponent of this view might point out that Chisholm's sheep-in-the-field case,¹⁰ Feldman's clever reasoner case,¹¹ as well as Goldman's Barn-façade case¹² are examples of Gettier cases that do not instantiate either one of those principles. There are a few problems with this alternative. First, it is not clear that the protagonist in all of those cases fails to know. Secondly, even if the protagonist in all of those cases fails to know, a case may be made that the protagonist in all of those cases relies on unknown premises essentially. Thirdly, this alternative runs the risk of over-generating Gettier cases. I will briefly discuss each of these issues in turn.

Ernest Sosa has famously argued that there is an important sense in which the protagonist of the Barn-façade case knows that that is a barn - the "animal knowledge" sense of "knowledge" -, while there is another sense in which he fails to know - the "reflective knowledge" sense of "knowledge."¹³ William Lycan also reports not having the intuition that the person in this case fails to know.¹⁴

¹⁰Chisholm [1977].

¹¹Feldman [1974].

¹²Goldman [1976].

¹³cf.Sosa [2007, p.96 fn.1] and Sosa [2011, p.82-95]. In Sosa [2011, p.92], Sosa tries to accommodate the fact that "many of us cannot believe that the fake barns subject knows at any level whatsoever, whether animal or reflective" by distinguishing still a further sense in which one can be truly said to know something - i.e., the "human knowledge" sense of "knowledge." Human knowledge, according to Sosa, comes in degrees and in its lowest degrees it corresponds to reflective knowledge; in its highest degree, human knowledge involves "scientific and even philosophical perspectives that enable defense of one's first-order belief as apt." Sosa argues that those who think that Henry doesn't know, full stop, have the "human knowledge" sense of "knowledge" in mind.

¹⁴cf.Lycan [2006, p.158,p.162-3].

Finally, Tamar Gendler and John Hawthorne have challenged the very reliability of intuitive judgments in response to fake barn cases.¹⁵ The point here is that there is wide disagreement as to whether the protagonist in Goldman's case in fact fails to know. This result contrasts sharply with the wide agreement among epistemologists that the protagonist in the original Gettier cases fails to know.

Even if the protagonists in all of those cases fail to know, it can be argued that they rely essentially on propositions they do not know. Someone will certainly complain: "Wouldn't that turn those cases into *inferential* Gettier cases?" Maybe. Whether or not a belief depending essentially on a proposition makes that belief inferential depends in part on how one thinks about inference. If only explicitly inferential beliefs count as being the product of inference (as when one says to oneself "A; therefore, B"), then the target belief in those cases is arrived at non-inferentially. But, if a belief depending essentially on a proposition in the sense sketched above counts as "inferential" in a broad sense of that word (as I think it is plausible to think it does), then the justified true belief in those cases *are* the result of inference and, thus, fall under the purview of GC. Once we accept this view of inference, it is easy to see that the target belief in those cases depends on at least one other proposition they fail to know: "The clock is working" in Russell's case; "That is a sheep" in Chisholm's case; "Nogot owns a Ford" in Feldman's case; "There are no barn look-a-likes around" in Goldman's case. Their true beliefs depend evidentially and causally on those beliefs they fail to know. Given GC, that explains why they fail to know.

The alternative that we are considering also runs the risk of over-generating Gettier cases. If we thought that we could have Gettier cases that do not instantiate either justification closure or fallibility but in which the protagonist has a justified true belief without knowledge, then it would be hard to see how we

¹⁵Gendler and Hawthorne [2005].

would rule out the possibility that, say, versions of the lottery case count as Gettier cases. After all, the protagonist of a lottery case arrives at a justified true belief that his ticket is a loser via inference from his knowledge that the odds that his ticket is a winner are slim and his knowledge that the lottery is fair. In cases like that the protagonist has an inferentially justified true belief which depends on things he knows; he fails to know the conclusion of his inference, but, still, this is clearly not a Gettier case. The issue here is not that other views of what a Gettier case is cannot explain *why I fail to know* in the lottery case, but, rather, that they cannot explain *why the lottery case is not a Gettier case*.

The initially plausible suggestion that instantiating fallibility and justification closure is sufficient but not necessary for a case to be a Gettier case turned out to be problematic. A second alternative would be to say that instantiating fallibility is necessary, while instantiating justification closure is not necessary for a case to be a Gettier case. This alternative is compatible with GC explaining why the protagonist of a Gettier case fails to know: GC applies to cases in which one's conclusion depends essentially on a falsehood, but it is silent about whether the case involves deduction or not. If, given the fallibility principle Gettier cases are such that their protagonist's conclusion depends essentially on a falsehood, then GC applies to all Gettier cases.

The third and last alternative says that the instantiation of justification closure, but not of fallibility, is necessary for a case to be a Gettier case. Perhaps someone would defend this alternative using Feldman's clever reasoner case whose protagonist allegedly knows all the premises essentially involved in his reasoning. However, as I said before, I doubt that the protagonist in that case does not rely essentially on an unknown proposition.¹⁶

We can also try to explain why some epistemologists might confidently insist that there are Gettier cases that do not fit GC's bill. One plausible explanation

¹⁶Shope [1983, p.211-6] detailed discussion of this case makes the same point.

for this inclination is that most of those epistemologists lose sight of the fact that Gettier cases are merely *one* class of cases in which the protagonist has a justified true belief but no knowledge. The suggestion is that these epistemologists tacitly assume that *all* cases of justified true belief with no knowledge are Gettier cases. Thinking that that is the case, whenever they stumble upon a case of justified true belief with no knowledge they think they have stumbled upon a Gettier case. This would explain their inclination to resist the suggestion that all Gettier cases are cases in which justification closure and fallibility are instantiated.

The problem with this train of thought is that many cases of justified true belief are *clearly* not Gettier cases. Here are at least two classes of cases that are not Gettier cases and in which the protagonist has a justified true belief but not knowledge - lottery cases and cases involving sorites series. I briefly considered lottery cases, so I will not repeat that point here. As for cases involving sorites series, consider a series of tiles going from clearly red to clearly yellow. In that case I know that the first tile is red and that the last one is yellow. It is plausible to think, however, that at some point in the series I will have a justified true belief that the tile in front of me is red/yellow, but I will not know that it is. This is a clear cases of justified true belief without knowledge, but it is far from being a Gettier case.

The upshot of this brief discussion is that, even if not perfect, the idea that a case G is a Gettier case only if *both* fallibility *and* justification closure are instantiated in G seems to be the least problematic suggestion about what makes a Gettier case a Gettier case.¹⁷ The good news is that, if fallibility is an essential component of a Gettier case, then GC explains why there is ignorance in *all* Gettier case.

¹⁷cf. Shope [1983] for a similar characterization of Gettier cases.

4.3 Two Gettier Problems

Gettier cases are said to give rise to the so-called “Gettier Problem.” As usually understood, “Gettier Problem” refers to a problem in conceptual analysis. Gettier (and virtually everyone since his paper was published) took his cases to be a refutation of the claim that the concepts *knowledge* and *justified true belief* are necessarily co-extensive. In the two cases discussed by Gettier, the protagonist instantiates the concept *justified true belief*, but, it is argued, not the concept *knowledge*. Call this worry raised Gettier cases the *Conceptual Gettier Problem*, for obvious reasons.

Some epistemologists trying to solve the Conceptual Gettier Problem have argued that a fourth condition must be added to the traditional analysis.¹⁸ Others have proposed analyses of “knowledge” that reject the traditionalist claim that justification is a necessary condition on “knowledge” and put forward a counterfactual condition instead.¹⁹

If Williamson is right and “knowledge” is unanalyzable, then there is a sense in which trying to meet the Conceptual Gettier Problem is a waste of time. On the other hand, even if he is right, it does not follow that Gettier cases are devoid of philosophical interest. In particular, Gettier cases raise what may be called the *Explanatory Gettier Problem*: what exactly explains why protagonists in Gettier cases fail to know? According to GC, the protagonists of Gettier cases fail to know because their belief relies essentially on at least one premise they do not know.

¹⁸For example, Klein [1971], Chisholm [1977], Goldman [1986], Lehrer [1990], Feldman [2002], and Pritchard [2012].

¹⁹For example, Nozick [1981], Dretske [1971].

Knowledge first epistemology has so far said little about the Conceptual Gettier Problem and nothing about the Explanatory Gettier Problem.²⁰ Knowledge-firsters can appeal to GC in order to solve the Explanatory Gettier Problem, for GC is a special case of KFK and the latter is entailed by Williamson's account of evidence. This is good news for knowledge first epistemology, for, GC, if true, provides knowledge-firsters with a solution to this important epistemological problem.

In sum, there seems to be two different kinds of problems raised by Gettier cases: a conceptual one and an explanatory one. Even though these two problems are related, epistemologists have focused most of their efforts on solving the conceptual problem.²¹ Since the main motivation for trying to solve this problem was the assumption that "knowledge" can be analyzed, the appeal of Gettier cases has diminished as doubts about whether this concept is really analyzable grew stronger. If I am right and Gettier cases really raise an explanatory problem, then those cases are epistemologically relevant even if "knowledge" is unanalyzable.

4.4 Some Caveats

I had better say a few words about what GC *does not* claim in order to preempt possible confusion. For instance, GC will be considered ludicrous if we fail to realize that it assumes that certain conditions are in place. Once those conditions are in place (and only then) GC applies. For example, we have no need to reach for GC as an explanation for why S fails to acquire inferential knowledge if S's reasoning is an instance of an invalid form of argument such as affirming the

²⁰But see Williamson [2013].

²¹Of course, nothing I said here should suggest that traditional epistemologists trying to meet the conceptual challenge have nothing to say about the explanatory problem. For example, if one thinks that "knowledge" can be factorized into "justification," "truth," "belief" and "x," then one might insist that what explains why protagonists of Gettier cases fail to know is that her belief is not x.

consequent or if he based his reasoning on a weak inductive argument. Since those are clear vices in reasoning, the fact that S relies on a fallacious form of argument or on a weak inductive argument explains, by itself, his lack of knowledge. What is more, I take it that this is how things should be, for we ought to apply GC only to Gettier cases, and neither of those defects in reasoning, by themselves, turn a case into a Gettier case.²² Likewise, even if the agent uses a deductively valid argument or a strong inductive argument (and even if she knows all the premises essentially involved in her reasoning), but she is drunk or otherwise “in a bad cognitive shape,” we have no need to appeal to GC in order to explain why she fails to acquire inferential knowledge. GC applies only to cases in which subjects are employing laudable forms of reasoning and are in good enough cognitive condition. If we forget about this caveat, we are liable to dismiss GC too quickly. Again, this is plausible because epistemologists do not tend to see cases in which the protagonist’s reasoning is defective in the ways suggested, or cases in which the protagonist is in bad cognitive shape as Gettier cases.

There is also the question of division of explanatory labor. When well understood, GC is an answer to the following question:

(Q1) Why do protagonists in Gettier cases fail to know?

The answer is that the reasoning they use to arrive at their justified true belief involves essentially at least one premise they do not know. Once that is settled, one could ask the further question:

(Q2) Why is it that reasoning which involves essentially propositions one does not know does not yield knowledge of its conclusion?

²²In fact, in those cases S may even know all the premises involved in his reasoning and fail to know the conclusion because his reasoning is fallacious or inductively weak.

GC, by itself, is not an answer to Q2. Given GC, however, we may reasonably expect that an appropriate answer to Q2 might come from at least two different places. Part of an answer to Q2 might come from our practice of evaluating the epistemic standing of our own inferences and the inferences of others, while another part of that answer might come from epistemological theories about what counts as evidence and what, therefore, may be used, with epistemic propriety, as a premise in reasoning.

If we are inclined to deem otherwise sound inferences defective because the subjects undergoing those inferences rely essentially on propositions they do not know, then that would indicate that we normally view knowledge of premises essentially involved in inference as necessary for knowledge of that inference's conclusion. If propositions one does not know are not part of one's evidence set, then that would explain why having one's reasoning essentially involve premises we do not know prevents one from knowing the conclusion of one's reasoning, for only what is part of one's evidence may be properly used as a reason (i.e., premise) in reasoning.

As we saw in chapter 3, we have reason to think that this kind of answer to Q2 might be made plausible: it seems that we *do*, in general, deny knowledge to subjects whose reasoning relies essentially on a premise they do not know, and, according to knowledge-first epistemologists, our evidence includes only propositions we know. Arguably, GC is made more plausible by the fact that it generates non-obvious expectations we have a reason to think are in fact satisfied.

4.5 A recent ancestor of the Gettier Conjecture

The “No-False-Grounds” solution to the Conceptual Gettier Problem is a recent ancestor of GC.

In an early reply to Gettier's paper, Michael Clark²³ proposed that the protagonists in Gettier cases fail to know because their true belief depends on at least one false belief. Alvin Goldman's account of inferential knowledge in Goldman [1967] also required that all beliefs involved in an inference be true if they are to yield knowledge of the reasoning's conclusion.²⁴ Gilbert Harman also defended a version of this approach to the conceptual problem raised by Gettier cases.²⁵ Recently, Richard Feldman,²⁶ William Lycan,²⁷ and Richard Foley²⁸ have also endorsed similar views. The unifying idea behind all these *no-false-grounds* (NFG) views is that subjects in Gettier cases fail to know because they rely essentially on at least one false belief.²⁹ Since those philosophers are trying to solve the Conceptual Gettier Problem they took the no-false-grounds condition to be a necessary condition on "knowledge."

Although NFG was originally a solution to the Conceptual Gettier Problem, the view can be easily turned into a proposed solution to the *Explanatory Gettier Problem*: necessarily, the protagonist of a Gettier case fails to know that p because her reasoning in support of p depends essentially on at least one false premise. Once understood in this way, NFG and GC become somewhat similar claims: since knowledge entails true belief, whenever one fails to satisfy NFG, one fails

²³Clark [1963].

²⁴cf., Goldman [1967, p.369-70]. Goldman explicitly takes his account of inference to be an improvement on Clark's view in Goldman [1967, p.364]. According to Goldman, if Clark had augmented his view to include beliefs about causal relations as grounds for inferential knowledge, then Clark's view would have been "almost equivalent" to his view.

²⁵cf.Harman [1973, 1980, 1986].

²⁶Feldman [2002, p.33-7].

²⁷Lycan [2006, p.153-8]. Lycan also endorsed a similar view back in the 1970's in Lycan [1977].

²⁸Foley [2012].

²⁹To be precise, Clark's version of the view suffers from a problem the other views do not. Since Clark did not distinguish between essential and inessential grounds, his view yields the wrong result in cases of *harmless ignorance* we discussed above. Clark is committed to saying that the subject in those cases fails to know because her reasoning depends (even if non-essentially) on a falsehood.

to satisfy GC as well. The difference between GC and NFG, of course, is that the converse of this claim is not true, for it is possible for one to believe truly and fail to know. Supposedly, however, NFG proponents will also require that one be justified in believing all premises essentially involved in one's inference, for one would not be justified in believing the conclusion of one's reasoning if that were not the case. If one were not justified in believing all the premises essentially involved in one's reasoning, NFG would not rule out the clearly absurd possibility of someone acquiring inferential knowledge on the basis of mere *lucky guesses*, as when S believes truly each of the premises essentially involved in his inference because of individual coin tosses. This possibility is absurd, because knowledge cannot be based on mere lucky guesses.³⁰

This shows that the NFG-based solution to the Explanatory Gettier Problem and GC are not so different after all. The former suggests that the protagonist of Gettier cases fails to know the conclusion of her reasoning because it depends essentially on at least one proposition she does not believe truly and justifiably.

³⁰Peter Klein has questioned (in a comment to an earlier draft of this chapter) the claim that justified belief in the premise one's conclusion depends essentially on is necessary for inferential knowledge. According to him, a modified version of a case he discusses in Klein [2008, p.37] is a counterexample to this claim. Here is the alleged counterexample: Weatherman believes that the average annual precipitation in Northwest Montana is about 13 inches because he believes that accurate records have been kept for over eighty years and the rainfall depicted in the number of years that records were kept averages to 13 inches. He also believes that he is very bad at counting the number of years records are kept, but, he has always been in the ballpark. The average rainfall is about 13 inches, but accurate records were kept for only seventy-nine years. According to Klein, the Weatherman knows that the average rainfall is about 13 inches even though this belief depends essentially on a false *and* unjustified premise. This is, of course, another alleged case of knowledge from non-knowledge. I have already said a few things in reaction to those cases in chapter 2 and I mean for the points I make there to carry over to the case with which we are now concerned. I do have to admit that I feel even less inclined to attribute knowledge in this case than in other alleged cases of knowledge from non-knowledge, though, because the protagonist is described as not even justified in believing the premise on which his conclusion relies essentially. On the one hand, I am at a loss if I am asked to argue in support of a claim I see as obvious. On the other hand, I suspect that Klein has in mind here a different sense of "evidence" (and perhaps "knowledge") than I have in mind in this essay. While I have in mind an "objective" sense of "evidence," Klein is presupposing a very subjective sense of "evidence" when discussing this case. Klein's subjective sense of "evidence" is fine and good, but it is not the sense of "evidence" with which I am mainly concerned here.

The latter suggests that the protagonist of Gettier cases fails to know the conclusion of her reasoning because it depends essentially on at least one premise she does not know. Since, plausibly, knowing that p is a way of believing that p justifiably, how different NFG and GC actually are hinges on how one understands “believing that p justifiably”. If, as it has been suggested,³¹ one is justified in believing that p only if one knows that p , then NFG and GC amount to the same thing.³² If, on the other hand, one can be justified in believing that p even if one does not know that p (like Gettier explicitly assumed to be the case), then NFG and GC are distinct accounts of the Explanatory Gettier Problem. The point here is that the difference between the two views, although significant, does not make these views irreconcilable. I discussed the concept “justification” in chapter 3, so I will not try to settle this issues here. I want, however, to show how a familiar variation of the general NFG theme entails KFK and, consequently, GC.

Even though the defeasibility account of “knowledge” (an alleged improvement on NFG and a popular “solution” to the Conceptual Gettier Problem in the 1970’s and 1980’s) entails KFK, this has not (as far as I know) being explicitly pointed out in the literature.

According to the defeasibility analysis of “knowledge”, undefeated justification is a necessary condition on “knowledge”.³³ One’s belief in the conclusion of one’s reasoning is undefeated only if there is no truth d such that the conjunction of d and the premises in one’s reasoning fails to justify one’s belief in the conclusion of that reasoning.³⁴ We can turn the defeasibility theory into a proposed solution

³¹cf. Sutton [2007], Williamson [2011].

³²Given the discussion about justification in chapter 3, if one is w -justified in believing that p , then one knows that p . As I argued there, the Gettiered subject is at most r -justified in believing the conclusion of his argument. One might be r -justified in believing something even if one’s belief falls short of knowledge.

³³cf. Lehrer and Thomas Paxson [1969], Lehrer [1990], Klein [1971, 1981].

³⁴The fully worked out defeasibilist account of “knowledge” has to distinguish *misleading defeaters* from *genuine defeaters*. See Klein [1979] for the distinction.

to the explanatory Gettier Problem. The resulting view entails KFK. Suppose, for a *reductio ad absurdum*, that this is not the case; that is, suppose that, for some S, S knows that p as the result of an inference even though she fails to know one of the premises essentially involved in her reasoning, and that all premises essential to her reasoning are undefeated. But, if S fails to know one of the premises essentially involved in her reasoning, say q, then “S doesn’t know that q” is a defeater of S’s justification for the conclusion, p, of her reasoning. This is a contradiction, since we are supposing that there is no defeater of S’s justification for p and we have uncovered such a defeater. We are then allowed to reject our assumption. It follows that the defeasibility theory entails KFK. Furthermore, since KFK entails NFG, KFK and the defeasibility theory are virtually equivalent accounts of inferential knowledge.

I can also offer a different argument in support of this entailment. If the defeasibility account of “knowledge” is correct, then it seems that there is a genuine defeater of S’s justification whenever S fails to know one of the premises essentially involved in her inference. The reason for that is straightforward: if S is in that situation, then the truth “S does not know premise x of her reasoning” is a defeater of S’s justification, because it entails “Either x is false, not believed, or unjustified”. Whatever makes this disjunction true also prevents S from knowing the conclusion of her reasoning.

The main result of this section is that NFG is not so different from GC after all. Be it because “justification” might be understood as entailing knowledge or because NFG+defeasibility entails KFK.

Chapter 5

Epistemic Luck

As I write this thesis, the most popular solution to the Gettier Problem comes from a view broadly construed as “anti-luck epistemology.” In this chapter, I analyze this solution and conclude that it fares worse than the solution based on KFK I offered in the previous chapter.

5.1 Anti-Luck Epistemology

Duncan Pritchard, the most prominent proponent of anti-luck epistemology, has recently argued that a correct definition of knowledge contains an anti-luck condition. This claim has at least two important pieces of information worth unpacking. First, in spite of admonitions to the contrary,¹ Pritchard thinks there is still hope for what he calls the “analytical project” (i.e., the philosophical task of analyzing the concept “knowledge” in terms of simpler concepts). He makes this point refreshingly explicit in the following passage:²

Until relatively recently, a key task assigned to the epistemologist was to offer an adequate definition of knowledge, one that was informative, non-circular, and which could suitably accommodate our salient epistemological intuitions. Call this the *analytical project*. This project has fallen into disfavor recently, with many arguing that it is a hopeless task. Given the lack of success that epistemologists have had on this score it is not surprising that a disconsolate mood should have set in amongst those working on this project. Nevertheless, such pessimism is premature. Indeed, I will be arguing that there is an adequate theory of knowledge available which fulfills

¹e.g., Williamson [2000].

²Pritchard [2012, p.247].

the remit of the analytical project.

Second, Pritchard argues that the correct definition of “knowledge” states that knowledge is incompatible with the true belief being the result of epistemic luck; he calls this the “anti-luck intuition.” What is more, he says that much of the literature on the Gettier Problem supports this claim:³

Consult any introductory text in the theory of knowledge, and you will find a statement of this intuition. If, for example, a commentator is asked to explain why mere true belief cannot suffice for knowledge, the standard response is to point out that in cases of mere true belief, unlike knowledge, one’s cognitive success can simply be a matter of luck. The role of this intuition in contemporary theory of knowledge is particularly apparent in the post-Gettier literature, where it is often stated that precisely the point of the Gettier cases is that they demonstrate that justified true belief is compatible with one’s cognitive success being merely due to luck. The failure of the tripartite account to accommodate the anti-luck intuition is thus meant to be a decisive strike against it.

In what follows, I will argue that Pritchard’s analytical project fails. The reason for this failure is straightforward: epistemic luck, as characterized by Pritchard, is not a necessary condition on knowledge. Counterexamples to Pritchard’s safety-based account of epistemic luck abound and we will look at one such example below. Unfortunately, the problems with the analytical project do not end there - they run deeper.

Safety is a modal notion and true beliefs that are safe are *modally robust*; that is, if a true belief is safe, then it is true not only in the actual world but it is also true in most or all close possible worlds in which the protagonist has the belief. As we will see, the analytical project also defines non-epistemic luck in terms of modal robustness. According to this view, acquiring a true belief by luck is like any other case of non-epistemic luck (e.g., someone winning a fair lottery): when any of those things happen, they happen in the actual world but fail to

³Pritchard [2012, p.247-8].

happen in most or all close possible worlds; in short, lucky events are *modally weak*. However, this definition of non-epistemic luck is incorrect, for we can have *modally robust lucky events*.

Things get worse still for the analytical project. Our account of inferential knowledge, KFK, sidesteps all the worries raised for this version of anti-luck epistemology and undergirds a better account of epistemic and non-epistemic luck.

5.2 The Appeal To The Concept Of Epistemic Luck

In his seminal paper, “An Analysis of Factual Knowledge,”⁴ Peter Unger suggested that one knows that *p* only if it is not an accident that one’s belief that *p* is true.⁵ Unger says that what explains why the subject in Gettier’s original cases does not know is because “it is entirely accidental that he is right about the matter in question, whereas, for him to know, it must be quite the opposite. It must be not at all accidental that he is right about the matter.”⁶

Unger’s discussion is often cited as the first attempt to connect the Gettier Problem (i.e., the counterexemplification of the claim that, necessarily, *S* knows that *p* if, and only if, *S* has a justified true belief that *p*) and the notion of luck or accident. Unger himself never uses the term “luck” to describe these cases, but his unanalyzed use of “accidental” is usually taken to be synonymous with “lucky” by the anti-luck epistemologist.⁷ Peter Klein also suggested that Gettiered subjects fail to know because it is merely a *felicitous coincidence* that their belief is true

⁴Unger [1968].

⁵Even before Unger, Bertrand Russell had alluded to the idea that knowledge excludes the possibility of luck when he asked “can we ever know anything at all, or do we merely sometimes by good luck believe what is true?” in Russell [1912, p.131].

⁶Unger [1968, p.165].

⁷Most prominently in Pritchard [2005, ch.5].

given the subject's evidence.⁸ More recently, Linda Zagzebski has suggested that one can use the notion of luck in a formula designed to produce Gettier cases: first, construct a case in which S forms a false belief in such a way that it has the epistemic property (or properties) required for knowing and that would usually lead to a true belief were not for *bad luck*; second, amend the case in a way such that an element of *good luck* cancels the first element of bad luck and makes S's belief true.⁹ It is safe to say that Zagzebski's "formula" for how to construct a Gettier case is the standard understanding of how luck interferes in those cases.¹⁰

Even though these epistemologists (and many others) embraced the ideology according to which Gettiered subjects fail to know because they are lucky to have formed a true belief, it was not until recently that anti-luck epistemologists started offering a careful analysis of epistemic luck itself. Influenced by Ernest Sosa's work on the notion of safety,¹¹ Duncan Pritchard¹² proposed an analysis of epistemic luck in terms of the notion of safety and went as far as to suggest that the very point of the Gettier Problem is to show that a belief being justified and true does not exclude the possibility that the belief is true by luck. Gettier cases vindicate what Pritchard calls the *anti-luck intuition*: "when one knows, one's cognitive success (that is, one's believing truly) is not a matter of luck."¹³ Moreover, in this picture, epistemic luck, like moral luck, is a *kind of luck*. Gettiered subjects, like lottery winners and Nagel's reckless driver, are all lucky.¹⁴ Gettiered subjects

⁸Klein [1971].

⁹Zagzebski [1994, p.69].

¹⁰Note that Zagzebski's formula fits the original Gettier cases only if she is talking about the same belief throughout, for in the original cases there is never a time when the conclusion of Smith's reasoning is false.

¹¹For Sosa's views on safety see Sosa [1999, 2007].

¹²Pritchard [2005, 2012].

¹³Pritchard [2012, p.247].

¹⁴For the claim that epistemic luck is a kind of luck see Pritchard [2005, ch.5] and Coffman [2007]. See Nagel [1979], Williams [1981] and the papers collected in Statman [1993] for a discussion of moral luck.

are lucky to have a true belief, reckless drunk drivers are sometimes lucky when they do not hit anyone, and lottery winners are lucky to have won the lottery. Here is how safety-based accounts of epistemic and non-epistemic luck are usually presented:¹⁵

(Generic Luck) If an event E is lucky for S, then (i) E is significant to S (or would be significant, were S availed of the relevant facts); and (ii) E occurs in the actual world but does not occur in a wide class of the nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world.¹⁶

(Epistemic Luck) If a contingent true belief that p is lucky for S, then S believes truly that p in the actual world, but falsely in one (or more) nearby possible worlds in which S forms the belief that p in the same way as she formed her belief that p in the actual world.¹⁷

To see how an appeal to Epistemic Luck above explains why the protagonist in Gettier cases fails to know, remember Gettier's Case I and II:

The Coin Case

Suppose that Smith and Jones have applied for a certain job. And suppose that Smith has strong evidence for the following conjunctive proposition: (d) Jones is the man who will get the job, and Jones has ten coins in his pocket. Smith's evidence for (d) might be the fact that the president of the company assured him that Jones would in the end be selected, and that he, Smith, had counted the coins in Jones's pocket ten minutes ago. Proposition (d) entails: (e) The man who will get the job has ten coins in his pocket. Let us suppose that Smith sees the entailment from (d) to (e), and accepts (e) on the grounds of (d), for which he has strong evidence. In this case, Smith is clearly justified in believing that (e) is true.

¹⁵Since I will be arguing that all luck is epistemic in the sense that it involves knowledge, I will call non-epistemic luck "generic luck" in order to avoid potential misunderstandings.

¹⁶Pritchard [2005, p.128-33].

¹⁷cf., Pritchard [2005, ch.6], and Pritchard [2012, p.253-256]. I will discuss the issue of how to extend this account to include necessarily true/false propositions below.

The Ford Case

Suppose that Smith has strong evidence for the following proposition: (f) Jones owns a Ford. Smith's evidence might be that Jones has at all times in the past within Smith's memory owned a car, and always a Ford, and that Jones has just offered Smith a ride while driving a Ford. Let us imagine, now, that Smith has another friend, Brown, of whose whereabouts he is totally ignorant. Smith selects three place-names quite at random, and constructs the following three propositions: (g) Either Jones owns a Ford, or Brown is in Boston; (h) Either Jones owns a Ford, or Brown is in Barcelona; (i) Either Jones owns a Ford, or Brown is in Brest-Litovsk. Each of these propositions is entailed by (f). Imagine that Smith realizes the entailment of each of these propositions he has constructed by (f), and proceeds to accept (g), (h), and (i) on the basis of (f). Smith has correctly inferred (g), (h), and (i) from a proposition for which he has strong evidence. Smith is therefore completely justified in believing each of these three propositions. Smith, of course, has no idea where Brown is.

In both cases Smith has a justified true belief, but epistemologists do not want to say that he knows in either case. They accepted that these cases refute the claim that, necessarily, S knows that p if and only if S has a justified true belief that p. To paraphrase Monty Python's John Cleese in the memorable "dead parrot" sketch, as soon as Gettier's paper was published the traditional theory of knowledge became an "ex-theory".

According to the anti-luck epistemologist, Smith is lucky to believe truly in both cases. Smith's belief that the man who will get the job has ten coins in his pocket is true in the actual world, but false in worlds in which Smith has eleven or nine coins in his pocket and in worlds in which neither Smith nor Jones gets the job and the person who does get the job does not have ten coins in his pocket. Similarly, even though Smith's belief that either Jones owns a Ford or Brown is in Barcelona is true in the actual world, it is false in many close possible worlds in which Jones does not own a Ford, but in which Brown is not in Barcelona.

5.3 Precisifying The Anti-Luck View

The initial pass at explaining why the Gettiered subject fails to know in the previous section is too crude. In this section I refine that approach further. Not all potential problems for anti-luck views will be addressed, but I will make the view precise enough for our purposes.

Necessary Truths. It seems that we could have Gettier cases in which the target true proposition involves a necessary truth. In those cases it is not clear how the subject's true belief could be lucky, since there is no world in which her true belief is false. To see that, consider the following modified Ford case: the case goes as described by Gettier except for the fact that, instead of inferring that either Jones owns a Ford or Brown is in Barcelona from his false belief that Jones owns a Ford, Smith infers that either Jones owns a Ford or, either Brown is in Barcelona, or ϕ (where " ϕ " is some complicated mathematical truth Smith believes only because he likes the way it sounds). Smith's inference has now the valid form $p \vdash p \vee (q \vee \phi)$, but, since ϕ is a necessary truth, so is $p \vdash p \vee (q \vee \phi)$.¹⁸ So, in this modified Ford case Smith still does not know the relevant true belief, but the anti-luck epistemologist cannot run her story about luck and explain why Smith fails to know, for there are no worlds in which the proposition Smith believes is false. Duncan Pritchard suggests the following solution to this problem:¹⁹

the safety principle, properly understood at any rate, is concerned not with the safety of a belief that p (i.e. the proposition actually believed), but rather with the safety of a relevant doxastic output of a belief-forming process (which, while being a belief that p in the actual world, could be a belief in a different proposition in a nearby possible world). In this way, possible worlds where this process results in a different doxastic output from that which results in the actual world - such as a belief that q rather than that p - can be relevant to the safety of the belief. On the face of it at least, this

¹⁸Of course, as in the original Ford case, the main disjunction here is exclusive.

¹⁹Pritchard [2009, p.45].

modification of the view enables it to deal with cases involving necessary propositions.

The suggestion here is that we should emphasize the (un)safety of the belief forming process rather than the (un)safety of the belief. Thus, with regard to our modified Ford case, the anti-luck epistemologist can point out that, even though Smith's belief is necessarily true, he used a belief forming process that is unsafe - the process of forming beliefs about complicated mathematical matters on the basis of how propositions in that domain "sound" to him.²⁰

Not All Luck Is Created Equal. Even though anti-luck epistemologists take epistemic luck to be a "kind" of luck akin to generic luck, they agree that one may sometimes know that p even though luck intervenes in one coming to believe that p. They acknowledge that luck does not always exclude knowledge. Since I nearly got hit by a truck a few hours ago, I am "lucky" to believe truly that there is computer monitor in front of me. I am lucky to have any beliefs at all, but this luck does not prevent me from knowing that there is a monitor in front of me. The idea here is that I am lucky to be *in a position to believe* that there is a monitor in front of me, rather than lucky to *believe truly* that there is a monitor in front of me. Following Pritchard, call this kind of luck "Doxastic Luck."²¹ Doxastic luck is a kind of luck that is compatible with knowledge. Similarly, it might be a matter of luck that I found conclusive evidence for, say, the location of some long lost treasure without it being a matter of luck that I believe truly that the treasure is at the location the map says it is. This kind of luck, too, does not prevent me from knowing. Again, following Pritchard, call this kind of luck "Evidential Luck."²²

Too Many Gettier Cases. Anti-Luck epistemology risks classifying cases that

²⁰See also Sosa [2007, p.26-7]. Sosa relativises safety to "bases" and calls it "basis-relative safety."

²¹Pritchard [2005, p.136].

²²Pritchard [2005, p.138].

are clearly *not* Gettier cases as Gettier cases. The theory runs the risk of overgenerating Gettier cases. To illustrate that, consider a case in which Smith believes truly that his lottery ticket is a loser. Suppose Smith knows his ticket is one of ten million tickets and that he also knows that the lottery is fair and that it has only one winner. Many people think that Smith does not know his ticket is a loser even though, given what he knows, he is justified in believing his ticket is a loser and it is true that it is. But, Smith's justified true belief that his ticket is a loser is unsafe (i.e., there are close possible worlds in which he believes it is a loser in the same way he comes to believe that in the actual world, but his belief is false in those worlds). So, it seems that anti-luck epistemology's explanation of why Gettiered subjects fail to know overgenerates and we are now forced to hold the implausible view that the lottery case is a Gettier case. This seems like the wrong result.

5.4 Against Anti-Luck Epistemology

5.4.1 Lucky But Safe True Belief: Against Epistemic Luck

There is a potential difficulty for this account of what goes awry in Gettier cases. What if I modify the original Gettier cases and make it the case that Smith would have ten coins in his pocket in all close possible worlds (say, because that is the exact change he gets when he takes any train in his town) or Brown was in Barcelona in all close possible worlds (say, because he is bedridden)? Supposedly, Smith would still have a justified true belief that does not amount to knowledge in those modified Gettier cases but his true belief would be modally robust - i.e., it would still be true in all or most close possible worlds. It is unclear what the anti-luck epistemologist would say in reply to this point.

Actually, there are many counterexamples to epistemic luck in the literature that make a similar point. For example, consider the following case offered by

Avram Hiller and Ram Neta, a case which shows that, contrary to what anti-luck epistemologists suggest, safe true beliefs are sometimes luckily acquired.²³

The Counterfactually Robust Nogot Case

Naira believes that Jones owns a Ford, and so she uses disjunction-introduction to infer, and so comes to believe, that either Jones owns a Ford or Brown will not win a Grammy award. As it turns out, Brown is not even a musician and has nowhere near the skills needed to win a Grammy, so the second disjunct is true. Indeed, we may suppose that Brown is so lacking in even the most rudimentary musical ability that the second disjunct is true in nearly all, if not all, nearby possible worlds in which Naira forms her belief in the same way (namely, by disjunction introduction). However, ... Naira has never given any thought to whether Brown can, or will, win a Grammy, and she does not believe the second disjunct to be true, or perhaps Naira is mistakenly convinced that Brown will win a Grammy - perhaps she overheard someone assert “Braun will win a Grammy”, and she mistook the person to be talking about Brown. Even when it turns out that Brown does not win that Grammy, still Naira failed to know that Jones owns a Ford or Brown will not win a Grammy. And she failed to know this even though in nearly all, if not all, the nearby possible worlds in which Naira forms the belief in the same way, the belief is true, for in nearly all, if not all, of those worlds, Brown will not win a Grammy.

While epistemic luck as safety cannot account for why Naira fails to know a truth she is justified²⁴ in believing, KFK has a straightforward story to tell about the case: Naira’s inference involves essentially the falsehood that Jones owns a Ford, thus she fails to know that either Jones owns a Ford or Brown will not win a Grammy award.

²³cf., Neta [2007, p.308]. See also Lackey [2006, p.288] and Coffman [2010, p.246] for other cases making basically the same point. Lackey’s case is particularly interesting because it is a variation on Goldman’s Barn-Façade case.

²⁴Naira is not only justified in believing that Jones owns a Ford, but she is also justified in believing that Brown will not win a Grammy award in the version of the case in which she hears someone say that “Braun” will win a Grammy award.

5.4.2 Lucky But Counterfactually Robust Events: Against Generic Luck

Even though the industry of counterexamples to generic luck is nascent and, therefore, not (as of yet) as prolific as the industry of counterexamples to epistemic luck as safety, it has the potential for growing fast. Jennifer Lackey²⁵ has offered the following case against generic luck as counterfactually weak events (i.e., events that hold in the actual world, but fail to hold in most close possible worlds):

The Buried Treasure Case

Sophie, knowing that she had very little time left to live, wanted to bury a chest filled with all of her earthly treasures on the island she inhabited. As she walked around trying to determine the best site for proper burial, her central criteria were, first, that a suitable location must be on the northwest corner of the island - where she had spent many of her fondest moments in life - and, second, that it had to be a spot where rose bushes could flourish - since these were her favorite flowers. As it happens, there was only one particular patch of land on the northwest corner of the island where the soil was rich enough for roses to thrive. Sophie, being excellent at detecting such soil, immediately located this patch of land and buried her treasure, along with seeds for future roses to bloom, in the one and only spot that fulfilled her two criteria. One month later, Vincent, a distant neighbor of Sophie's, was driving in the northwest corner of the island - which was also his most beloved place to visit - and was looking for a place to plant a rose bush in memory of his mother who had died ten years earlier - since these were her favorite flowers. Being excellent at detecting the proper soil for rose bushes to thrive, he immediately located the same patch of land that Sophie had found one month earlier. As he began digging a hole for the bush, he was astonished to discover a buried treasure in the ground.

Vincent is clearly lucky to have found the treasure, but it is not the case that he fails to find the treasure in any of the close possible worlds. Vincent does not decide to go plant roses in the northwest corner of the island by chance. He does that in virtue of his deep seated desire to honor the memory of his mother and his love for that particular area of the island. He is lucky to have found Sophie's

²⁵Lackey [2008, p.261].

treasure, even though this is a counterfactually robust event (i.e., an event which holds in all close possible worlds).

The last two sections spell trouble for advocates of the analytical project of defining “knowledge” who take an anti-luck condition (understood as a condition ruling out modal weakness) as necessary for knowledge. The analytical project is carried out to fruition only if one’s proposed definition is exceptionless. However, the counterfactually robust Nogot case seems to show that a safety condition is not necessary for epistemic luck. Lackey’s buried treasure case shows that the anti-luck epistemologist is wrong about the nature of luck itself.

Is there still room for luck in epistemology? I think there is, but the role I see for luck in epistemology brings no solace to the analytical project. I will get back to this in section 5.5.3 below. Right now, I want to discuss briefly a different reaction one might have to the result that safety does not capture epistemic luck.

5.5 Gettier-Lucky Knowledge?

The counterfactually robust Nogot case shows that safe beliefs may be Gettiered. One might give up on the idea that knowledge requires safe belief, but insist - surprisingly enough - that knowledge is compatible with the kind of epistemic luck we find in Gettier cases. In short, one might insist that, when well understood, knowledge is compatible with the kind of luck Pritchard takes to be incompatible with knowledge.

Stephen Hetherington has recently argued this latter point.²⁶ According to him, there are Gettier cases in which the Gettiered subject knows the target proposition. In what follows we will look at Hetherington’s case for Gettiered-lucky knowledge and show that it is anything but inconclusive. The case Hetherington discusses is *not* a Gettier case. Not only that, we can give an account of

²⁶cf.Hetherington [2011],Hetherington [2013].

epistemic luck different from unsafe belief and that gives the result that epistemic luck is incompatible with knowledge.

5.5.1 The Argument for “Gettierism”

Before we turn to Hetherington’s alleged example of Gettier-lucky knowledge, let me first state the traditional argument for the claim that knowledge is incompatible with Gettier-luck. I will take this to be more or less the argument Hetherington is attacking, even though he never presents it in the same explicit way I will present here. Furthermore, I do not take this argument to necessarily represent the inference epistemologists (or anyone else for that matter) actually draws when confronted with Gettier cases. The argument is meant as a mere rational reconstruction of the argument epistemologists accept, explicitly or not. Premise one states that, for any subject, if she is in a Gettier situation, then her justified true belief is, from an epistemic point of view, luckily acquired. In symbols:²⁷

$$1. \quad \Box(\forall s)(\forall p)(Gs \rightarrow \text{JTB}_{sp} \ \& \ Lp)$$

Premise two says that, if S’s true justified belief is epistemically lucky, then S does not know. This is the premise Hetherington challenges:

$$2. \quad \Box(\forall s)(\forall p)((\text{JTB}_{sp} \ \& \ Lp) \rightarrow \neg K_{sp})$$

Finally, premises one and two entail the conclusion that an agent in a Gettier situation fails to know. Following Hetherington, call this claim “Gettierism.”

$$3. \quad \Box(\forall s)(\forall p)(Gs \rightarrow \neg K_{sp})$$

Hetherington attacks premise two with a case in which the agent allegedly has Gettier-lucky knowledge. I now turn to this case.

²⁷ “G_x” stands for “x is in a Gettier-like circumstance;” “JTB_{xy}” stands for “x truly and justifiably believes that y,” “L_x” stands for “the belief that x is epistemically lucky” and “K_{xy}” stands for “x knows that y.”

5.5.2 An Alleged case of Gettier-Lucky Knowledge

Hetherington argues that the following case is a Gettier case.²⁸

The Noah Case

In a normal environment and under good lighting conditions Noah forms the belief that he is seeing a tree in front of him. He believes that on the basis of the normal sensory evidence acquired through his reliable faculty of vision. He also believes the following proposition, which he takes to be the sole explanation of how he finds himself believing that he sees a tree. Call it “Naive Realism” (NR): ‘If I had to explain how it is true that I am seeing a tree, I would reach for this description: [then follows a brief statement of a naive realism about perception - in effect, what most philosophers would term a ‘folk theory’ of perception]. That is all I would need. It is a full explanation.’

We are asked to suppose (i) that NR is false; (ii) that NR is “part of Noah’s evidence or reasoning” for his belief that he sees a tree in front of him; and, (iii) that Noah knows nothing or close to nothing about “the scientific details of how the world, courtesy of any detail not observed by him, is making his belief true.”²⁹ From the description of the case and from (i)-(iii), Hetherington takes it to follow (iv) that Noah’s justified true belief involves epistemic luck of the sort epistemologists traditionally take to be incompatible with knowledge (rather than evidential or doxastic epistemic luck), for Noah believes (falsely) that NR is the sole explanation for how his belief came to be true. If this were right, then (v) Noah is in a Gettier situation. The problem for defenders of Gettierism is that claims (i)-(v) entail (vi) that Noah knows that he sees a tree in front of him. This argument, if sound, shows that Noah knows he sees a tree in front of him, even though his true belief was luckily acquired. I think this argument is unsound. I will argue that Noah knows that he sees a tree, but that his knowledge is not Gettier-lucky. As I will argue in the next section, instead of denying that Noah

²⁸Hetherington [2011, p.80].

²⁹Hetherington [2011, p.80-1].

actually knows he sees a tree in front of him, we should deny that this case is in fact a Gettier case.

Now, remember our discussion in chapter 4 of what makes a case a Gettier case. There I argued that C is a Gettier case only if the S's justified true belief relies essentially on a proposition S fails to know. Since Noah's justified true belief that he is seeing a tree in front of him depends on the proposition NR he fails to know, one might think that this is in fact a Gettier case. This would be a mistake, however, for at least three different reasons: (i) it is possible that "Noah sees a tree in front of me" does not depend essentially on NR; (ii) it is also possible that, even though "I see a tree in front of me" depends essentially on NR, Noah is not Gettiered since essential dependence on a falsehood is only a necessary condition on a case being a Gettier case. In the remainder of this section, I will raise doubts as to whether "I see a tree in front of me" depends essentially on NR or not. I will conclude that there is reason to believe that it does not.

First, remember that there are cases of harmless ignorance (see chapter 2 above). Remember the case we used to illustrate that idea:

The Harmless Ignorance Case

Smith arrives through reasoning to the belief that (p) someone in his office owns a Ford. His premises are (q) that Jones owns a Ford, (r) that Jones works in his office, (s) that Brown owns a Ford and (t) that Brown works in his office. As it turns out s is false while q, r and t are true.

Intuitively, Smith knows that someone in his office owns a Ford even though the reasoning he used to arrive at that belief includes a falsehood. However, it is generally understood that this falsehood is harmless, in that it does not prevent him from knowing that someone in his office owns a Ford. Smith's belief in the falsehood s is neither causally nor evidentially essential to Smith's belief in the truth that p. Smith's justified true beliefs that q, r, and t could have easily

caused Smith to believe that p , q , r and t also justify him in believing that p . In sum, Smith's belief that s is evidentially and causally dispensable or irrelevant, meaning that he would still know that p even if he did not believe that s .

Examples of harmless ignorance abound. One can gain knowledge on the basis of an enumerative induction which contains only one false premise. For example, if I saw hundreds of black ravens but believed, of a bird that was in fact a boat-tailed crackle, that it was black and a raven, I clearly still know that all ravens are black. Of course, I do not know that all ravens are black if I do not succeed in recognizing ravens any better than chance. There will also be borderline cases in which we cannot be sure the number of false beliefs about ravens in a crackle infested place prevents one from knowing or not (e.g., does S know that all ravens are black if she saw sixty two black ravens and took thirty eight boat-tailed crackles to be black ravens?) But my example is not a borderline case. In the example I am considering the one false belief about the crackle is a harmless falsehood and, thus, does not prevent me from knowing that all ravens are black.

Noah's case and cases of harmless ignorance seem to share one important feature. In particular, Noah's false belief that NR seems both causally and evidentially inessential to his true belief that he sees a tree in front of him. This makes Hetherington's case structurally similar to the Harmless Ignorance case, not to Gettier's original cases. Noah's belief in NR is causally dispensable or inessential, because one can plausibly imagine that Noah would still believe he sees a tree even if he did not believe that NR was the case - say, because Noah was well versed in cognitive sciences and would not appeal to anything like NR to explain how or why he sees the things he thinks he saw.

Noah's belief in NR seems also evidentially dispensable or inessential. Noah would clearly be justified in believing he sees a tree in front him even if NR were not in his evidence set. Hetherington himself says that Noah's tree-belief "is

correct about the world in a normal way, and he has quite normal evidence in support of it - normal sensory evidence, plus the “everyday” acceptance of NR.” But, if Noah’s belief that he sees a tree depends for its positive epistemic status on his sensory evidence, not only on NR, then we may ask - like we did with respect to Lehrer’s Harmless Ignorance case - whether Noah’s sensory evidence could turn his true belief into knowledge all by itself. I think it is clear that, in most cases, Noah’s sensory evidence would be good enough to give him knowledge. I say “most cases,” because one can imagine variations of the case in which Noah has no explanation of how or why he sees a tree, variations in which Noah has a scientifically sound explanation of that fact and, finally, variations in which Noah has really bad explanations of it. Let us call the first kind of variation of the Noah Case “Neutral Variations,” the second kind “Good Variations” and the third kind “Bad Variations.” Maybe, in bad variations of the case his sensory experience is not enough to turn his belief into knowledge, because his bad explanation would defeat whatever epistemic support he got from his sensory evidence. For example, if Noah thought that what explained why he saw a tree was that he was God and that, for God, seeing a tree is “creating” a tree with His mind, then we might want to say that Noah’s view on these matters is so outlandish that the justification he gets from his sensory evidence is undermined by his view and that he fails to know he sees a tree. In any event, given Hetherington’s description of the case, it is plausible for us to think that Noah would not be in such a scenario in many if not all nearby possible worlds to the one he is in. Since Noah seems to be in a world, for all practical purposes, identical to the actual world, then in most of those close possible worlds, Noah is in a situation similar to the one in which most of us find ourselves. Thus, since I can know that I see a tree in front of me on the basis of my sensory experience alone, so can Noah. Like Smith in the Harmless Ignorance case, Noah’s false belief is evidentially inessential or dispensable. We seem, therefore, to have strong reason to think that Noah’s case

is a case of harmless ignorance. It is obvious that Noah's true belief is safe. So, Pritchard too could deny that this is a case in which there is epistemic luck.

5.5.3 What is lucky about Gettier-Lucky belief?

Hetherington takes epistemic luck to be the following:³⁰

(H-Luck) The belief in question is true in a way that is quite different from how the believer's evidence indicates its being true. The idea is that it is lucky that the belief is true in the way it is, given the very different way in which the evidence indicates its being true.³¹

There are a few problems with H-luck. First, it is not clear how it relates to generic luck. And, given that it is the anti-luck epistemologist's own desideratum that those accounts of luck are so related, this is not good. For instance, it is far from clear what the connection is between H-luck and the kind of luck paradigmatically involved in fair and large lotteries. Second of all, even if we do not try to reconcile H-luck with an account of non-epistemic luck, it also faces problems as a characterization of Gettier-luck.

First of all, it is far from clear what exactly Hetherington means by "the belief x is true in a way w ." "Belief" in this locution supposedly refers to the propositional content of beliefs - as in "Liz's belief that she won the lottery is true" - rather than to the belief states themselves - as in "Liz's belief that she won the lottery brought tears to her eyes" - for it makes no sense to talk about a belief state "being true" in this or that way. Strictly speaking, belief states are neither true nor false and, therefore, cannot be "true in a way w ". Therefore, this locution has to mean something like "the proposition p is true in a way w ." But, if this is what Hetherington has in mind, then what does he mean when he

³⁰Hetherington [2011, p.78].

³¹Hetherington [2011, p.79].

says that some proposition p is true in a way w ? Even though propositions are the rightful bearers of truth and falsehood, it seems to make as much sense to talk about a proposition being true “in a way w ” as it makes sense to talk about a belief state being true “in a way w ,” for all propositions seem to be true in one and the same way - by, roughly, “matching” the facts they are about or represent. There is only one way in which the proposition “ a is F ” can be true: by a being F . Of course, one might want to say that “Some F s are G s” is “made true in different ways” - say, by a being F and G or by b being F and G and so on. But that would restrict the application of H-luck to beliefs with that kind of content. Clearly, however, beliefs with all kinds of content can be Gettier-lucky, so that will not work. This problem with Hetherington’s use of “the belief x is true in way w ” pervades all his discussion of Gettier-luck. Nowhere in that discussion does he say what he means by that locution. The closest he ever gets to making it more precise is in a footnote.³² In this footnote he says that he is using the locutions “true in a way w ” “as non-technically as is allowable for [his] purposes” and that he will “call upon truthmaker theory” to make the notion more technical in a later chapter. The problem is that, as far as I can see, there is no technical account of the locution “truth in a way w ” in the later chapter he mentions in the footnote. At the very least, the topic is never directly addressed. True, Hetherington summarizes David Armstrong’s truthmaker theory³³ in chapter five, section four, but there he never cashes out the promissory note he issued in the footnote I just mentioned - i.e., he never gives “true in a way w ” a more “technical treatment” using a truthmaker theory.³⁴ So, for all Hetherington tells us, it is far from clear what he means by the key expression in his only explicit account of Gettier-luck. And, to make matters worse, we saw that the only plausible

³²Hetherington [2011, p.93, fn37].

³³Armstrong [2004].

³⁴It is also worth noting that Armstrong himself never uses the locution “true in a way w ” in Armstrong [2004].

reading of the locution “the belief x is true in a way w ” artificially limits the application of intuitive luck to beliefs with a certain kind of content. For these reasons it seems safe to say that Gettier-luck is very likely not H-luck.³⁵

We saw above that whatever Gettier-luck is, it is of the same kind as generic luck. I showed above that safety-based accounts of epistemic and generic luck face counterexamples. In the remainder of this section I will look at a different accounts of generic luck and see if this type of luck fares better than the safety-based view does. If it does, then we might hope to derive an account of Gettier-luck. However, if it does not work, then the central tenet of anti-luck epistemology according to which Gettier-luck is a kind of luck is not satisfied, and we should be a lot less confident that it can generate the right account of epistemic and Gettier-luck.

E. J. Coffman has presented a definition of luck that cashes out generic luck in terms of chance, lack of control and significance conditions.³⁶ He adds to those conditions a condition requiring that the lucky agent be a sentient being.

(C-Luck) S is lucky with respect to E at t iff (i) S is sentient at t ; (ii) E has some objective evaluative status for S at t ; (iii) there was just before t a large chance that no event sufficiently similar and equal in significance to E would occur at t ; and, (iv) E lies beyond S 's direct control at t .

Condition (i) excludes from the domain of luck inanimate objects like trees and

³⁵Further evidence for the idiosyncratic nature of H-luck comes from yet another footnote (Hetherington [2011, p.78,fn8]) in which Hetherington discusses the idea that Gettierized beliefs are “true in a way that is not appropriately linked to the believer’s evidence.” There he suggests that this idea should be “supplemented” by a truthmaker theory, in particular by the theory offered by Heathcote [2006]. Curiously, Heathcote deals almost exclusively with the subset of Gettier-like cases in which the target proposition is an existential generalization - suggesting that Heathcote, like Hetherington, restricts epistemic luck to a small class of beliefs with existentially generalized contents. Maybe Hetherington had Heathcote in mind when he suggested H-luck, but that hardly shows that this notion is (or should) be widely held.

³⁶Coffman [2007, p.396].

rocks. In fact, for Coffman,³⁷ “luck’s domain includes all and only sentient subjects, individuals that have inner lives that can go better or worse for them.” Thus, in reply to the suggestion that we sometimes speak truly when we use sentences such “It was lucky for your family portrait to have survived the house fire” or “It was lucky for that rock to have survived water erosion,” Coffman argues that these sentences in fact express, respectively, a true propositions about the relation between a sentient being and a nonsentient object and a proposition about the extremely low probability of that event taken place. So, according to this translation schema, if I use “It was lucky for your family portrait to have survived the house fire” in a context C, I express, in C, the proposition that you are lucky vis-à-vis the fact that the portrait survived the fire; whereas, if I use “It was lucky for that rock to have survived water erosion” in C, I thereby express the proposition that the rock’s surviving the water erosion was highly unlikely.

One worry about the claim that luck applies only to sentient beings is that it excludes from the realm of luck entities like groups or organizations. We sometimes speak truly when we utter sentences like “It was unlucky for that criminal organization that the body they dumped into the east river floated” or “It was lucky for FC Barcelona that two of the best defense players from the Manchester United scored two goals against their own team.” One can also imagine the following situation. Suppose ACME Corporation is going bankrupt. The board members meet and they all agree on a desperate measure to try to save the company: they will buy a lottery ticket from the state lottery. The lottery jackpot is five hundred million dollars - enough money to keep ACME running for at least another year. As it turns out, ACME’s ticket is the winner. I think it is clear that it is lucky for ACME to have won the lottery. This is true even though no one (particular) sentient being is lucky to have won the lottery.

Coffman might want to reply to this kind of worry by saying that the group

³⁷Coffman [2007, p.387].

or organization is only *derivatively lucky* in the sense that the sentient beings that make up the group are lucky and that the group is lucky *in virtue of* being made up by those lucky individuals. That is a possible reply to this worry, but I am not sure what its merits really are. One problem with it is that it entails that, by changing the individuals making up the organization, one will end up with a different organization. But that is implausible - Pepsi Corporation is the same company today as it was when it was created, even though all the founding members have died. In any event, I will assume for the sake of argument, that only sentient beings can be lucky or unlucky.³⁸

Condition (ii) in C-Luck says that an event E needs to have “some objective evaluative status for S at t” in order for S to be lucky with respect to E at t.³⁹ This condition gets support from cases such as lottery wins. Winning the lottery objectively improves one’s life in paradigmatic cases. The problem is that a variation of that case can be turned into a counterexample to this condition on luck. That an event E needs to have “some objective evaluative status for S at t” in order for S to be lucky with respect to E, entails that S cannot be lucky with respect to E, if S’s wellbeing is neither improved nor worsened by E. Unfortunately, one can very easily imagine cases in which the obtaining of E neither improves nor worsens S’s wellbeing, but S is lucky that E obtains. Suppose Liz’s wealth amounts to more than fifty billion dollars. One day Liz buys a lottery ticket in a fair and large lottery for fifty cents, because she pities the person selling her the ticket at the subway station (maybe the person told

³⁸Note that I am *not* taking issue with the claim that groups have mental states. On the contrary, I do think groups can know all kinds of things. For example, physicists and biologists know all kinds of things about the universe and the living things in the known universe. The point against Coffman is just that, intuitively, groups are not sentient beings, even though their individual members are. But, since Coffman thinks sentience is a necessary condition on something being lucky for something else, it follows that his account mistakenly rules out the possibility that groups are lucky.

³⁹Pritchard [2005, p.132] and Levy [2009, p.490] also accept that something equivalent to (ii) is a necessary condition on luck.

Liz he gets a small fraction of the fifty cents for every ticket he sells). Suppose further that the lottery's prize is only *fifty-one cents*. Now, winning the main prize in this lottery is not important to Liz - she is, after all, a multi-billionaire: she bought the ticket only because she wanted to help the person selling it, and she stands to gain only one cent if she wins. It does not seem to follow from any of this, however, that Liz is not lucky to have won the lottery. She beat the odds! But this is what condition (ii) would have us say, for, by assumption, winning the lottery does not make Liz's wellbeing any better or worse than it did before she won it. The intuition seems to be that Liz is lucky because she beat the odds that were stacked against her, not because winning had any "objective evaluative status" for her. Since there is nothing special about this particular example one can multiply them as much as one likes. Condition (ii) is not a necessary condition on luck.⁴⁰

Condition (iii) says that S is lucky with respect to E at t only if there was just before t a large chance that no event sufficiently similar and equal in significance to E would occur at t. Since Coffman takes this condition to be "roughly equivalent" to the safety-based account of luck I criticized above,⁴¹ my criticism carries over to C-Luck.

Condition (iv) in Coffman-luck says that S is lucky with respect to E at t only if E lies beyond S's direct control at t. According to Jennifer Lackey, the

⁴⁰Peter Klein suggested that winning the lottery *does* improve Liz's wellbeing in that she acquired a capacity she did not have before winning the lottery. According to Klein, she now has the capacity to, say, pay the ransom of her children in a situation where kidnappers demand all her billions *plus one cent*. I do not think this shows what Klein claims, however. Consider: the case goes as described in the main text, except for the fact that, on her way to deliver the one cent prize to Liz, the person carrying the money loses the penny and, knowing that Liz does not care for the money anyway, never bothers to give Liz any other penny. Now, should we say that Liz is not lucky to have won the lottery because she never got her one cent prize and thereby acquired the capacity to spend all her billions plus one cent? I do not think so. But this is exactly what we would expect if Klein's suggestion were on the right track. So, I conclude that, contrary to Klein's suggestion, Liz's wellbeing was not improved by her winning the lottery.

⁴¹Coffman [2007, p.6].

following case is a counterexample to this claim:⁴²

The Demolition Worker Case

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

The point of the example is that Ramona is lucky to have blown the warehouse even though she has control over whether to press the button or not. Lackey concludes that one can be lucky with respect to an event even if one has control over that event.

A natural response to the case is for one to say that, although the situation Ramona is in is riddled with luck, the luck is not where Lackey puts it. Consistent with the lack of control requirement on luck, one can argue that Ramona was lucky to be in control of the detonation, but that, once in control of the detonation, she was not lucky to have exploded the warehouse. Ramona was lucky to be in a position to cause the detonation, but, once so positioned, causing the detonation was not a matter of luck for her. That is, even though bad luck intervened in the form of a rat chewing the electrical wires, good luck also intervened in the form of Ramona's colleague hanging his jacket at the right place and allowing electrical current to flow again; Ramona went from having no control to having control over the detonation. The point is that, at the moment when Ramona exploded the warehouse she was in control of the detonation.⁴³

⁴²Lackey [2008, p.258].

⁴³Coffman [2009, p.502].

Lackey [2008, p.259] anticipated this reaction to her case: “although an event may be within a given agent’s control, that the agent has such control can itself be largely a matter of luck, and hence the event resulting from this control can be lucky as well.” Hence, Lackey’s response to me would be that, since Ramona acquired control over the detonation in a lucky way, the warehouse blowing up when she pushed the button is also lucky. The general principle behind this response is that, if one is lucky to be in control of acting in a certain way and acting in this way brings about a certain outcome, then one’s relation to this outcome *inherits* the luck one enjoyed from being in control of acting in the way that brought the outcome about. Coffman aptly called this inheritance thesis “The Luck Infectious Thesis” (LIT)⁴⁴ and stated it thusly:

(LIT) If you were lucky to be free to A and you A-ed, then you are lucky that you A-ed.

Given (LIT), Ramona was lucky to have exploded the warehouse, because she was lucky to be free to explode the warehouse. What is more, according to Lackey,⁴⁵ the fact that Ramona “would regard the resulting explosion as an event whose occurrence is extraordinarily lucky” is evidence that (LIT) is true. However, whether what Ramona would say supports the idea that she is in fact lucky, depends, among other things, on whether she can competently draw the distinction between *being positioned to A* and *A-ing* and also on whether she can be reasonably expected to apply this distinction competently.⁴⁶

Not only does (LIT) enjoy little support, but also counterexamples to (LIT) can be easily generated. Suppose Liz’s psychology professor flips a coin and, as a result, assigns Liz to the control group of a psychological experiment. Suppose

⁴⁴Coffman [2009, p.503].

⁴⁵Lackey [2008, p.256].

⁴⁶Coffman [2009, p.503].

further that, once Liz gets into the room with an older participant, this person has a heart attack. Now, if Liz is a certified first responder, it is not a matter of luck that her fellow participant receives the proper emergency care and survives his heart attack, even though it was a matter of luck that Liz was there to save him and she did.⁴⁷

I do not believe this reply against (LIT) secures the claim that the lack of control condition is a necessary condition on luck, though. To see that, consider the following case.⁴⁸

The (Almost) Rigged Lottery Case

Liz works for the FBI improving the protection of governmental computerized systems against cyber attacks. One day, while working as a freelancer for the New York Lottery, she sees herself in a situation where she could rig the lottery in her favor and pocket ten million dollars. All it takes is a couple of short lines of code. What is more, since she is great at what she does, she knows she could erase all digital footprints and make the lottery look legit. She also knows she could use a friend as the “winner”, so that no suspicion would be raised. Further, she knows she could rig the lottery now, two weeks before the drawing, or up to five minutes before the drawing takes place. Liz never rigs the lottery, though; she is not a criminal. She does take note of these facts and takes steps towards improving the security of the Lottery’s cyber security instead.

Liz has direct control over what ticket would be the winner, but she chooses not to exercise that control and rig the lottery in her favor.⁴⁹ Now for the kicker: suppose Liz had a ticket for the lottery and that it is the winner. She did not

⁴⁷Coffman [2009, p.503].

⁴⁸This is a modified version of a case discussed by Coffman [2009, p.503-4].

⁴⁹It could be argued that Liz did not in fact have direct control over the lottery, because she did not take the necessary steps to rig it. She could have rigged it, but, since she did not, she did not have direct control over the outcome. I think this reply presupposes an excessively narrow account of direct control. If true, it would entail that anyone who is in a position to ϕ (i.e., anyone who knows all the individually necessary and collectively sufficient steps to ϕ -ing and is so situated that, were she to carry out those steps, nothing would prevent her from ϕ -ing), but chooses not to ϕ fails to have direct control over ϕ -ing. This is an absurd result. I do have direct control over whether I type an “a” between “(” and “)” right now, but I chose not to exercise that control. The fact that I decided not to add that letter between the parentheses does not show that I do not have direct control over typing it there.

rig the lottery, so she won fair and square. However, Liz is lucky that she won the lottery even though she had control over the lottery result. Thus, the lack of direct control is not a necessary condition on luck.

Something Coffman says in reply to a similar case suggests the following rejoinder to this case.⁵⁰ Liz was lucky to have won the lottery *fairly*. That is a different event from Liz's winning the lottery (full stop). She was free to rig the lottery and make herself the winner, but she was not free to do anything that would have resulted in a legitimate win. Once we distinguish between Liz's *fair* win and her win (full stop), we see that she did not have control over whether she would win *fairly* or not; she only had control over whether she would win (full stop). Therefore, we can uphold the lack of control requirement on luck, for Liz was lucky to have won the lottery *fairly*, but not lucky to have won the lottery (full stop).

This is a mistaken reply to the case, and it does not establish that lack of control is a necessary condition on luck. First of all, notice that I could have described the case in a different way, a way that entails that Liz was lucky even though she did have control over the outcome of the lottery. Liz did not merely have control over whether she would win (full stop), she had control over *whether she would win fairly*. When she decided not to rig the lottery, she made it possible for her to win the lottery fairly; she could not have won the lottery fairly if she had decided to rig the lottery. So, Coffman cannot say that Liz did not have any measure of control over her fair win. In the usual case, the lottery winner has no control whatsoever over his or hers win. Normally there is nothing an agent can do to influence the result of a state run lottery. But, the case I am discussing is not the usual case. Here, there is something Liz can do to influence the result. Once she finds herself in a position to rig the lottery in her favor, and to do so without anyone ever knowing, she is in a position to decide not only who wins,

⁵⁰Coffman [2009, p.504].

but also whether the win will be fair or unfair.

Compare: my religion does not allow me to take drug A, but if I do not take drug A I have 1% chance of surviving. Suppose my doctor leaves me alone in a room with many flasks of drug A and that I know I could take the right dosage of the drug, cure myself from the disease threatening to kill me, and that no one would ever know I took it. Now, if I do not take the drug and survive the disease, I am lucky to have survived, not lucky to have survived “fairly” - where “surviving fairly” means “surviving without breaking the rules of my religion”. As in (Almost) Rigged Lottery Case, I also have control not only over whether I survive (full stop), but over whether I survive without breaking the rules of my religion. By choosing not to take the drug, I have also chosen that, *if I survive*, then I will survive without breaking the rules of my religion.⁵¹ The upshot is that Coffman’s rejoinder in defense of the lack of control condition fails and we can conclude that this is not a necessary condition on luck.

The general moral is that C-Luck is a hopeless account of generic luck and, thus unsuitable as a model of epistemic luck.

We have to recognize, however, that something *like* the chance condition on C-luck must be a necessary condition on luck. There seems to be a sense in which an event being lucky depends in part on that event being improbable, unusual, or *chancy*. Can we appeal to anything besides counterfactual weakness to explain the connection between luck and chanciness? Yes. Besides not having control over her ticket being the winner, someone who plays the lottery does not have the right to expect her ticket to be the winner. S, at t, has the right to expect event E to hold at a later time t_n only if S is in a position to know at t that E will hold at t_n . So, once Liz in (Almost) Rigged Lottery Win decided not to interfere with the lottery, she was not in a position to know that her ticket would be the winner and,

⁵¹Of course this does not mean that I have control over whether I will survive, but it does show I think, that I have control over whether I survive without breaking the rules of my religion, *if I survive at all*.

thus, did not have the right to expect it to win. The way in which chance matters for luck, then, is more subjective than the requirement of counterfactual weakness suggests. Counterfactual weakness concerns what is objectively likely, when in fact what is needed is a condition telling us what the subject could rationally, given what he knows, expect. This shift allows us to explain why Vincent in the Buried Treasure case is lucky to have found the treasure. Vincent was lucky to have found the treasure because he did not have the right to expect to find it.

Roy Sorensen has argued that luck judgments are implicitly of the form “Given q , S was lucky that p .”⁵² The “ q ” in “Given q ” is something S knows. Hence, I can say that getting a flat tire on my way to catch a flight is bad luck *given that I know that I am driving a brand new car*, but that the same event amounts to good luck given that I know both that I am driving a brand new car *and that my flight will crash*.⁵³ In this case, then, to say that “It will be lucky that I got a flat” is to say that things will turn out better than I now may reasonably expect.⁵⁴ I can state this epistemic condition on luck in the following way:

(Chance) S is lucky with respect to E at t only if S is not in a position to know at t that E will hold at time t_n ($t_n > t$).

First, I need to say something about the notion of “being in a position to know”. Here is what I mean by this expression: one is in a position to know that p if and only if one would know that p if one were to believe that p on the basis of one’s total available evidence.⁵⁵ Consequently, according to Chance, being lucky amounts to a kind of ignorance. One is lucky to win a fair and large lottery for which one bought only one ticket because one was not in a position to know one would win before one’s number was drawn. One also could not have reasonably

⁵²Sorensen [1998, p.187].

⁵³Sorensen [1998, p.187-8].

⁵⁴Sorensen [1998, p.188].

⁵⁵Here too I presuppose a time and a world index.

expected to win before learning the result. One is lucky that the meteorite did not hit one's house because one was not in a position to know that it would not before it hit someone else's house. One also could not have reasonably expected that the meteorite headed to one's neighborhood would not fall on one's house. I was lucky to have gotten the grant because the competition was fierce and, therefore, I was not in a position to know I had gotten it before the official result was in, and I could not have reasonably expected to win it. What is more, if one were asked, right before the relevant event happened, whether it would hold or not, one would not be in a position knowledgeably answer that question given the evidence one had available to one at that time. This explains why it seems appropriate for the subject in any of those situations to reply "I don't know" to the question "Will E hold?"

Also, remember that we are taking statements of the form

"S is lucky that E"

to be elliptical for something like

"Given k, S is lucky that E occurred,"

and in which "k" refers to a contextually relevant body of knowledge which does not include the proposition that E has occurred. The idea is that, if E is lucky, then one's total knowledge before one learned that E occurred did not exclude the possibility that $\neg E$. In that sense, if I say, today, that Liz is lucky she won the lottery, and everyone in our context knows Liz won the lottery ten years ago, my statement was elliptical for a statement whose "k" refers not to the body of knowledge that we share in this context of utterance, but to the body of knowledge that we (including Liz) shared *before* we learned Liz's ticket was the winner. So, when I say, today, that Liz is lucky she won the lottery, I assert something that is made true by the state of our knowledge at a time t before any of us (including

Liz) learned she had won; at t our knowledge was such that it did not eliminate the possibility that she had lost the lottery.⁵⁶

This account of what seems appropriate for the lucky subject to say fits nicely with the knowledge norm of assertion:⁵⁷ the reason why it is appropriate for the lucky subject to assert that he does not know that she will win the lottery, in response to the question “Will you win the lottery?” is because, given her evidence, not only does she fail to know that she will win, but she can easily know that she does not know whether she will win or not, since she knows the lottery odds and she knows that the lottery is fair.

Note too that Chance can swiftly explain why Vincent, in the buried treasure case, is lucky to have found the treasure: given Vincent’s evidence at any time before he found the treasure, he was clearly not in a position to know the treasure was buried where it was - as a matter of fact, he was not even in a position to know there was a treasure anywhere! It is also blatantly obvious that it is appropriate for Vincent to say, in response to the question “Will you find any buried treasure in the place you select to plant your roses?” that he does not know that he will, since he has no reason to think anyone buried any treasure anywhere.

Chance can also make sense of the common feeling of surprise that usually accompanies situations in which one is the subject of good or bad luck. No one will ever see a winner of a fair lottery who is *not* surprised when she learns that she won the lottery. Everyone *is* very (happily) surprised when a car passenger miraculously survives a head on collision with a truck. And so on and so forth. In all these cases, the surprised subject (be they the lucky person or some observer who is in a similar epistemic position as the lucky person was in right before the

⁵⁶This applies also in a case in which Claire the clairvoyant foresees her lottery win. In that case, when Claire says to herself “I am lucky to win the lottery”, her statement is elliptical for “Given *k*, I am lucky to win the lottery” where “*k*” refers to the body of knowledge Claire had *before* she foresaw her win.

⁵⁷Williamson [2000].

relevant event took place) is not in a position to know that the relevant event will hold. This creates a justified expectation in the lucky subject that it is likely that the event will not hold - for all she knows, the event might not hold. But, then, when the event does hold, this expectation is frustrated, and the natural reaction is for one to be surprised that the event held.

So, Chance is a plausible candidate for the role of generic luck.

If we take seriously the anti-luck epistemologist's suggestion that epistemic luck is just the special case of generic luck, then I can propose the following version of epistemic luck:

(Gettier-Luck / Epistemic Luck) S is Gettier-lucky to believe truly that p at t only if S is not in a position to know at t that p is true at a later time t_n .

This is a necessary condition on a belief being true as a matter of luck. For a true belief to be Gettiered, it also needs to satisfy the other conditions I suggested in chapter 4 - namely, the true belief depends essentially on some unknown proposition from which it is deduced.

Unlike the safety-based version of epistemic luck, this principle does not make reference to the proposition S believes as being contingent or necessary and, thus, has no difficulty dealing with either kind of content. No codicil mentioning the method in which the belief is acquired is needed. Remember also what I said above about S being a position to know that p: if S is in a position to know that p, then S has the evidence she needs to come to know that p, and were S to believe that p on the basis of that evidence, nothing would prevent her from knowing that p.

So, in the modified Ford case above, Smith was lucky to have believed truly that either Jones owns a Ford or either Brown is in Barcelona or ϕ , (where " ϕ " is a mathematical truth Smith believes only because it "sounds good") because, given his evidence, he is not in a position to know that proposition is true. A

similar thing can be said about the original Gettier cases, the Barn-Façade case, and whatever other case one cares to classify as a case in which epistemic luck is present.

Epistemic luck is also compatible with different explanations of why the subject in a particular situation is not in a position to know. It is compatible, for example, both with a no-false-grounds account of why the subject is not in a position to know and with a safety account of why the subject is not in a position to know. So, one could, in principle, say that Henry is lucky to believe truly that that is a barn, because he is not in a position to know that it is, and then either add “and Henry is not in a position to know because his belief is unsafe” or “and Henry is not in a position to know because his reasoning essentially involves a falsehood.”

In order to assess this definition of Gettier-luck is correct, let us look at some cases and see how it accounts for the luck involved in those cases.

Consider Gettier’s coin case. Luck intervenes in that case because Smith is not in a position to know that the person who will get the job has ten coins in his pocket.

One might worry that this gets things backwards. “We want to appeal to the notion of luck to explain why we fail to know in Gettier cases, but now you are appealing to the lack of knowledge in Gettier situations to explain why there is luck in Gettier cases.”

There are at least two reasons why inverting the order of explanation in this way should not worry us. First, we always knew Gettierized agents fail to know because their epistemic position is, in some sense, *not good enough*. Some epistemologists called “luck” the factor weakening their epistemic position. The problem is that, luck itself is partially an epistemic phenomenon. Non-epistemic accounts of luck are either false (e.g., the counterfactual weakness condition), not sufficient to account for all instances of luck (e.g., the lack of direct control condition), or

both. Since we now know luck is partially an epistemic phenomenon, we will have to look some place else for a full characterization of the epistemic position of Gettierized agents. (As I showed in chapter 3, I can explain why Gettiered subjects fail to know without appealing to the notion of epistemic luck.) This brings us to the second reason why inverting the order of explanation is not worrisome. The result that luck is partially an epistemic notion is a problem only for views that take the absence of Gettier-luck to be not merely incompatible with knowledge but to be constitutive of knowledge.⁵⁸ I make no such claim here, however. This is not a good outcome for Pritchard (some would say “it is bad luck for Pritchard”!), since his analytical project proposes that a condition eliminating luck is necessary for a correct analysis of knowledge. My proposed definition of luck would make the definition of “knowledge” circular and, thus, inadequate for the analytical project.

As a result, not only does Hetherington’s case for Gettier-lucky fail, but we have also now seen why it *could not* have succeeded: S would be lucky with respect to believing truly that p only if S were not in a position to know that p, hence, in order to Gettier-luckily know that p, S would have to fail to be in a position to know that p! Therefore, I can safely conclude that Hetherington’s case for Gettier-lucky knowledge is ultimately incoherent, even if (and that is a big “if”!) prima facie plausible.

Finally, this account of epistemic luck does not go well with the version of anti-luck epistemology I am considering, and the reason is obvious: epistemic luck, when plugged into a proposed definition of the concept of knowledge, makes that definition circular. This result does not mesh well with Duncan Pritchard’s analytical project. So much the worse for the analytical project. Chance and Gettier-Luck/Epistemic Luck account for all the cases of luck I am interested in, including the cases that caused trouble for the safety-based account of those

⁵⁸Steglich-Petersen [2010, p.376-7].

notions. I have plenty of reason to reject the safety-based account of luck; even if that costs me the analytical project. As far as I can see, the failure of the analytical project in showing that a safety-based account of epistemic luck is a necessary condition on “knowledge” is more inductive evidence for the claim that knowledge is not an analyzable concept.

Bibliography

- Peter Achinstein. *The Book of Evidence*. Oxford University Press, 2003.
- Jonathan Adler. *Belief's Own Ethics*. MIT Press, 2002.
- Aristotle. *Metaphysics*, volume The Complete Works of Aristotle v.1. Princeton University Press, 1991.
- Aristotle. *Posterior Analytics*. Oxford University Press, 2nd edition, 1994.
- David M Armstrong. *Belief, Truth and Knowledge*. Cambridge University Press, 1973.
- David M Armstrong. *Truth and Truthmakers*. Cambridge University Press, 2004.
- Matthew Benton. Two More for the Knowledge Account of Assertion. *Analysis*, 71(4):684–687, 2011.
- Jessica Brown. Impurism, Practical Reasoning and the Threshold Problem. *Noûs*, 48(1):179–192, 2014.
- Roderick Chisholm. *Theory of Knowledge*. Prentice Hall, 1977.
- Michael Clark. Knowledge and Grounds: A Comment on Mr. Gettier's Paper. *Analysis*, 1963.
- E.J Coffman. Thinking about luck. *Synthese*, 158:385–398, 2007.
- E.J Coffman. Does Luck Exclude Control. *Australasian Journal of Philosophy*, 87(3):499–504, September 2009.

- E.J. Coffman. Misleading Dispositions and the Value of Knowledge. *Journal of Philosophical Research*, 35:241–258, 2010.
- Juan Comesaña and Holly Kantin. Is Evidence Knowledge? *Philosophy and Phenomenological Research*, 80(2):447–454, 2010.
- Claudio de Almeida. What Moore’s Paradox is About. *Philosophy and Phenomenological Research*, 62(1):33–58, January 2001.
- Claudio de Almeida. Benign falsehoods: a Lesson About Inferential Knowledge. *Unpublished manuscript*, 2003.
- Claudio de Almeida. On Useful Falsehoods, 2004. URL <http://certaindoubts.com/?p=195>.
- Claudio de Almeida. *Moorean absurdity: an Epistemological Analysis*. Oxford University Press, 2007.
- Claudio de Almeida. Epistemic closure, Skepticism and Defeasibility. *Synthese*, 2011.
- Keith DeRose. Assertion, Knowledge, and Context. *The Philosophical Review*, 111:167–203, 2002.
- Keith DeRose. *The Case for Contextualism: volume one*. Oxford University Press, 2009.
- René Descartes. Rules for the Direction of the Mind. In Dugald Murdoch John Cottingham, Robert Stoothoff, editor, *The Philosophical Writings of Descartes*, volume I. Cambridge University Press, 1985.
- Dylan Dodd. Against Fallibilism. *Australasian Journal of Australasian Journal of Philosophy*, 89(4):665–685, December 2011.

- Fred Dretske. Epistemic Operators. *The Journal of Philosophy*, 67(24):1007–1023, 1970.
- Fred Dretske. Conclusive Reasons. *Australasian Journal of Philosophy*, 49(1): 1–22, May 1971.
- Fred Dretske. *Knowledge and the Flow of Information*. MIT Press, 1981a.
- Fred Dretske. The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40 (3):363–378, 1981b.
- Fred Dretske. Is Knowledge Closed Under Known Entailment? The Case Against Closure. In Matthias Steup and Ernest Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, 2005.
- Jonathan St. B. T. Evans and Keith Frankish. *In Two Minds: Dual Processes and Beyond*. Oxford University Press, 2009.
- Haskell Fain and A. Phillops Griffiths. On Falsely Believing that One Doesn't Know. *American Philosophical Quarterly: Monograph Series*, 1972.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, 2012.
- Richard Feldman. An Alleged Defect in Gettier Counter-Examples. *Australasian Journal of Philosophy*, 52(1):68–9, May 1974.
- Richard Feldman. *Epistemology*. Pearson, 2002.
- Branden Fitelson. Strengthening The Case For Knowledge From Falsehood. *Analysis*, 70(4):666–669, October 2010.
- Richard Foley. *When Is True Belief Knowledge?* Princeton University Press, 2012.

- Tamar Szabó Gendler and John Hawthorne. The Real Guide to Fake Barns. *Philosophical Studies*, 124:331–352, 2005.
- Edmund Gettier. Is Justified True Belief Knowledge? *Analysis*, pages 121–123, 1963.
- Carl Ginet. Knowing Less by Knowing More. *Midwest Studies in Philosophy*, 5: 151–162, 1980.
- Alvin Goldman. A Causal Theory of Knowing. *The Journal of Philosophy*, 1967.
- Alvin Goldman. Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73(20):771–791, 1976.
- Alvin Goldman. *Epistemology and Cognition*. Harvard University Press, 1986.
- Alvin Goldman. *Knowledge in a Social World*. Oxford University Press, 1999.
- Alvin Goldman. Philosophical Naturalism and Intuitional Methodology. In *Proceedings and Addresses of the American Philosophical Association*, 2010.
- Alvin Goldman. *Reliabilism and Contemporary Epistemology*. Oxford University Press, 2012.
- Alvin Goldman and Erik Olsson. Reliabilism and the value of Knowledge. In Adrian Haddock, Alan Millar, and Duncan Pritchard, editors, *Epistemic Value*. Oxford University Press, 2009.
- Paul Grice. *Studies in the Way of Words*. Harvard University Press, 1991.
- Gilbert Harman. *Thought*. Princeton University Press, 1973.
- Gilbert Harman. Reasoning and Evidence One Does Not Possess. *Midwest Studies in Philosophy*, 5:163–182, 1980.
- Gilbert Harman. *Change In View*. MIT Press, 1986.

- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004.
- John Hawthorne. The Case for Closure. In Mathias Steup and Ernest Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, 2005.
- John Hawthorne and Jason Stanley. Knowledge and Action. *Journal of Philosophy*, 105(10):571–590, 2008.
- Adrian Heathcote. Truthmaking and the Gettier Problem. In Stephen Hetherington, editor, *Aspects of Knowing: Epistemological Essays*. Elsevier Science, 2006.
- Stephen Hetherington. *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. Oxford University Press, 2001.
- Stephen Hetherington. *How to Know*. Blackwell, 2011.
- Stephen Hetherington. There Can be Lucky Knowledge. In Mathias Steup, John Turri, and Ernest Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, 2nd edition, 2013.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- Keith J. Holyoak and Robert G. Morrison. Thinking and Reasoning: A Readers Guide. In Keith J. Holyoak and Robert G. Morrison, editors, *The Cambridge Handbook of Thinking and Reasoning*, chapter 1, pages 1–9. Cambridge University Press, 2005.
- Frank Jackson. A Causal Theory of Counterfactuals. *Australasian Journal of Philosophy*, 55(1):3–21, May 1977.
- Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- Immanuel Kant. *Critique of Pure Reason*. Macmillan, 1950.

- Thomas Kelly. Evidence, August 2006. URL <http://plato.stanford.edu/entries/evidence/>.
- Peter D. Klein. A Proposed Definition of Propositional Knowledge. *The Journal of Philosophy*, 1971.
- Peter D. Klein. Misleading "Misleading Defeaters". *The Journal of Philosophy*, 76(7):382–386, July 1979.
- Peter D. Klein. *Certainty: A Refutation of Scepticism*. University of Minnesota Press, 1981.
- Peter D. Klein. Skepticism and Closure: Why the Evil Genius Argument Fails. *Philosophical Topics*, pages 212–236, 1995.
- Peter D. Klein. Warrant, Proper Function, Reliabilism and Defeasibility. In Jonathan L. Kvanvig, editor, *Warrant and Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Rowman and Littlefield, 1996.
- Peter D. Klein. Useful False Beliefs. In Quentin Smith, editor, *New Essays in Epistemology*. Oxford University Press, 2008.
- Peter D. Klein. What Makes Knowledge the Most Highly Prized Type of Belief? In T. Black and K. Becker, editors, *The Sensitivity Principle in Epistemology*, pages 152–169. Cambridge University Press, 2012.
- Keith Allen Korcz. Recent Work on the Basing Relation. *American Philosophical Quarterly*, 34(2):171–191, 1997.
- Keith Allen Korcz. The Epistemic Basing Relation, 2010. URL <http://plato.stanford.edu/archives/spr2010/entries/basing-epistemic/>.
- Hilary Kornblith. *Knowledge and its Place in Nature*. Oxford University Press, 2002.

- Saul Kripke. *Philosophical Troubles*. Oxford University Press, 2011.
- Jonathan Kvanvig. *Warrant and Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*. Rowman and Littlefield, 1996.
- Jennifer Lackey. Pritchard's Epistemic Luck. *Philosophical Quarterly*, 56:284 – 289, 2006.
- Jennifer Lackey. Norms of Assertion. *Noûs*, 41(4):597–626, 2007.
- Jennifer Lackey. What Luck Is Not. *Australasian Journal of Philosophy*, 86:255 – 267, 2008.
- Keith Lehrer. Knowledge, Truth, and Evidence. *Analysis*, 25:168–175, 1965.
- Keith Lehrer. *Theory of Knowledge*. Westview Press, 2nd edition, 1990.
- Keith Lehrer and Jr. Thomas Paxson. Knowledge: Undefeated Justified True Belief. *The Journal of Philosophy*, 66(8):225–237, April 1969.
- Neil Levy. What, and Where, Luck is: A Response to Jennifer Lackey. *Australasian Journal of Philosophy*, 87:489–497, 2009.
- David Lewis. Counterfactual Dependence and Time's Arrow. *Noûs*, 13(4):455–476, 1979.
- John Locke. *An Essay Concerning Human Understanding*. Oxford University Press, 1975.
- William Lycan. Evidence One Does Not Possess. *Australasian Journal of Philosophy*, pages 114–26, 1977.
- William Lycan. The Gettier Problem Problem. In Stephen Hetherington, editor, *Epistemology Futures*. Oxford University Press, 2006.

- Ishani Maitra and Brian Weatherson. Assertion, Knowledge, and Action. *Philosophical Studies*, 149(1):99–118, 2010.
- Norman Malcolm. Knowledge and Belief. *Mind*, 61(242):178 – 189, 1952.
- Brian McLaughlin. Is Role-Functionalism Committed to Epiphenomenalism? *Journal of Consciousness Studies*, 13:39–66, 2006.
- G. E. Moore. Moore’s Paradox. In Thomas Baldwin, editor, *G. E. Moore: Selected Writings*. Routledge, 1993.
- Thomas Nagel. Moral Luck. In *Mortal Questions*. Cambridge University Press, 1979.
- Avram Hiller; Ram Neta. Safety and Epistemic Luck. *Synthese*, 158:303–313, 2007.
- Robert Nozick. *Philosophical Explanations*. Belknap Press of Harvard University Press, 1981.
- L A Paul and Ned Hall. *Causation: A User’s Guide*. Oxford University Press, 2013.
- Alvin Plantinga. *Warrant: The Current Debate*. Oxford University Press, 1993.
- John Pollock and Joseph Cruz. *Contemporary Theories of Knowledge*. Rowman and Littlefield, 1999.
- Duncan Pritchard. *Epistemic Luck*. Oxford University Press, 2005.
- Duncan Pritchard. *Knowledge*. Palgrave Macmillan, 2009.
- Duncan Pritchard. Anti-Luck Virtue Epistemology. *The Journal of Philosophy*, 2012.
- Bertrand Russell. *The Problems of Philosophy*. Oxford University Press, 1912.

- Bertrand Russell. *Human Knowledge: Its Scope and Limits*. Routledge, 1948/2009.
- Israel Scheffler. *Conditions of Knowledge*. Scott, Foresman and Company, 1965.
- Assaf Sharon and Levi Spectre. Dogmatism Repuzzled. *Philosophical Studies*, 148:307–321, 2010.
- Robert Shope. *An Analysis of Knowing*. Princeton University Press, 1983.
- Roy Sorensen. Dogmatism, Junk Knowledge, and Conditionals. *Philosophical Quarterly*, 38:433–454, 1988.
- Roy Sorensen. *Blindspots*. Oxford University Press, 1998.
- Roy Sorensen. *A Brief History of the Paradox*. Oxford University Press, 2003.
- Roy Sorensen. Epistemic Paradoxes. In *Stanford Encyclopedia of Philosophy*. Stanford University, 2012. URL <http://plato.stanford.edu/archives/spr2012/entries/epistemic-paradoxes/>.
- Ernest Sosa. How Must Knowledge Be Modally Must Knowledge Be Modally Related to What Is Known? *Philosophical Topics*, 1999.
- Ernest Sosa. *A Virtue Epistemology*. Oxford University Press, 2007.
- Ernest Sosa. *Knowing Full Well*. Princeton University Press, 2011.
- Daniel Statman, editor. *Moral Luck*. State University of New York Press, 1993.
- Asbjørn Steglich-Petersen. Luck as an Epistemic Notion. *Synthese*, 176:361–377, 2010.
- Jonathan Sutton. *Without Justification*. MIT Press, 2007.
- John Turri. The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, 89(11):37–45, March 2011.

- Peter Unger. An Analysis of Factual Knowledge. *The Journal of Philosophy*, 1968.
- Peter Unger. *Ignorance: A Case for Scepticism*. Oxford University Press, 1975.
- Jonathan Vogel. Are There Counterexamples to the Closure Principle? In *Doubting: Contemporary Perspectives on Skepticism*. Kluwer, 1990.
- Ted A. Warfield. Knowledge from Falsehood. *Philosophical Perspectives*, 19: 405–416, 2005.
- Brian Weatherson. What Good are Counterexamples? *Philosophical Studies*, 115 (1):1–31, 2003.
- Brian Weatherson. Only Knowledge is Evidence, 2011. URL <<http://tar.weatherson.org/2011/04/22/only-knowledge-is-evidence/>>.
- Matthew Weiner. Must We Know What We Say? *Philosophical Review*, 114(2): 227–251, 2006.
- Bernard Williams. *Moral Luck*. Cambridge University Press, 1981.
- John Williams and Mitchell Green. *Moore's Paradox: new essays on belief, rationality and the first person*. Oxford University Press, 2007.
- Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, 2000.
- Timothy Williamson. Knowledge First Epistemology. In Sven Bernecker and Duncan Pritchard, editors, *The Routledge Companion to Epistemology*, chapter 20, pages 208–218. Routledge, 2011.
- Timothy Williamson. Gettier Cases in Epistemic Logic. *Inquiry*, 56(1):1–14, 2013.
- Crispin Wright. (Anti-)Sceptics Simple and Subtle: G.E. Moore and John McDowell. *Philosophy and Phenomenological Research*, 65:330–348, 2002.

Linda Zagzebski. The Inescapability of Gettier Problems. *The Philosophical Quarterly*, 1994.

Linda Zagzebski. What is Knowledge? In John Greco and Ernest Sosa, editors, *The Blackwell Guide to Epistemology*. Blackwell, 1999.