



# Knowledge Closure and Knowledge Openness

A Study of Epistemic Closure Principles

Levi Spectre

©Levi Spectre, Stockholm 2009  
ISBN 978-91-7155-973-9  
Printed in Sweden by US-AB, Stockholm 2009

# Contents

Preface.....	vi
Chapter 1: An Introduction to the Closure Debate in Epistemology.....	9
1.1 The Closure Debate .....	14
1.2 Layout of the Argument .....	17
1.3 Closure and its Motivation .....	19
1.4 What is the Closure Principle? .....	23
1.5 Debating Closure.....	29
1.5.1 Hawthorne's Arguments.....	29
1.5.2 Nozick on Closure.....	32
1.5.3 Dretske on Closure and Heavyweight Propositions .....	35
1.5.4 Section Summary.....	37
Chapter 2: Scepticisms and Closure.....	38
2.1 Scepticism and Closure.....	38
2.1.1 Cartesian Scepticism, Moorism, and what Comes In-Between.....	41
2.2. Cartesian Scepticism.....	43
2.2.1 Immodest Scepticism.....	44
2.2.2. Modest Scepticism.....	47
2.3 Mundane Scepticism .....	53
2.4 Live Scepticism.....	58
Chapter Summary .....	59
Chapter 3: The Kripke Dogmatism Puzzle and Epistemic Ascent .....	62
Chapter Three Outline .....	62
3.1 The Threat of Dogmatism.....	63
3.1.1 Junk Conditionals and Junk Knowledge .....	66
3.1.2 Junks or Knowledge? Dogmatism Repuzzled.....	68
3.1.2.1 A Short Excursion to the Bahamas.....	77
3.2 Epistemic Ascent.....	79
3.2.1. Dogmatism and Easy Knowledge and Bootstrapping .....	81
3.2.2 Epistemic Ascent and Knowledge of True Belief .....	83
Chapter Summary .....	87
Chapter 4: Evidence and Open Knowledge.....	90
Chapter outline .....	90
4.1 The Watch Case and its Probabilistic Analysis. ....	91

4.1.1. Evidence and Probabilities.....	91
4.2. Probabilistic Argument for Evidence and Knowledge Openness .....	96
4.3 Evidence Openness from Principles.....	99
4.3.1. Equivalence, Consistency and Addition .....	104
4.3.2. Equivalence and Distribution.....	109
4.3.3. EAD, EDIS and the Logic of Evidence .....	110
4.3.4. Carnap's Matrix .....	111
Chapter Summary .....	114
<b>Chapter 5: Open Knowledge – Costs and Benefits .....</b>	<b>119</b>
5.1. From Evidence to Knowledge.....	119
5.1.1. Premise (V) of the Open Knowledge Argument .....	124
5.1.1.1. An Objection to the Use of (V).....	127
5.1.1.2. An Objection to (NED) .....	132
5.2. Knowledge Without Evidence.....	137
5.3. Knowledge as Evidence .....	140
5.3.1. Knowledge as Evidence: Jeffrey Conditionalization.....	142
5.3.2. Knowledge as Evidence: Standard Conditionalization.....	144
5.3.3. Knowledge as Evidence: Williamson's Account .....	148
5.3.3.1. Safety and Chance .....	149
5.3.3.2. Lottery Propositions .....	152
5.3.3.3. The Principal Principle and Practical Deliberation .....	158
5.3.3.4. Fallibilism .....	162
5.3.4. Knowledge as Evidence: Conclusion .....	165
5.4. Evidential Knowledge .....	166
5.4.1. The Benefits of Epistemic Openness .....	166
5.4.2. Denying Closure: Not as Bad as You Think.....	168
Chapter Summary .....	175
<b>Chapter 6: Concluding Remarks .....</b>	<b>179</b>
<b>Appendix: A Short Overview of Bayesianism.....</b>	<b>186</b>
A Bayesian Framework .....	187
Convergence Theorems.....	192
Deep convergence .....	197
<b>Bibliography .....</b>	<b>199</b>

# Chapter 1: An Introduction to the Closure Debate in Epistemology

*Those are my principles, and if you don't like them... well, I have others.*

*Groucho Marx*

We know that *Socrates was troubled* if it is true that *Socrates lived* and true that, *if he lived, he was troubled*. We know this since we have excellent reason to believe that basic deductive inferences patterns are truth-preserving. If the premises of valid inference patterns are true, we have a logical guarantee that the conclusion is true as well. Such is the case with the *modus ponens* inference pattern, a candidate as good as any for being valid. The study of logic deals, among other things, with the identification of such patterns as well as with proof of their validity. Yet, when propositional attitudes govern the scope of the premises of provably valid inferences, even with regard to *factive* attitudes<sup>1</sup> such as knowledge, there is no *logical* guarantee that the conclusion will also fall within their scope ( $K(p) \wedge K(p \rightarrow q) \not\vdash K(q)$ ). Regardless of the logical guarantee, however, the prevalent inclination is to think that these inferences actually do preserve truth from premises to conclusion. If I know that *grass is green*, and if I know that *if something is green it has a color*, I know that *grass is colored*. We cannot prove that I know this, but a denial would seem very odd. What else would it take for me to know that *grass is colored*?

Not all cases inspire the same confidence as this grass example. My wife knows, since we have talked about this extensively, that *if she gets extended time-off from her job, we will go to the Bahamas*. She also knows, since her trusted boss has told her as much, that *in the case of her husband's unfortunate death, she will receive extended time-off*. Does she then know that, *if I die, we will go to the Bahamas*?<sup>2</sup> Perhaps she does, or perhaps when she thinks about this, knowledge no longer governs the scope of the premises of this valid inference. The point is that, even though the truth of the premises guarantees the truth of the consequence, knowledge (though tightly con-

---

<sup>1</sup> Factive propositional attitudes are those that entail the truth of the propositions they govern.

<sup>2</sup> The example is inspired by Sorensen's tennis example (1988b: 405). He uses it to make a point about slippery-slope arguments.

nected with truth) is not obviously preserved. Yet something is irregular about my wife having these two items of knowledge, even before we placed them as premises of the inference. What seems odd, among other things, is precisely that, assuming she has no intention of taking her dead husband to the Bahamas but knows the two conditionals, it would seem that she should count as knowing that *if I die, we are going to the Bahamas*. The oddness, then, seems to count against the assumption that she could know the two conditionals to begin with. So, despite any trepidation evoked by my wife's morbid inference, the idea that if one knows that  $p$ , and one knows that *if  $p$  then  $q$* , one will know that  $q$ , seems rather basic. It seems that we would be better off denying she knows the premises than denying that she knows the conclusion (while knowing the premises).

Even with regard to *non-factive* attitudes such as belief, we are inclined to think that valid inferences allow us to lend equal or greater credence to their conclusions as we lend to their premises. For good reason, I believe that *you will soon be calling me*. I also have good reason to believe that *if you do, you are at home*. Do I have reason to believe that *you are at home*? Even though we might have less confidence regarding this question as we have regarding truth or knowledge, the intuitive straightforward answer is affirmative. If this is true regarding attitudes that do not entail truth, then regarding knowledge, which does entail truth, our inclination seems to be on firm ground. The temptation is strong to think that this kind of knowledge preservation through valid inference *always* holds, that is, it should be considered to be a principle: For all propositions  $p$  and  $q$  and subjects  $S$ , if  $S$  knows that  $p$  and knows that  $p$  implies  $q$ , then  $S$  knows that  $q$ .<sup>3</sup> This claim is commonly known as *the principle of closure*.<sup>4</sup>

The view that I will defend throughout this manuscript is that there is good reason to question the validity of the closure principle, that is, that there is good reason to think it does not generally hold. In fact, though I do not deny that closure is intuitive, I will claim that the motivation I have given so far for the principle of closure is misguided. But motivation and intuition aside, the reason to suspect knowledge closure is that, unlike truth, knowledge involves properties that may not be preserved across logically

---

<sup>3</sup> This is a rough statement of the principle. Many epistemologists reject it as stated but accept weaker variants.

<sup>4</sup> The principle of closure for knowledge has several other names: "Deductive closure," "knowledge closure," "principle of transmission of knowledge by deduction." Most common, I think, it is simply "the principle of closure" which should be considered here to relate to knowledge unless stated otherwise. For instance, I might say that I am concerned with justification closure, evidence closure, or rational belief closure, in a certain context.

What it means for a claim to hold as a principle is that it holds with absolute generality. Usually its statement is given in terms of a conditional statement such that by necessity if the antecedent holds, the consequent will hold. I do not think that anything turns on whether the principle is stated as a universally quantified conditional or as a necessary, or strict, conditional.

valid inferences. Though it may not seem apparent, I can have evidence for the truth of a proposition  $p$ , and know that  $p$  implies  $q$ , and yet my evidence for  $p$  will not count in favor of  $q$ . If knowing requires having good enough evidence, I may have evidence for the premise of a valid inference – enough to know it to be true – and yet fail to have the kind of evidence I need in order to count as knowing its conclusion. This, at least in rough outline, is the reason why I think we need to be suspicious of the claim that we always know the truth of conclusions of valid inferences if we know the truth of their premises.<sup>5</sup> I question, then, the claim that when I infer a proposition from something I know and do this correctly, I know that the proposition I inferred is true. The thesis is that the failure of evidence to be preserved from premises to conclusion of valid inferences gives us good reason to think that closure fails.

I must admit at the outset that the closure principle is highly intuitive, and why this is so will concern me in the current and later chapters. And yet, the main issue is to discern theoretical considerations that bear on the main question while steering away from intuitions. Before discussing these issues in detail, I would first like to ascertain what is at stake in determining whether closure is a valid principle. In this regard, at least five issues are worthy of mention.

First, as the examples above suggest, inferences from known premises are central in everyday situations. When Sherlock Holmes investigates crime scenes and questions suspects and witnesses, his goal is to know the identity of the perpetrator(s) of a crime. He seldom, if ever, gains direct access to this knowledge (though after the identity of the criminal becomes known, the reader often receives more direct confirmation). Rather, his method is to extend his knowledge by inferring from knowledge he has gained through interview, inspection, and extensive background information. Though the character of Holmes is fictional, the method is real and has a central place in science and everyday reasoning. If closure is not valid, we might need to reevaluate the confidence we place in the knowledge we regularly take people to possess. If knowledge is not always extendable by proper inference, what good is knowledge, and indeed, what good is proper inference?

Second, a major challenge to the idea that we know many propositions to be true is presented in the form of Cartesian skepticism. Skeptics of this variety typically argue that, if I know that  $q$  follows from  $p$  (theoretical knowledge that not only do they not dispute but indeed rely on), if I don't know that  $q$ , I don't know that  $p$  either. This is the closure principle stated in a *re-*

---

<sup>5</sup> Though my reasons for making the distinction are different, like John Hawthorne (2004a) I will later claim that there is a major distinction that is important for the debate concerning closure, namely, the distinction between valid inferences with more than one premise, and valid inferences that have only one. But for my purposes the distinction does not matter since the reason I offer to suspect closure - the intransitivity of evidence - does not turn on this distinction. The account, then, offers the same response to both forms of closure principle.

*ductio* argument. Although they have no proof, they typically claim that we do not know that we are not being systematically deceived in such a way that our beliefs, which we typically take to be true, are really false. Such would be my predicament, for instance, if I were a brain in a vat being fed sensory input that made it appear to me that I had a body when, in reality, I am bodyless. Many other things that I currently believe would be false; as a matter of fact, I would hardly have any true beliefs about the environment I would falsely believe I occupy (though as Descartes famously argued, I would still know, since I would know that I am thinking, that I exist in some environment). If the closure principle is valid, the skeptic argues, I am ignorant of a vast body of beliefs I take to be part of my knowledge. This argument collapses if knowledge is not closed, provided, that is, that closure fails in the right way, i.e., it fails with regard to the cases that the skeptics utilize. If closure indeed fails in the right way, we can meet this skeptical challenge if we can say something about why and how it fails in this way.

Third, the closure principle plays a vital role in many central issues and paradoxes of epistemology besides Cartesian skepticism, including the Lottery and the Preface paradoxes, the Easy Knowledge problem (or bootstrapping problem), the self knowledge and semantic externalism problem, some of the Gettier cases, the Surprise Examination paradox (or hangman paradox), the warrant transfer questions, the Knowability paradox, the Knower paradox, Kripke's epistemic dogmatism puzzle, and Kripke's puzzle of belief.<sup>6</sup> These problems have closure of one type or another in full focus or at least as one of the premises necessary to generate them.<sup>7</sup> Knowledge openness<sup>8</sup> - the idea that knowledge is not closed under proper inference - can both dissolve<sup>9</sup> at least some of these problems as well as explain why they have such a strong grip on us (by appeal to closure's intuitive tow). If closure is not valid, we might have been looking at the wrong issues trying to solve problems that are only apparent.

---

<sup>6</sup> Additional issues include Live Skepticism, Mundane Skepticism (chapter 2) and Epistemic Ascent (chapter 3).

<sup>7</sup> The connection between closure and Cartesian skepticism is widely noted and will be considered in chapter 2. For the involvement of closure in Cohen's Easy Knowledge problem see his (2002), (2005), and Hawthorne's (2004a) as well as Chapter 3. For its involvement in the Bootstrapping problem see Vogel (2000, 2007) (and Chapter 3). For closure in Gettier cases see his (1963). For the Knowability paradox and closure Brogaard and Salerno (2006). Surprise Exam and the Knowability paradoxes Kaplan and Montague (1960). The connection between the Semantic-externalism/self-knowledge problem and closure is noted in Brown (2004). Kripke's epistemic dogmatism and closure see Cargile (1995), Sharon and Spectre (forthcoming), and Chapter 3. For the involvement of closure in Kripke's belief puzzle, see Frances (1999).

<sup>8</sup> The term "open knowledge" was coined (as far as I know) by Nozick (1981: 208).

<sup>9</sup> More carefully, whether or not non-closure can dissolve a given paradox depends on how closure fails. It may very well be that many of the problems remain since knowledge is closed with regard to cases that bring about a given paradox. I will not explore all the problems I mention so I will simply say that I think the hangman paradox does not turn on closure failure, Kripke's puzzle of belief I am less certain about.



Fourth, the centrality of closure is also evident when we take a brief look at some of the main views advanced in recent (and less cent) epistemology. Advocates often take the validity of closure and the lack of tenability of skepticism as premises of arguments for their respective views. In other words, it is common to take a non-skeptical approach that allows knowledge to be closed as a dual objective, and this dual objective has become a central desideratum in contemporary epistemology. In this respect, these views are what we may call *closure driven*. Such are some of the main arguments for epistemic Contextualism (Stewart Cohen 1988, David Lewis: 1996); for Dogmatism (G. E. Moore: 1959, Jim Pryor: 2000); for epistemic Relativism (John McFarlane: 2005); for Contrastivism (Jonathan Schaffer: 2004); and for Subject Sensitive Invariantism (John Hawthorne: 2004a, Jason Stanley: 2005).<sup>10</sup> Skepticism, which has already been mentioned, is also closure-driven but has only one theoretical objective, not two.

The fifth and final key issue in which closure is involved concerns epistemic logic. Since any normal modal logic has  $\mathbf{K}$  ( $=\Box(\alpha\rightarrow\beta)\rightarrow(\Box\alpha\rightarrow\Box\beta)$ ) as an axiom,<sup>11</sup> non-closure of knowledge means that no normal modal logic for knowledge is tenable. Trying to idealize epistemic agents (something we must do in any case, even if closure holds) will not result in a normal modal logic for knowledge unless we idealize along the lines on which closure fails.<sup>12</sup> (Perhaps epistemic logic can be limited to sub-domains such as mathematical knowledge, which are normal even if normal epistemic modal logic is generally ill fitted for knowledge.) The issue at hand, therefore, is whether epistemic logic is even possible.

---

<sup>10</sup> Note that these views can be combined with non-closure. In fact an open knowledge contextualist view has been advanced by Mark Heller (1999). I am inclined to think that a defensible view can be offered along the lines of an open knowledge Subject Sensitive Invariantist view as well. This is to say that the practical environment of a thinker can determine what this thinker knows, but this shift does not have to vary systematically with what closure predicts. Some comments in this manuscript will be relevant for such a proposal.

<sup>11</sup> “One common way of presenting a particular normal modal system is by stipulating that it contains, in addition to the PC-tautologies, a certain specified collection of other formulae. These formulae, which must contain  $\mathbf{K}$  but otherwise can be any collection we choose, are then said to be axioms of the system...” (G.E. Hughes and M. J. Cresswell, 1968: 5) In other words, an alternative label to this essay could be: “Does epistemic modal logic rest upon a mistake?” I prefer the present title since I do not discuss the matter of what logic adequately captures central features of epistemic states and dynamics. Nevertheless, much of what will follow be relevant. I owe this point to Assaf Sharon and Daniel Rönnedal.

<sup>12</sup> If  $\mathbf{K}$  is taken as an axiom of epistemic modal logic, the following is a corollary: if  $\alpha\vdash\beta$ , then  $(\mathbf{K}\alpha\rightarrow\mathbf{K}\beta)$ . Without idealization, most epistemologists will not accept  $\mathbf{K}$  (I have Lewis in mind as a possible exception, see Hawthorne (2002) for a detailed argument). If we idealize along the lines of conclusive evidence (i.e. modal logic for all cases where knowledge is based on evidence that entails the known proposition), a weaker  $\mathbf{K}$ -like axiom would be valid even by the open knowledge advocate’s lights. From this perspective, this is why there is no problem (as far as I can see) with mathematical epistemic logic. The evidence, which I take to be proofs in this case, entails the known mathematical propositions.

These are the reasons that make the closure of knowledge such a central question.<sup>13</sup> Only some of these reasons have been fully acknowledged but the issue has been debated, though less extensively than one might expect. Before turning to consider what the closure principle is, its basic motivation, and what I take to be the issue on which the question turns, I wish to say a few things about how the debate has been conducted in the literature, to be followed by a statement of my aims in this manuscript and how I will try to achieve them.

## 1.1 The Closure Debate

*What you know to follow from your knowledge is something you know is true.* This is the basic and unquestionably intuitive idea that formulations of closure are meant to express. Indeed, an air of paradox beclouds a failure to know consequences you know to follow from your knowledge. Unsurprisingly, then, a large majority of contemporary epistemologists consider one or another weak version of this principle to be valid. Some (e.g. Stewart Cohen) view it as “something like an axiom about knowledge,” (2005: 312) others, like Richard Feldman, think that doubting its validity is a very bad idea – “one of the least plausible ideas to come down the philosophical pike in recent years” (1999: 95). Most, however, such as Anthony Brueckner (1985) and Jonathan Kvanvig (2006: Chapter 4), take a positive and more moderate line. Brueckner, although he sides with closure, thinks that it ultimately depends on theoretically informed opinion. While believing that closure governs our use of the term *knowledge*, Kvanvig denies that it sounds paradoxical to say that one knows that one has hands but does not know that one is not a handless brain in a vat (or at least, he thinks it is not as bad as some theorists, e.g., DeRose (1995), make it out to be).

Contemporary open knowledge advocates (explicit deniers of the knowledge closure), are a slim minority. Notably, Fred Dretske (1970), Robert Nozick (1981), and Robert Audi (1991), who denies justification closure, have either advanced theories that entail open knowledge or have proposed (supposed) counterexamples to the principle.<sup>14</sup> Among the examples pro-

---

<sup>13</sup> There is a danger that a subject such as the present one is too central. The danger is that in order to adequately respond to a question as the present one, one would have to consider too large a portion of epistemology. Indeed I have had to neglect many issues that relate less directly to the line of thought I found myself following. For instance, I have had to only briefly refer to central alternative responses to skepticism and to leave out much of the discussion I planed to include regarding ampliative justification and knowledge, Bayesian epistemology (which I find congenial to open knowledge), epistemic compartmentalization, the weaknesses of other positions with regard to their attempts to leave knowledge closed, and more.

<sup>14</sup> Other theorists who are less known for their advocacy of open knowledge are Colin McGinn (1984); Gilbert Harman and Bret Sherman (2004); Stephan Yablo (MS); Mark Heller

posed as counterexamples to the principle of closure, Dretske (1970) has in/famously claimed that if, on the basis of looking at a zebra, one comes to know that *there is a zebra in the zoo pen*, and if one also knows that *if there is a zebra in the pen one is not looking at a mule disguised to look like a zebra*, it does not follow that one *knows* that *the animal in the pen is not a disguised mule*. The factivity of knowledge (the principle that knowledge entails truth), however, does of course entail the truth of the latter proposition.<sup>15</sup> Knowing through memory that *one's car is parked in the driveway* might not allow one to know by inference alone that *it has not since been stolen and taken away from the driveway* (Jonathan Vogel, 1990).<sup>16</sup> Nozick (1981) claims that his open knowledge account has the advantage of having a way to answer Cartesian skepticism.

In general, the benefits of open knowledge originate in its ability to account for a conservation of knowledge in face of epistemically questionable consequences that are known to follow from it. Open knowledge, then, allows for a detachment of unknown consequences from known premises, which would be destroyed by *reductio* if closure were valid. But since closure is intuitive, it would be theoretically better to account for the supposed counterexamples in some manner that leaves closure intact, provided one could also account for why, despite appearances, one would be wrong to think that the consequences are not known while the premises are.<sup>17</sup> If we could not account for why we know consequences of propositions we know despite appearances to the contrary, closure would leave us exposed to radical skepticism.<sup>18</sup>

---

(1999); Krista Lowler (2005); Bryan Frances (1999); Alvin Goldman (1976); Stephan Mainzen (1998), and James Cargile (1995). Some theorists who think knowledge is open have not written on the subject, and there are probably others of whom I am either unaware or fail to remember.

Professor Martin Kusch has suggested to me that Ludwig Wittgenstein is committed to open knowledge at least as far as "hinge propositions" are concerned, and Mark Kaplan (forthcoming) claims that J. L. Austin also held an open knowledge view. Thanks to Martin Gustafsson for bringing this to my attention.

<sup>15</sup> Open knowledge advocates do not typically contest the validity of *modus ponens*. Vann McGee (1985) who does contest it would be a closure denier only on some variants of the principle. He need not deny closure on the principle I will claim is the one at the center of a proper debate.

<sup>16</sup> Vogel does not present this as a counterexample but rather as an improved version of the kind of example Dretske has proposed that is only a *prima facie* case against closure (he claims). Vogel himself is one of the prominent defenders of closed knowledge.

<sup>17</sup> Some closure advocates admit that intuitively and unreflectively, at least, premises of arguments that conclude with the negation of skeptical propositions seem to be known while the consequences are not. In fact otherwise it would be hard to account for why skepticism is challenging.

<sup>18</sup> A central response to skeptical challenges, e.g. contextualism, has been to add some index relative to which the truth of knowledge ascriptions of the premises come and go together with the truth ascriptions of the knowledge of the conclusions. It is then a further question whether relative to this index knowledge is closed or open.

Those who agree that we know many propositions to be true—non-skeptics—are left with the choice of trying to account for why, despite appearances to the contrary, one does know propositions of the sort that have caused controversy, or account for the loss of knowledge of the premises despite the initial judgment that they had been known, or alternatively, account for why closure, despite our intuition, is invalid.

Unfortunately, despite the centrality of this issue, few have proposed direct arguments for closed knowledge that go beyond intuitive considerations, although two arguments proposed by John Hawthorne against open knowledge will be considered in a later section of this chapter. Most arguments proceed on the *assumption* that knowledge is closed. In fact, the understanding that a theory entails open knowledge is commonly taken to be a decisive argument against it, regardless of its theoretical merits.<sup>19</sup> Yet, as we saw in the last section, good reasons justify inspecting the question directly with the best tools at our disposal rather than simply settling for uninformed intuitions, unless we find we have no alternative.<sup>20</sup> That investigation is the purpose of most of the discussion in this manuscript.

I will not attempt to fill all the theoretical lacunae in the direct arguments for closure or openness, but I will take steps in that direction. Posing some challenges to closed knowledge, challenges that give reason to believe that knowledge is open despite appearances to the contrary, is the course I will follow. Specifically, the claim I will be defending is that solid reasons back the thesis that knowledge is not closed.<sup>21</sup>

---

<sup>19</sup> Such is the way many theorists respond to Nozick's tracking account of knowledge.

<sup>20</sup> My view regarding the role of intuitions in epistemology is that they are indeed central. However, when appealing to intuitions, we have to be relatively confident that there is no further analysis that can clear the way to view things in light of more basic cases. As will become clear in this chapter as well as Chapter 4, there are appealing principles that must be rejected in light of considerations that show them to be uninformed and indeed false.

<sup>21</sup> Maitzen (1998) claims that knowledge has been shown by Kaplan and Montague (1960) to be open. The knower paradox, he claims, establishes that knowledge cannot be closed. Kaplan and Montague use a Gödel type diagonal sentence construction of the form  $K(\neg p) \leftrightarrow p$  to derive a contradiction assuming closure, factivity, and knowledge of factivity. The details need not concern us here. If Maitzen were right the present discussion would be of substantially diminished interest.

Space precludes me from considering Maitzen's argument. A clue to the fact his conclusion is questionable is the ongoing debate regarding which of the premises is to be given up. Specifically, see the debate between Cross (2001) and Uzquiano (2004). Professor Pagin has shown me that Maitzen's way of formulating his argument commits him to the unwelcome result that one can prove  $p$  and will be in no position to know that  $p$  is true. I do not believe that closure or non-closure can be proven.

## 1.2 Layout of the Argument

The main argument of this manuscript takes the form of a challenge that includes several sub-arguments. These sub-arguments might seem at first isolated, but will seem less so in light of the argument in Chapters 4 and 5. By the end of Chapter 5, I believe that the formidable benefits of views that do not assume closed knowledge will have become evident. The overall argument, then, has two parts: the first poses challenges to closed knowledge, and the second is the appreciation of the theoretical benefits accruing from open knowledge.

As for the sub-arguments, after locating and explaining which closure principle should stand at the center of a proper debate, I use two of John Hawthorne's arguments against open knowledge to show that Fred Dretske's and Robert Nozick's open knowledge accounts are problematic. Nozick's account, I claim, is inconsistent. The remainder of Chapter 1, then, is an attempt to state where the debate stands at present. In Chapter 2, I advance three types of skeptical challenge to closed knowledge: Cartesian Skepticism, Mundane Skepticism, and a skeptical challenge that Bryan Frances (2004) has proposed, for which he coined the term Live Skepticism. In articulating these challenges, I try to mark some of their advantages and shortcomings from a skeptical perspective, as well as note their relation to the principle of closure. However, I do not survey the responses that have been offered to skepticism, nor do I try to say anything about the potential ones. Instead, I use them as a guide for what I see to be the underpinning of the challenge to closed knowledge. This underpinning is further articulated through another set of problems for closed knowledge: Saul Kripke's Dogmatism puzzle is the first, and the second is Epistemic Ascent, which covers the problem of Bootstrapping proposed by Jonathan Vogel (2000) and the Easy Knowledge problem introduced by Stewart Cohen (2002, 2005), as well as some other related challenges. I criticize the standard solution to the Dogmatism puzzle proposed by Gilbert Harman (1973), Roy Sorensen (1988), and John Hawthorne (2004a). My claim is that, contrary to the accepted wisdom, the Dogmatism puzzle remains unresolved, and is an instance of the general conundrum of closure generating unwarranted Epistemic Ascent. This is the main line of my argument in Chapter 3.

The underpinning of all these challenges, I argue next, is the openness of evidence, that is, that if one has evidence  $e$  for a proposition  $p$ , and one knows that  $q$  follows from  $p$ , one does not necessarily have evidence for  $q$ . To establish the openness of evidence, I set out two arguments on probabilistic grounds as well as one on non-probabilistic grounds, that is, on basic principles of evidence. I reflect on the historical origins of open evidence, mainly through the work of Carnap and Hempel. Evidence openness is then advanced as the main argument against knowledge closure, leading to a better understanding of the challenges to knowledge closure set out previ-

ously in the manuscript. In Chapter 5, I respond to rejoinders to the argument for open knowledge from open evidence. Some of these rejoinders help to refine the argument, including a response to Hawthorne's arguments that, in Chapter 1, had been found to be decisive against the most notable open knowledge views (Dretske's and Nozick's). Chapter 5 also critiques Timothy Williamson's knowledge account, which rejects one of the premises of the argument for open knowledge. The chapter also offers a preliminary defense of a certain form of justification closure that is compatible with open knowledge. I then claim that we can explain the oddness of open knowledge and solve all the previous challenges if we understand open knowledge on the basis of open evidence. This is the second part of the overall argument that states the benefits of open knowledge. I claim that, having solved these problems and understanding why knowledge could be open and why it might seem to be closed while respecting much of the intuition regarding closure through justification closure, we should take open knowledge as a serious theoretical contender.

Obviously, some important issues are not considered in this manuscript. I provide no theory of knowledge, evidence, or justification, and I am not concerned here with how open knowledge can be combined with existing theories of knowledge, although I think open knowledge can be combined with many of them. Williamson's theory is an exception. I discuss his view insofar as it pertains to the issues at hand, mainly because his view is almost diametrically apposed to the view of open knowledge that emerges here. As noted, I refrain from criticizing other replies to the different challenges I posed, though I doubt any one theory could deal with all of them. Most notably, I do not criticize epistemological Contextualism or Subject Sensitive Invariantism, at least not in a systematic and substantial way. I leave those issues for another occasion.

My thesis, as I said, is that open knowledge should be taken seriously. My own view is that knowledge is open, though I am not completely convinced of this. I am also inclined to view the entire situation as paradoxical. I see the intuitive force of closure, but also the import of the arguments that lead me to believe it to be invalid, as well as the theoretical benefits accruing from open knowledge. When and if it turns out that there is a need to revise our thinking about the arguments for open knowledge and also renounce its benefits, this itself will help to advance the understanding of these central issues.

In order to proceed, we need a better understanding of what the closure principle is, what the surrounding debate and its proper terms are, and what might be its driving motivation.

### 1.3 Closure and its Motivation

To gain a better understanding of the closure principle, I want to introduce a twist into the story about closure as articulated in my opening remarks, namely, that inferences of the basic *modus ponens* variety seem to preserve knowledge and give us reason to believe knowledge is closed. This formulation is misleading. Even if this is how we tend to reason about closure, there is good reason to think that all of us, friends and foes of closure alike, should think that knowledge is open with respect to these basic inference patterns. If these patterns are open, they surely cannot support knowledge closure.

When a subject knows that  $p$ , she might not be absolutely certain that  $p$ . This can be represented by some measure of probability in the unit interval, such that the probability of  $p$  is high but less than 1. As we will soon see, the basic idea does not have to be stated probabilistically. What matters is that, whatever measure of (epistemic) uncertainty we allow a subject to have while knowing that a proposition is true, this uncertainty can accumulate. Thus, one might know that  $p$  with some degree of uncertainty, and one might also know that *if  $p$  then  $q$*  with the same degree of uncertainty. The rational degree of uncertainty of the *modus ponens* conclusion,  $q$ , will be greater than the uncertainty one is required to have for either one of these two premises. Taking  $\text{Pr}(\bullet)$  to be a rational credence function for a subject  $S$ , we might have the following case:  $\text{Pr}(p)=0.9$ ,  $\text{Pr}(p \rightarrow q)=0.9$ , and as the case may be,  $S$ 's rational credence for  $q$  should be 0.81.<sup>22,23</sup> Whether or not we want to place a threshold for knowledge, surely we do not want to commit ourselves to the possibility that one will know a proposition when the rational degree of confidence in that known proposition is lower than 0.5. Moreover, it seems highly questionable that it is rational for a subject to be very confident that  $\neg q$  is true while she knows that  $q$ .<sup>24</sup> But if we take *modus ponens* to be closed with respect to knowledge, we will have to say something like this.<sup>25</sup>

Like most subjects, I know many things. I know that I was born December 31<sup>st</sup> 1967, I know that my mother is married to my father, that I will eat lunch today, etcetera. If I list these things and then think of whether I know

---

<sup>22</sup> Here I am appealing to the possibility of independence. Also the case can be dramatized a bit more if we take each premise to be 0.7 and with independence we would have  $\text{Pr}(q)=0.49$ . But as we will see nothing of substance will turn on the values we chose as long as they fall short of 1.

<sup>23</sup> Christensen advances similar arguments with respect to belief. See Christensen (2004: Chapters 1 and 2). For him the basic question with regard to rational belief closure is whether beliefs are binary or modeled on probabilities. The basic question, I think, is different. It is whether rational doubt is allowed to accumulate, or in the case of knowledge, the issue turns on whether known beliefs are allowed to be rationally/epistemically less than certain. Also, as will become clear, the same argument can be restated assuming that fallibilism is true. The discussion here is closer to Hawthorne (2004a:181-5).

<sup>24</sup> It is also epistemically questionable that one would know that  $q$  while being rationally required to be highly confident that  $\neg q$ .

<sup>25</sup> We might just have used *modus tollens* or some of the other basic inference patterns.

that the conjunction of the items on my list is true, I may realize that it is highly likely that at least one of the propositions on my list is false. Nothing seems strange or irrational about my doubt regarding this conjunction. In fact, I think something would probably be wrong with me if I had the same level of confidence in the conjunction as I have for each of the many conjuncts on my list. This point has been reiterated as a lesson regarding rational belief at least since Kyburg presented his Lottery paradox (1961),<sup>26</sup> and Makinson (1965) presented the Preface paradox. If I believe that each of  $p_1, \dots, p_n$  on my list is true, and assuming the number of items on the list,  $n$ , is sufficiently large, I ought not believe the conjunction  $p_1 \wedge \dots \wedge p_n$ . This Preface paradox type case suffices to show that the transition from a conjunction of many items of belief to a belief in the conjunction of those items is not innocent. Innocence is somewhat of an understatement, since I may even realize that to believe that this conjunction is true is not rational for me because, say by induction, I know that I have very good reason to doubt its truth. Since the evidence for this conclusion is often at my disposal, my ability to know the conjunction appears dubious whether or not I have such a realization.

One might suppose that the Preface non-closure and the *modus ponens* closure are epistemically compossible.<sup>27</sup> That is, since the rational credence of a known proposition is high, and since the *modus ponens* case is one where I merely know  $p$  and know that *if  $p$  then  $q$* , doubt will not accumulate, at least not to the point where the premises are known and the conclusion is not. Alas, this initially attractive view fails. If one accepts that, in the preface case, I can know that each of the items on my list  $p_1, \dots, p_n$  is true, while failing to know (since even for me it is not rational to believe) that  $p_1 \wedge \dots \wedge p_n$  is true, *modus ponens* is not closed.

In a review of Roy Sorensen's *Blindspots*, Peter Pagin (1990) has shown why. We are assuming that I know that each of  $p_1, \dots, p_n$ . Now suppose I reason as follows:<sup>28</sup>

---

<sup>26</sup> Kyburg develops this argument further, among other places, in his (1970) paper, "Conjunctivitis." Hawthorne (2004a: 46-50) connects these considerations with closure but comes to different conclusions than those that are reached here.

<sup>27</sup> Roy Sorensen thought so, at least implicitly. See Peter Pagin (1990). This mistake is quite frequent, unfortunately.

<sup>28</sup> In this manuscript the double arrow " $\Rightarrow$ " denotes a priori strict implication. I will use " $\vdash$ " as well as a standard underline to denote logical entailment (in the meta language) and " $\rightarrow$ " for material conditionals. Other symbols are either standard or will be introduced in due course.

In the present context I focus on strict implication since although it is not a matter of logical entailment that *I am thinking*  $\Rightarrow$  *I exist*, if this implication is strict and true, it will have the same implications as a logically valid one. That is, it will be knowable *a priori* and will not add any new doubt or uncertainty that was not in the other premises of the argument. The demand that it be *a priori* knowable is to set this implication apart from propositions such as *if this is water, it is H<sub>2</sub>O*. Those will be strict implications which I can be less than certain of and that can add rational doubt when conjoined with other propositions. As the examples in the main text will make clear, a double arrow will not entail that the implication is known *a priori* but merely that it can be known by reasoning alone.



$$(1) \quad \frac{p_1 \quad p_1 \Rightarrow (p_2 \rightarrow (p_1 \wedge p_2))}{p_2 \rightarrow (p_1 \wedge p_2)}$$

by *modus ponens* closure, I know that:

$$(2) \quad p_2 \rightarrow (p_1 \wedge p_2)$$

and since I know that  $p_2$  is true (it's on my list), I now know both premises of a *modus ponens* argument. So by a second application of *modus ponens* closure, on (2) and  $p_2$ , I know the conclusion of (3).

$$(3) \quad \frac{p_2 \quad p_2 \rightarrow (p_1 \wedge p_2)}{p_1 \wedge p_2}$$

Clearly, by going on in this manner I will know  $p_1 \wedge \dots \wedge p_n$ , which we assumed I do not know. Our problem is that accumulation of doubt, the requirement that I cannot rationally be required to believe a proposition and know its negation, together with the idea that if I know that  $p$  implies  $q$  I will know that  $q$  if I know that  $p$ , form an inconsistent triad. Notice that no probabilities were appealed to in the course of Pagin's argument. The crucial step was already made when we assumed at the outset that rational uncertainty, doubt, or risk, can accumulate.<sup>29</sup>

But if we do not restrict ourselves to non-probabilistic reasoning, we can verify that, although per hypothesis each of the propositions  $p_1$  through  $p_n$  has a very high probability, the proposition  $p_1 \wedge \dots \wedge p_n$  can have a very low probability. To locate the problematic step that I make in my reasoning, we may also verify that the step from (1) to (2) does not introduce any new doubt that was not in  $p_1$  to begin with. It is my bringing in  $p_2$  in (3) and placing it as premise, together with my knowledge that  $p_2 \rightarrow (p_1 \wedge p_2)$  (knowledge that inherited its level of uncertainty from  $p_1$ ) that brings in a new level of doubt. Hence, just like in the simple case above of two items of knowledge, it was the *modus ponens* inference that brought in a new level of uncertainty. Following this line of reasoning, we are bound to conclude that it is not the

---

<sup>29</sup> A similar conclusion can be reached on a fallibilist assumption. I refrain from using fallibilism here since I think that the phenomena of rational uncertainty is wider. For the argument with regard to fallibilism see Chapter 5.3.3 pertaining to Williamson's view.

case that whenever I know that  $p$  and know that *if  $p$  then  $q$* , I know that  $q$ , and this is the case even if I know that I performed a valid inference. Somewhere, I went from known premises of a *modus ponens* inference to an unknown conclusion.

Though other approaches to this problem might be possible, we have good reason to suspect that *modus ponens* closure is too strong a principle. Some theorists claim that we need to weaken the closure principle since I might know that  $p$  and know that  $p$  implies  $q$ , and yet might fail in *putting my items of knowledge together*.<sup>30</sup> But as the accumulation of doubt argument would naturally lead us to believe, this is the wrong line along which closure should be weakened. Even if I put the items together and reason impeccably, I will fail to know the conclusion of some *modus ponens* arguments. But if *modus ponens* is not closed with respect to knowledge, then *modus ponens* and other basic inference patterns do not give us motivation for closure (as I misleadingly suggested in the opening remarks).<sup>31</sup>

To reject closure of knowledge more generally, however, would be premature. Let us consider more abstractly how we stumbled into the inconsistency. On the one hand, the conjunction  $p \wedge q$  trivially follows from  $p$  and  $q$ . On the other hand,  $\text{Pr}(p) > \text{Pr}(p \wedge q) < \text{Pr}(q)$  (assuming independence and that  $0 < \text{Pr}(p), \text{Pr}(q) < 1$ ). If my rational belief that  $p$  and  $q$  is less than certain, then basic inferences I know to be valid can carry me from rationally held premises to conclusions I should not believe. Two ways immediately suggest themselves as solutions to our problem. The first is to jettison the idea that rational doubt can accumulate. But regrettably, no non *ad hoc* reason is forthcoming, possibly leading us to think that, once doubt is present, it will not accumulate. A second way is to claim that no epistemic doubt is present

---

<sup>30</sup> Lewis (1996: 442-3) cashes out the phrase “putting items of knowledge together” in terms of epistemic compartments. We can be compartmentalized in such a way that we do not know that we have legs (since we are considering skeptical possibilities) but we know we are bushwalking (and using our legs in so doing). He then suggests that a subject knows what is known in any of her compartments. I find that this suggestion is problematic since it places severe constraints on what we can know in different compartments about our ignorance, or alternatively, that our theoretical reasoning about knowledge is deemed uninteresting. Here is why. Assuming we know that we are ignorant of having legs in our philosophical compartment and we know that we are bushwalking and in doing so using our legs in the bushwalking compartment, assuming Lewis' suggestion that we know what is known in any compartment, we would count simultaneously as knowing and not knowing we have legs (by factivity). So either we cannot know we do not know we have legs (in the philosophical compartment), or our knowledge pertains solely to what we know in a given compartment. The first option is the unreasonable restriction on knowledge of ignorance and the second is the unreasonable result that our knowledge of what we know or don't know pertains to what is known within a compartment. Surely more needs to be said here (and I have said more elsewhere - Spectre (unpublished manuscript a)), but to pursue it would take me too far away from the current discussion.

<sup>31</sup> This claim I think is well supported even if we do somehow manage to find a way to leave *modus ponens* closed. For the claim to be well supported I need not convince anyone that it is not closed. Suffice it that we have good reason to think so.

to begin with. Though his reasons are different, Timothy Williamson's account amounts to something resembling this suggestion.<sup>32</sup> I will argue that this view does not solve our problem in Chapter 5.

An entirely different reaction to the inconsistency, which I believe to be preferable, is to leave *modus ponens* and the other valid inference patterns behind, that is, to admit that they are not closed. What we need are inferences that do not allow doubt or uncertainty to accumulate. Besides the problem of finding the right way to weaken the principle of closure, the drawback of this strategy is that a central everyday intuitive motivation for the principle of closure is lost. As things stand, a subject may know that  $p$ , know that  $p$  implies  $q$ , put these items of knowledge together, reason impeccably, and fail to know that  $q$ .<sup>33</sup>

With that said, it is time to locate a closure principle that can stand at the center of a proper debate.

## 1.4 What is the Closure Principle?

How should we weaken the closure principle? What got us into trouble in the previous section was the accumulation of rational uncertainty. What we want, then, is to ensure in some way that, in the course of an inference that is closed, doubt or uncertainty will not accumulate. In contrast to focusing on how to conjoin or "put together" items of knowledge, we need to be more vigilant about what kinds of items of knowledge can be unproblematically conjoined. Thus, in place of *modus ponens* closure, we may offer the following principle:

- (CP1) For all  $p$  and  $q$ , and subjects  $S$ , if  $S$  knows that  $p$  and knows that necessarily if  $p$  then  $q$ , and  $S$  puts these items of knowledge together,  $S$  knows that  $q$

where "knows" relates to a single time throughout. The reason we might think that (CP1) will not fail where *modus ponens* did is that (CP1) guarantees that, besides the doubt in  $p$ , no doubt will be added in the consequence  $q$

---

<sup>32</sup> I think that it is unfortunate that he does not highlight this aspect of his knowledge account. A significant benefit of his view is that it supports the closure of basic valid inference patterns such as *modus ponens*.

<sup>33</sup> The reason for this has been pointed out in different ways by Kyburg (1961, 1970) and Makinson (1965) and has had a great influence on theories of rational belief, but less so, I suspect, with regard to theorizing about knowledge. Although appreciating this neglect should, I think, influence our thinking about the motivation of knowledge closure, fortunately the main issues regarding its consequences vis-à-vis closure have been taken into account by some epistemologists. The reasons that have been given for restricting closure, restrictions that avoid these consequences, however, are different from those given here.

since one knows the consequent follows necessarily from the antecedent. Nevertheless, we are not in the clear since (CP1) neglects the doubt that might accumulate due to the fact that one might not be completely and rationally certain that, *necessarily, if p then q* is true. In other words, (CP1) will not solve our problem.

An implication can be epistemically less than certain in two ways. First, the implication that figures in the *modus ponens* inference might not be a strict one and might, for instance, be a material implication. Second, even if the implication is strict, it might not be known with epistemic certainty, for instance, I might know that a strict implication is true on the basis of a logician's testimony. (CP1) addresses the first kind of uncertainty, but not the second. And if in any of our *modus ponens* inferences we have an implication of one of these two sorts, doubt will have room to accumulate. To see this more clearly consider the following case. Suppose I know some proposition  $p_1$  is true, and that I learn from a logician we may call *number 1* that  $p_1$  logically entails  $p_1 \wedge p_2$ . Also, from logician *number 2* I learn that  $p_1 \wedge p_2$  logically entails  $p_1 \wedge p_2 \wedge p_3$ , ... and from *logician 99* that  $p_1 \wedge \dots \wedge p_{n-1}$  logically entails  $p_1 \wedge \dots \wedge p_n$ . Assuming that these testimonies are independent, I know by the same reasoning as before that the likelihood of at least one of the 99 logicians' testimonies being erroneous is very high. (CP1) seems suspect since the same kind of reasoning that led us to suspect *modus ponens* closure can be emulated using the logician case.

Suppose we impose a further restriction to overcome this difficulty:

- (CP2) For all  $p$  and  $q$ , and subjects  $S$ , if  $S$  knows that  $p$  and knows *a priori* that *necessarily if p then q*, and  $S$  puts these items of knowledge together,  $S$  knows that  $q$ .

(CP2) seems to solve the problem that the logician case raises, but still leaves much to be desired. First, it highly restricts the application of closure. Only when we know *a priori* that  $p$  strictly entails  $q$  will (CP2)'s antecedent be satisfied. Hence, this leaves us with a principle that says nothing about many cases we might think are closed. Second, I may know that all 99 logicians are better *a priori* reasoners than I am. To claim that I would know that  $p_1 \wedge \dots \wedge p_n$  is true due to my own limited reasoning, but that I would fail to know it is true on reasoning I know to be superior to my own, appears suspect.<sup>34</sup>

The issues I will discuss do not turn on the kind of doubt, uncertainty, or risk that stems from possible failures in the reasoning process itself, though

---

<sup>34</sup> Maria Lasonen-Aarnio raises similar concerns in her (2008) paper. She uses successive deductions to the conclusion that risk will accumulate using single-premise closure similarly to the way it mounts on its multi-premise version. Hence, she claims, single- and multi-premise closure rise or fall together.

a good argument against closure along these lines is possible.<sup>35</sup> Let us then bracket this last problem and focus on the first. Regarding the first problem, one useful course might be to look more closely at any motivation for closure we can find to see whether it fits the conclusion we have reached, and consider whether we can improve and close in on the principle that should stand at the centre of a proper closure debate. In other words, the closure principle we have reached might not capture the idea we set out to express.

In set theory, a set  $D$  is said to be closed under a given operation  $C$ , provided that for every object  $x$ , if  $x$  is a member of  $D$ , and  $x$  is  $C$ -related to any object  $y$ , then  $y$  is a member of  $D$ . What is required, then, is an operation on knowledge that has a chance of leaving it closed without uncertainty accumulation. In light of the previous section, a good candidate for closure can be formulated by taking  $D$  to be a set of propositions known by a subject, and stating that  $D$  will be closed, under knowledge that an operation of *a priori* necessary inference  $C$  has taken place on a given member of  $D$ :

If  $C^* \in D$  and  $p \in D$ , then  $q \in D$

where  $D = \{p: p \text{ is known by } S \text{ at } t\}$ ,  $C^*$  is the proposition that  $S$  has carried out operation  $C$  ( $=$  *inferring a priori*  $q$  from  $p$ ) and  $p$  and  $q$  are proposition variables. But notice that this kind of closure relates to members of the set  $D$  and not to subsets of  $D$ . Using “ $K$ ” to denote the knowledge operator, i.e.  $K(q)$  ( $=q$  is known by  $S$  at  $t$ ), we can set aside the sets in favor of a more conspicuous and familiar formulation:

(CP3)  $K(p \wedge (p \Rightarrow q)) \Rightarrow K(q)$

This principle is a candidate for a correct closure principle. It evades the above problems associated with the Preface paradox essentially by assuming that knowledge of the premise survives conjunction of an item of knowledge and an item of knowledge of a necessary implication.

With this point in mind, I wish to refer back to something mentioned earlier. I said that if we put two items of knowledge together, say the two premises of a *modus ponens* inference, we have no guarantee (even according to

---

<sup>35</sup> The argument should, it seems, take the form of a dilemma. One horn equates conjunction introduction closure ( $K(p) \wedge K(q) \Rightarrow K(p \wedge q)$ ) with any other restricted closure principle on the grounds that there is an accumulation of doubt involved in both. This horn will lead to the Preface paradox for knowledge. The other horn upholds the distinction between the two kinds of closure principle by demanding that one knows that the inference is valid. But as Dag Prawitz argues, this demand would lead to an infinite regress since it amounts to the claim that in order to know the conclusion by inference one needs to know that it follows from the premises. Presumably one would have to make some deductions in demonstrating that the inference is right. But that would only allow for knowledge that the inference is valid if one knows that further inference is valid, etc. For an elaboration on this argument see Prawitz (2009). I take this opportunity to thank Professor Prawitz for bringing his argument to my attention and for several stimulating discussions.

many who accept closure) that one will know the conclusion. Now, however, we have apparently found a way to understand the phrase “putting together” so as to distinguish cases that are closed from cases that are not. If I can put my items of knowledge together and know them in conjunction, closure says that the conclusion will be known as well. If in putting the items of knowledge together, knowledge does not survive, bets are off. In fact, we might as well strengthen the principle so that it can include cases where the implication is not strict and yet the items of knowledge continue (as a conjunction) to be known. In other words, what matters is that the inference is made from a proposition that falls in its entirety within the scope of a single knowledge operator. Thus, a better candidate for a closure principle is:

$$(CP4) \quad K(p \wedge (p \rightarrow q)) \Rightarrow K(q)$$

(CP4) might not appear to establish clearly that the *necessary a priori* implication is known *a priori* or that the inference from the premise to the conclusion needs to be based on *a priori* and correct reasoning, but only the latter problem should be a matter of concern. The premise of (CP4) assumes that the implication and the (empirical) proposition are known, i.e., they fall within the scope of the same knowledge operator. We also have an explanation now of why it seems that closure is supported by *modus ponens*. In thinking that *modus ponens* is closed, we might have missed an important distinction between a subject S who knows that *p* and knows that *if p then q* ( $K(p) \wedge K(p \rightarrow q)$ ), and a subject S' who knows that *p and that if p then q* ( $K[(p) \wedge (p \rightarrow q)]$ ). Only subject S' will know by closure that *q*. S has no such guarantee.

But now we may realize that the implication does not do any real work here. The work is all done by the scope of the knowledge operator. So in order to have a clear idea about what closure principle is at issue and assuming (as I will throughout) that there is no problem with the *a priori* reasoning itself, we can state the principle at issue:<sup>36</sup>

$$(CP) \quad \text{For all propositions } p \text{ and } q \text{ and subjects } S, \text{ if } S \text{ knows that } p \text{ and infers } q \text{ } a \text{ priori from } p \text{ (correctly), then } S \text{ knows that } q.$$

Several writers, e.g. Hawthorne (2004)<sup>37</sup> and Gettier (1963), have proposed closure principles for knowledge and justification that amount to the

---

<sup>36</sup> There are those who stipulate in formulating closure that a subject must believe that *q* on the basis of the inference. See, for instance, Hawthorne (2004a: 34).

<sup>37</sup> Hawthorne (2004a: 39) uses Lewis Carroll's argument regarding *modus ponens* to saddle the open knowledge advocate with the tortoise unreasonable position. But the single-premise closure leaves room for the same position. As I see it, his argument goes against single premise closure as much as it does with respect to closure deniers. Thus one might know that *p* and know that *if p then q*, but not know that *p and that if p then q* and thus will not know that *q*. If one would want to claim that one does know that *q* in all these cases one would be tacitly

same thing. What matters is that the reasoning is made from a single item of knowledge (or belief) to a new item of knowledge (or belief). And what “single” means here is that it is a proposition that is governed by the knowledge (or belief) operator. While repudiating multi-premise closure, some theorists formulate closure with two premises (necessarily, if  $(Kp \wedge K(p \rightarrow q))$ , then it follows that  $Kq$ ). Some even confusedly call it single-premise closure.<sup>38</sup> However, this formulation is in fact an instance of multi-premise closure as one might verify by Pagin’s argument above. It would be true if the premises can always be known as a conjunction, i.e. if  $Kp \wedge K(p \rightarrow q)$  entails  $K(p \wedge (p \rightarrow q))$ . But this is to beg the question as the validity of multi-premise closure is already assumed. The relevant individuation conditions of ‘premise’ are determined by the range of the relevant operator (in this case knowledge) since the degree of rational or epistemic uncertainty applies to the operator. Of course not everyone is guilty of this kind of conflation. Williamson (2000: 117) and Hawthorne (2004: 32-4), for instance, are careful to correctly distinguish single- from multi-premise closure. The crucial point I mean to stress is that the former does not entail closure of basic inferential modes such as *modus ponens*. We are thus left with a debate concerning (CP), which is a weaker principle than many have supposed.

A general concern with restricting closure to single items of knowledge is that the question of whether knowledge is closed or open becomes uninteresting. We have reached a weak principle that leaves out many of the everyday inferences that regularly inform our reasoning and action. In a sense, deductive operations do not result in general extension of knowledge (an idea that several theorists—e.g. Williamson (2000: 117)—have claimed is the main motivation for closure) if multi-premise closure is rejected. This is the case if expansion of knowledge is understood to depend on combinations of known propositions to form new, logically stronger conclusions than any one of the previously known propositions. I must admit that this is a serious concern, though one that should trouble both sides in the single-premise closure debate.<sup>39</sup>

To show how weak the principle is and how, on many matters, open and closed knowledge advocates will find themselves on the same side, let us

---

assuming multi-premise closure (that if one knows that  $p$  and that  $q$  one will know that  $p$ -and  $q$  by proper inference).

<sup>38</sup> For just one example see Fumerton (2006: 24-5).

<sup>39</sup> Together with the open knowledge advocate the closed knowledge advocate is bound to admit that one can know that  $p$  and know that  $q$  follows from  $p$ , one may consider these items of knowledge together and yet fail to know that  $q$ . This means that arguments that appeal to the intuition that one does know that  $q$  in the case of (CP) cannot be appealed to since it is just as intuitive that *modus ponens* closure is valid. In fact, many theorists state the closure principle as I have stated the *modus ponens* closure principle. Notice there is no forthcoming reason why a subject should lose her knowledge of premises as a result of failing to know the conclusion of a proper *modus ponens* inference. Thus the move of trying to retain closure through denial of knowledge of the premises, does not seem promising here.

return to the troubling case of my wife going to the Bahamas with her deceased husband. The weakened closure principle will allow my wife to know both that *if she gets extended time-off we are going to the Bahamas*, and that *if I die, she will get an extended vacation*. For this case to pose a challenge to (CP), she would have to know both conditionals in conjunction. This union could apparently easily fail and, therefore, she would not know (thankfully) that she will go to the Bahamas with her deceased spouse should I come to my demise. Perhaps one would like to claim that she does in fact know this since she knows that I will not die in the near future. I am sympathetic to this reply, but my sympathy is not of a theoretical nature (it is because it entails my survival). The point is that we should be doubtful about finding a reason why she would not know the pair of conditionals in isolation, and hopeful of finding a reason why she would not know them in conjunction, whether or not she needs to know that I will not die in order to know the first conditional.<sup>40</sup> A stronger closure principle than the one I have advocated would obviously be in direct conflict with this judgment about my wife's knowledge. But this should be considered a problem for the stronger closure principle, not for the reply that both open knowledge advocates *and* (CP) advocates have the resources to offer.<sup>41</sup>

In appreciating how weak (CP)'s weakness (though not all agree that there is no stronger valid closure principle), we should not think that the question of its validity is insignificant. Admittedly, some issues (such as the one regarding my wife's epistemic state) do not depend on the question of whether (CP) is valid, but the issues I listed earlier are highly significant and most of them do depend on its validity.<sup>42</sup>

---

<sup>40</sup> An open knowledge advocate could also accept, though she need not, (something that a CP advocate would have to deny) that the conditionals can be known in conjunction with an inability to come to know the objectionable conclusion on the basis of a proper inference. Whether this is desirable is another question.

<sup>41</sup> Although knowing the conditionals in isolation does not mean that (CP) advocate must know them in conjunction, there is a different method by which my wife will know that *if I die, we are going to the Bahamas*. Assume that she knows that *if she gets extended time-off, we will go to the Bahamas*. She can infer from this knowledge that *if (if I die, then she will get extended time off), then, (if she gets extended time off, we will go to the Bahamas)* and by (CP) she will know this. But only by knowing the antecedent of this conditional in isolation and putting it together with her knowledge just gained by (CP), will she know the consequent by a second application of (CP) on the conjunction of these two items of knowledge. Only then will she know that *if I die, we are going to the Bahamas*. This knowledge will indeed require knowing that I will not die and yet it is still an option for the closed knowledge advocate to deny this on the basis of claiming that the conjunction is something she will fail to know. I consider this case more extensively in Chapter 3.1.2.1, p. 77.

<sup>42</sup> I am undecided regarding whether Kripke's belief puzzle can be stated in terms of (CP) (for rational belief). The same goes for the semantic externalism problem.



## 1.5 Debating Closure

With a better understanding of the principle at issue, we may now turn to a characterization of the closure debate. Non-skeptics who accept CP—closure advocates—claim that, whenever a subject  $S$  knows that a proposition  $p$  is true,  $S$  infers *a priori*  $q$  from  $p$ , then even in cases where  $S$ 's evidence (or reasons) for believing that  $p$  are insufficient for knowing that  $q$  taken on their own,  $S$  can know that  $q$  without augmenting  $S$ 's evidence for  $q$  or revising  $S$ 's reasoning. The idea here is that knowledge is always expandable by competent inference from a single item of knowledge, provided that it can be maintained. Open knowledge advocates—deniers of (CP)—claim that this is not always the case. Turning things around, the open knowledge view can be seen as a commitment to the following claim:  $S$  may know that  $p$ , and *a priori* infer  $q$  from  $p$ , yet  $S$  may fail to know that  $q$ . Closure advocates deny that this is an epistemic possibility. If one fails to know  $q$ , closure advocates maintain, one cannot know that  $p$  once the *a priori* inference to  $q$  from  $p$  has been competently made.<sup>43</sup> This, then, is the crux of the debate.

To see the workings of this debate, I will first turn to Hawthorne's argument targeting the idea that the epistemic states that open knowledge advocates are committed to are not compossible, even according to their own lights. I will then show that, regarding Nozick and Dretske's view, Hawthorne is right.

### 1.5.1 Hawthorne's Arguments

Advocating knowledge openness, closure deniers stand in opposition to two kinds of closure endorsers – skeptics and optimists. Skeptics often argue that, since one does not know some proposition  $q$  that is known to follow from some other proposition  $p$ , one does not know  $p$ . Optimists claim that both  $p$  and  $q$  are known, either simpliciter (e.g. G.E. Moore), or with reference to different contexts of ascription (Contextualists), or to different practical environments the subject is in (Subject-Sensitive Invariantists).<sup>44</sup> What both skeptics and optimists agree upon is that, if  $q$  is properly derived from a known proposition,  $q$  is known. This is the contention that advocates of the openness of knowledge reject, as reflected in the denial of (CP).

To defend closure against examples advertised by its deniers, Hawthorne argues that, in the interpretation of their view that closure deniers would have us adopt, these examples conflict with other, more basic and weaker

---

<sup>43</sup> For further reference, where  $O(x,y)$  is the operation of an ideal reasoner inferring *a priori*  $y$  from  $x$  and  $t$  is a time earlier than  $t'$ , we may characterize the debate with the following principle:  $\forall p \forall q [((p \in D)_t \wedge O(p,q)_t) \rightarrow ((q \in D)_{t'} \vee (p \notin D)_{t'})]$ . Open knowledge advocates deny it and closure advocates endorse it.

<sup>44</sup> This is not to say that Contextualism or Subject Sensitive Invariantism entail closure of knowledge.

epistemic principles. The advocate of knowledge openness, he claims, is forced to reject these highly compelling principles, so that denying closure on the basis of these examples is tantamount to denying weaker principles as well. Thus, if his arguments are cogent (as I will claim they are), Hawthorne manages to show that anyone who advocates the weaker and more basic principles is not an open knowledge advocate, at least not a consistent one. Nevertheless, the same reasons that motivate the denial of closure tell against the weaker principles Hawthorne puts to task. Thus, anyone rejecting closure for the right reasons will also reject the weaker principles on which Hawthorne's argument relies. Or so I will argue below (Chapter 5.1). Let us look at Hawthorne's arguments.<sup>45</sup>

Exposing the deeper connections and further commitments of closure denial, Hawthorne's arguments helps articulate what I take to be the proper grounds for epistemic openness.

The following are Hawthorne's weaker principles (2004a: 41), which he would have the closure denier adhere to, at least initially:

*Equivalence (EQ)*: Necessarily, if S knows that  $p$ , and S knows that  $p$  is *a priori* equivalent (or logically equivalent) to  $q$ , then S knows that  $q$ .

*Addition (AD)*: Necessarily, if S knows that  $p$ , then by competently inferring  $p$  or  $q$  from  $p$ , S thereby knows  $p$  or  $q$ .<sup>46</sup>

*Distribution (DIS)*: Necessarily, if S knows that  $p$  and  $q$ , S knows  $p$  and S knows  $q$ .

Indeed, all three principles seem highly plausible. To see how they lead to the same conclusion as CP let us, following Hawthorne, get back to Fred Dretske's well known zebra case (1970: 1015-6). Seeing a zebra-looking animal in the pen labeled "Zebra," one knows that the animal in the pen is a zebra (call this proposition  $Z$ ). But, presumably, one does not know that this animal is not a mule disguised to look like a zebra ( $\neg DM$  for short). So, according to the closure denier (in this case, Dretske) one knows  $Z$ , and knows that  $\neg DM$  follows from  $Z$ , but does not know  $\neg DM$ . And indeed, as I have claimed above, this is the heart of the debate. The open knowledge advocate is committed to the claim that some cases such Dretske's zebra case are epistemically possible.

Now Hawthorne's first argument that such cases are epistemically impossible runs as follows. We are assuming that S knows that:

---

<sup>45</sup> Hawthorne (2004a) and (2005). I do not present or attempt to answer all of Hawthorne's arguments in support of closure, only the ones that I take to be most forceful and to pose the greatest challenge for epistemic openness of the kind I think is best motivated.

<sup>46</sup> Further clauses can be added to these principles, see Hawthorne (2004a: 39), but for simplicity we omit them here. Nothing in our argument turns on this simplification.

(4)  $Z$  [assumption]

(5)  $Z \Rightarrow \neg DM$  [assumption]

By AD, S can infer and come to know:

(6)  $Z \vee \neg DM$  [AD,(4)]

Assuming that S is familiar with basic logical operations, she can know that:<sup>47</sup>

(7)  $(Z \vee \neg DM) \Leftrightarrow \neg DM$  [PL,(5)]<sup>48</sup>

By EQ it now follows that S knows that

(8)  $\neg DM$  [EQ,(6),(7)]

Thus to avoid the implausible consequences of the example, closure deniers must also deny AD or EQ (or both). Since other counter-examples to closure share the form of this one, the same problem will arise for them as well.<sup>49</sup>

Hawthorne also employs a parallel argument using DIS. Assuming that S knows that (4) and (5) are true and that familiarity with basic logical operations enables S to know that:

(9)  $Z \Leftrightarrow (Z \wedge \neg DM)$  [PL,(5)]

By EQ, S knows:

(12)  $Z \wedge \neg DM$  [EQ,(4),(9)]

DIS entails that, knowing (11), S is in a position to know:

(13)  $\neg DM$  [DIS,(12)]

---

<sup>47</sup> Hawthorne (2004a: 41, note 99) notes that, strictly speaking, that a thing is a zebra does not logically imply that it is not a painted mule. I do not think it even follows *a priori*. This is a problem for the adequacy of the example and not so much for the issue at hand.

<sup>48</sup> "PL" will stand for basic operations of propositional logic.

<sup>49</sup> The examples considered here all contain only pairs of propositions since we are only dealing with cases concerning CP. Hawthorne thinks there are some prospects for maintaining "Multi Premise Closure" as well (2004a: 186).

Again, the conclusion that closure deniers aim to avoid is reached by principles weaker than closure. If these consequences mandate rejection of closure, they should also warrant rejection of these weaker principles. Hawthorne's argument successfully shows that, to avoid the undesirable consequence of closure, either EQ or both AD and DIS are to be discarded. He concludes that closure is pretty much "non-negotiable,"<sup>50</sup> in other words, the cases that open knowledge proponents are committed to are epistemic impossibilities.

Hawthorne offers what is perhaps not only the cleverest but also the strongest argument in defense of epistemic closure. Nevertheless, exposing the deeper connections and further commitments of closure denial, Hawthorne's argument helps articulate what I take to be the proper grounds for epistemic openness. Why this is will be the subject of Chapters 4 and 5 below. At this stage, I will consider Nozick's and Dretske's open knowledge accounts to see how well they can cope with Hawthorne's argument.

### 1.5.2 Nozick on Closure

According to Nozick's "tracking" analysis of knowledge, S knows that  $p$ , iff, the following conditions are met:<sup>51</sup>

- (i) S believes that  $p$  is true [By some method M]
- (ii)  $p$  is true
- (iii)  $not-p > not-(S \text{ believes that } p)$  [by method M]
- (iv)  $p > S \text{ believes that } p$  [by method M]<sup>52</sup>

Nozick's subjunctive conditionals are analyzed, for heuristic purposes at least, by considering the worlds closest to the one in which S's purported knowledge is assessed, or what he terms its *neighborhood*. So, for instance, (iii) is true if in all the closest worlds in which  $not-p$  is clearly the case, S does not believe that  $p$  is true. On such an analysis, if  $p$  entails  $q$  and  $q$  en-

<sup>50</sup> Hawthorne, (2004a: 112).

<sup>51</sup> I am using ">" as a symbol for subjunctive conditionals.

<sup>52</sup> The appeal to methods of beliefs is meant to circumvent cases such as the following. By seeing her grandson, a grandmother knows that he is healthy. Had her grandson been ill, however, her family would have prevented her from knowing this (perhaps they would claim that he is traveling abroad). Had the grandson not been healthy, then, she would still have believed he is. Intuitively, Nozick realizes, seeing that her grandson is healthy, the grandmother knows as much even though the unqualified subjunctive conditional (iii) is false. Only if she would believe it *using the same method* – seeing him – in the closest worlds where he is unhealthy, would she not count as knowing this in the actual world. Without appeal to methods condition (iv) would entail a skeptical commitment. Suppose I see a cow and believe it is a cow. Still there are neighboring possible worlds in which I do not form this belief, say because I am looking the other way, my eyes are closed or it is dark outside. Avishai Margalit using a different example made a similar point. Nozick discusses his example (1981: 180).

tails  $p$ , then  $S$  knows that  $p$  iff  $S$  knows that  $q$ . The reason this must hold is that if  $p$  and  $q$  logically (or conceptually) entail each other, then they are true in the same worlds (at least in the standard picture that Nozick appeals to). Realizing this, he explicitly endorses an equivalence principle (1981: 690, note 60). Now, since  $S$  knows that  $q$  is equivalent to  $p$ ,  $S$  will not believe  $q$  if she does not believe  $p$  (Nozick does not contest the closure of rational beliefs). So, in a world in which  $q$  is true,  $S$  would believe that it is true, and in worlds (close ones at least) in which it is false,  $S$  does not believe that it is false. Thus, if  $S$ 's belief that  $p$  is truth-tracking (as it must be if  $S$  knows that  $p$ ), then  $S$ 's belief in  $q$  is truth-tracking. So, if she knows  $p$  and knows that  $p$  and  $q$  are equivalent,  $S$  knows  $q$ . On Nozick's account, therefore, (EQ) is valid.

At the same time, (DIS), according to Nozick, is false. He claims, for example, to know that he works in Emerson Hall (at Harvard University) *and* is not a brain in a vat, but denies that he knows that he is not a brain in a vat (Nozick, 1981: 229). In all neighboring worlds in which Nozick does not work at Emerson Hall he does not believe that he works there (perhaps he is working someplace else or taking time off). Yet, in the neighboring worlds in which he is a brain in a vat, he still believes that he works in Emerson Hall. Hence, the belief in the conjunction is what Nozick calls "sensitive". Taken in isolation, however, belief in the second conjunct is not sensitive. Hence, one may know  $p$  *and*  $q$ , yet fail to know  $q$  on Nozick's account. It seems, therefore, that Nozick is not a target of Hawthorne's second argument, an argument that he in fact anticipates (Nozick, 1981: 228)

Nozick, however, is threatened by Hawthorne's first argument, since on his account both (EQ) and (AD) (roughly, that one can add disjuncts to propositions one knows without losing knowledge) are valid. The reason (AD) is valid according to the tracking theory is the following. If  $S$  knows  $p$ , then her belief in  $p$  must be sensitive (satisfies all conditions but, most importantly for our purposes, (iii) and (iv)). Now since when  $p$ -*or*- $q$  is false, both  $p$  and  $q$  are false, if  $p$  is believed sensitively then, in worlds in which  $p$  is false,  $S$  does not believe  $p$  and therefore does not believe  $p$ -*or*- $q$ . In worlds where  $p$  is true,  $S$  believes  $p$  and therefore also that  $p$ -*or*- $q$ . Thus, if  $p$  is believed sensitively, so is  $p$ -*or*- $q$ .<sup>53</sup> Hence, on Nozick's account (AD) is valid (Nozick, 1981: 237).

But if Nozick accepts (AD) and (EQ), then Hawthorne's first argument can be advanced against him. Insofar as he considers it unintuitive and false to regard the conclusions of the problematic closure inferences as known,

---

<sup>53</sup> This is a bit of a shortcut. One may argue that not in all worlds in which  $p$ -*or*- $q$  is false,  $S$  does not believe that this disjunction is false.  $S$  might, in those worlds, believe  $q$  independently of  $p$ . But the point to notice is that in such counterfactual situations,  $S$  does not believe the disjunction on the basis of inferring it from  $p$  (since in those worlds  $S$  does not believe that  $p$ ). Hence  $S$ 's belief in  $p$ -*or*- $q$  is arrived at by a different method. So AD is a truth-tracking method for Nozick. Regarding methods see previous footnote.

while he accepts (AD) and (EQ), Nozick seems to embrace inconsistent commitments. For unless Hawthorne's argument poses an unforeseen problem, the advocate of knowledge openness must either reject (EQ) or both (AD) and (DIS).<sup>54</sup> Nozick's failure to see the contradiction in his account is somewhat strange. On the one hand he claims:

Also, it is possible for me to know  $p$  yet not know the denial of a conjunction, one of whose conjuncts is *not-p*. I can know  $p$  yet not know (for I may not be tracking)  $\text{not}(\text{not-}p \ \& \ SK)$ . (Nozick, 1981: 228)

On the other hand, Nozick does accept both (AD) and (EQ), and a few pages later states:

Knowledge, almost always,<sup>55</sup> will be closed under existential generalization. Similar remarks apply to inferring a disjunction from a disjunct. (Nozick, 1981: 236)

But clearly, since from  $p$  one can deduce  $p \vee \neg SK$ , Nozick is committed to the claim that one would know this. However, just a few pages earlier he had denied that one could know a logically equivalent proposition while accepting (EQ). Even without the aid of Hawthorne's argument, then, Nozick's account is inconsistent.<sup>56</sup>

To make things clear, suppose Nozick knows  $p$  (= I am in Emerson Hall) and infers  $p$ -or-not-SK (= I am in Emerson Hall or I am not (in the tank on Alpha Centauri now and not in Emerson Hall)) from  $p$ . By the disjunction addition principle, he counts as knowing  $p$ -or-not-SK. But he also knows that  $p$ -or-not-SK is equivalent to  $\text{not}(\text{not-}p \ \& \ SK)$  and hence "the person who knows  $p$  will know"  $\text{not}(\text{not-}p \ \& \ SK)$ , which is precisely what he denies: "I know I am in Emerson Hall now, yet I do not know that: it is not

---

<sup>54</sup> One might suggest that to retain consistency Nozick could drop the possible world account of subjunctive conditionals that is inessential to his tracking account of knowledge. It would be surprising, however, if a different account of subjunctive conditionals would assign different truth values to the relevant conditionals. But this question will have to await more direct consideration. In any event as noted in the text, Nozick explicitly endorses transmission of knowledge over disjunction introduction. It remains to be seen if this is just a mistake or an essential part of his account of knowledge. My suspicion is that it is the latter that is the case.

<sup>55</sup> From the counter-example to the closure of knowledge under existential generalization presented in note 68 on page 693 it is clear that that by saying "almost always" Nozick does not question the validity of this principle of addition in any respect pertaining to the present discussion.

<sup>56</sup> The first to make this point with regard to Nozick's account was Saul Kripke. In correspondence Kripke pointed out to me that in his unpublished manuscript on Nozick's epistemology, he makes a similar point to the one made here. Yet he also says, "I never noticed the specific contradiction to what he says in page 228, nor the fact that I now see there is an argument for adding a disjunct based on Nozick's concept of method, and even explicitly given by Nozick... It looks to me as if you are right, as far as I can see now, so this isn't a slip on Nozick's part, as I had indeed thought." I am grateful to Professor Kripke for this exchange.

the case that (I am in the tank on Alpha Centauri now and not in Emerson Hall).” (1981: 228).<sup>57</sup>

### 1.5.3 Dretske on Closure and Heavyweight Propositions

Dretske’s rejection of epistemic closure resembles Nozick’s.<sup>58</sup> Like Nozick, Dretske considers the rejection of closure a consequence of the correct theory of knowledge. Since knowledge depends on having a belief based on reliable procedure, and since a procedure may be reliable with respect to a belief and unreliable with respect to something entailed by that belief, closure must be denied. Thus, perception of a zebra normally allows one to know that an animal is a zebra, since in the closest possible worlds in which there is no zebra, one does not perceive a zebra. Yet, it does not allow knowledge that the animal is not a disguised mule, because in the closest worlds in which it is (unbeknownst to one) a disguised mule, one presumably believes that it is not.

Dretske does not specify what his position is with regard to (AD), but he explicitly endorses (DIS).<sup>59</sup> The considerations stemming from his theory of knowledge on which his rejection of closure is based seem to commit Dretske to (at least) (EQ) and (AD), so he too is vulnerable to Hawthorne’s cost-raising arguments.

More recently, Dretske has advanced another line of defense for knowledge openness. Our mundane everyday items of knowledge, he notes, entail a host of “heavyweight propositions” that, presumably, we are not able to know. That I drank wine last night implies that physical objects exist, that the world was not created five minutes ago and a wealth of similar heavyweight propositions. If closure were valid we would have to know that these propositions are true – knowledge, Dretske claims, we do not seem to be capable of acquiring.<sup>60</sup>

---

<sup>57</sup> Since he knows that not-SK follows from *p*, he cannot accept the claim that he does know this sentence for he would then know (by the equivalence principle) that not-SK.

<sup>58</sup> Dretske precedes and to a great extent anticipates Nozick’s account of knowledge entailing epistemic openness. See Nozick in Sosa and Kim (2000: 100) for a summary of the relatively slight differences between his and Dretske’s account.

<sup>59</sup> At least he once did (see his 1970: 1009). In a recent exchange with Hawthorne, Dretske does not take this avenue replying to the latter’s arguments by rejecting its premises as I propose. Specifically, Dretske does not meet Hawthorne’s challenge to take a stand on (DIS) (see Hawthorne, 2005: note 10), which suggests either that he does not fully endorse the line I will propose for the open knowledge advocate, or that he fails to see that it entails the rejection of both (DIS) and (AD).

<sup>60</sup> Interestingly, in the 2005 exchange, Hawthorne’s arguments target Dretske’s theory of knowledge (particularly the counterfactual condition on reliable belief-formation) and the reasons for rejecting closure that stem from this theory (the only exception is the argument employing (EQ) and (DIS) (Hawthorne, 2005: 31), to which, interestingly, Dretske does not

Unfortunately, Dretske does not clarify exactly what counts as a heavy-weight proposition.<sup>61</sup> It seems that this category is defined for him by his notion of ‘conclusive reason,’ which, in turn, runs into a host of counter-examples that Hawthorne presents.<sup>62</sup> Moreover, the arguments and cases we have been looking at, which will be analyzed more thoroughly in the coming chapters, suggest that whether a proposition can be known through inference from known premises independent of its weight<sup>63</sup> does indeed not depend on any generalizable feature of the proposition in itself but on its relation to the evidence.<sup>64</sup> Hawthorne’s arguments do not hinge on the example he employs. Without appealing to specific features of heavyweight proposition, subsequent arguments will show that the logic of evidence provides reasons for rejecting (AD) and (DIS). Though I will not be considering this claim directly, the best open knowledge account will apparently not rely on counterfactual features of knowledge or evidence but on the nature of actual evidence that supports known propositions.

Although sensitive to the right intuitions regarding the central cases, both Nozick and Dretske seem to miss the centrality of actual evidence in their reflections on closure, although Nozick’s requirement regarding actual method does in some way move towards correctly reflecting complex epistemic situations. By focusing on the question of evidential support, the present account of open knowledge will provide not only an argument against epistemic closure and a defense of epistemic openness, but also an explanation of why closure fails, if it does.

Without the pretense of offering an exhaustive account of evidential support, the basic idea will be the following: To get things right about knowledge, at least about some knowledge (or even more modestly, about knowl-

---

reply by suggesting rejection of (DIS)). As it is not founded on sensitivity conditions of knowledge, the account of open knowledge proposed here is not susceptible to such arguments. Hawthorne’s only remark about the evidential considerations raised by Dretske (2005) (to which he is replying) is: “Better that we appreciate (with Dretske) the force of this problem than rush too quickly to try to solve it” (40). But, if the argument of chapter 4 and 5 are right, correctly understood, the problem of evidential transmission (or the lack of it) is anything but peripheral to the question of closure.

<sup>61</sup> Heavyweight propositions, he says, are “out of range: we *cannot* see (hear smell, or feel) that they are true” (Ibid, 20, my emphasis). Hawthorne takes it that a proposition P is heavyweight “just in case we all have some strong inclination to think that P is not the sort of thing that one can know by the exercise of reason alone and also that P is not the sort of thing that one can know by use of one’s perceptual faculties (even aided by reason)” (Ibid, 33).

<sup>62</sup> See Hawthorne (2005).

<sup>63</sup> In a recent paper Dretske defends his rejection of closure by invoking the idea that information is not closed under entailment. This forces him to make the unintuitive claim that a signal can raise the probability of some proposition to one and still the proposition might be not only unsupported by the signal, but might, moreover, be “not available to signaling” (2006: 412).

<sup>64</sup> Dretske invokes the idea of failures of transmission of warrant proposed by Wright and Davies (Dretske 2005: 15), curiously neglecting to specify if and how he intends to differ from their (at least in intent) closure-preserving account (Wright 2000: 157, for reasons to doubt the feasibility of preserving closure while acknowledging transmission failure see Silins 2005: 89-95).



edge that is not based on conclusive evidence), we must focus on the evidence and distinguish it from what it supports or, to put it in Dretske's terms, from the information it conveys. Perception of a zebra-looking animal conveys the information that the thing is a zebra. It raises the probability that the proposition "this is a zebra" is true. And yet, it does not raise the probability that the proposition "this is not a disguised mule" is true. As a matter of fact, it raises the probability that this latter proposition is *false*. Given that something looks like a zebra, there is a greater probability on one's evidence that it is a mule made to look like a zebra than if it looks to one like, e.g. a matchbox. That is why, if someone embarks on a mission to locate zebra-look-alike mules, she ought to search among the zebra-looking animals. For similar considerations, if you suspect your car might have been stolen, the place to look is where you parked it, since if it was stolen, chances are it was stolen *from there*. While memory of parking the car in the driveway supports the belief that the car is in the driveway, it raises the probability that this belief is true, it does not support the belief that the car has not been stolen from your driveway. Given that you parked it there, the probability that the car has been stolen from the driveway is *raised*. The claim, then, is that both Dretske and Nozick (as reflected in the principles they support) failed to rest their own open knowledge account on firm ground and gave a problematic explanation of the cases that they proposed as reasons to doubt closure. If a reason may be stated for questioning the validity of closure, it is the openness of evidence and not any necessary subjunctive condition on knowledge. Or so, at least, I will argue below.

#### 1.5.4 Section Summary

We have seen that the most prominent open knowledge accounts – Dretske's and Nozick's – are in a dire state. Both accounts of open knowledge lack the necessary resources to respond to Hawthorne's challenge. Open knowledge, then, does not seem to have much going for it.

Nevertheless, I think it fair to say that even if one strongly objects to open knowledge, it is of some importance to understand what open knowledge is if it is not as Dretske or Nozick propose and, in the last two chapters below, much is devoted to such an account. Among other things, I hope to show that there is an open knowledge view that can respond to Hawthorne's argument, and that understanding how this reply works makes clearer a fundamental challenge that closed knowledge faces (Chapters 4 and 5).

In order to examine this view of open knowledge and see the challenge it poses to closed knowledge, we will first need to consider more customary challenges to closed knowledge. The first is skepticism of various sorts.

## Chapter 2: Skepticisms and Closure

In this chapter I shall start with describing a new type of gloss that has recently been put on an old argument to generate different epistemological theories such as; Contextualism, Mooreianism, Subject Sensitive Invariantism, Relativism, Open knowledge, Skepticism and Idealism. Though the theories themselves will not be a topic of the current discussion, what is of importance here is to see the central role that the principle of closure plays in different skeptical challenges. This will justify my direct focus on the principle as well as the idea that by rejecting it there is a way to avoid several skeptical challenges (which is not to say that other ways of meeting them are impossible). Once the argument can be put in a more general framework of open evidence (Chapter 4), we will not only have a unified way to respond to these different skeptical challenges but also have a better understanding of where skepticism goes wrong according to an open knowledge view.

### 2.1 Skepticism and Closure

You know you have hands. You know that if you have hands you are not a handless brain in vat. Do you know you are not a handless brain in a vat? Visiting the city zoo, you see a zebra in the pen. You know that if something is a zebra, it is not a mule. Do you know that the animal in the pen is not a mule disguised to look like a zebra? Having seen the report on *CNN*, you know that it snowed in New York yesterday. You also know that if it snowed in New York, then anything indicating that it hasn't is misleading. Do you also know that if the *Daily* reports that it hasn't snowed in New York yesterday, it's misleading? A first theoretically uninvolved answer to all these questions, it seems, is 'no.' In fact, if the typical answer were 'yes,' it would not be clear why the sort of skepticisms that are generated from these cases ever bothered anyone. Yet despite the initial inclination the principle of knowledge closure entails a positive answer. In each of the cases, if you know the first two propositions, you must know the third. But then again, how can you know that you are not a BIV (Brain In a Vat), that the animal you see is not a disguised mule, or that the *Daily* is misleading if its report conflicts with the one you heard on *CNN*?

The cases above exemplify a tension between three common intuitions:

- (i) What is known to follow from an item of knowledge is itself known. (*Closure*)
- (ii) For many propositions  $p$  (that you have a hand; that there is a zebra before you; that it snowed in New York yesterday), you know that  $p$ . (*Non-Skepticism*)
- (iii) For some  $q$  that logically follows from  $p$  (that you are not a handless brain in a vat; that the animal before you is not a mule cleverly disguised to look like a zebra, that the *Daily* made a mistaken report if it did not report as *CNN*), you do not know that  $q$ . (*Skepticism*)

Each of these three intuitions has a hold on us, yet their (simple) conjunction entails a contradiction: *Closure* (i) and *Non-Skepticism* (ii) entail the negation of *Skepticism* (iii). In each of the cases above the prevailing response has been to deny *Skepticism* (which is the idea advocated most famously by G.E. Moore), while accepting *Closure*. More sophisticated responses offer some way of accepting (i), (ii) and (iii) but not all at once. The offer has been to contextualize, to relativize (to practical interests, to contexts of assessment), to idealize, etcetera. In contrast to this widely accepted strategy, Dretske proposed to relevant-ize, i.e. to deny (i) by claiming that the unknown propositions of (iii) are indeed unknown (but irrelevant). Since (ii) conflicts with (iii) only given (i), consistency can be maintained if the skeptical cases can somehow be deemed irrelevant. We can mitigate the unintuitive cost of denying (i) by specifying which cases of knowledge failure would entail loss of knowledge of the propositions that are known to entail them (surely everyone would agree that some cases of knowledge are dependent on knowing some of the known consequence, e.g. since  $p$  follows from  $q$ , that if you don't know that  $q$  you don't know that  $p$ ). Thus, we may hold on to intuitions (ii) and (iii) by denying the *universality* of (i) – viewing closure not as a *principle* that holds for knowledge generally, but only for some subset of known propositions. Yet with the help of Hawthorne's arguments we have seen that characterizing the subset in the way Nozick and Dretske have, can lead to inconsistencies.

It is not these more recent problems, but primarily the closure intuition (i) that lead theorists in other directions. Most notably Stewart Cohen (1988) and David Lewis (1996) claim that (i) through (iii) hold, but only for fixed contexts of knowledge ascription. Along the same lines, a recent view proposed by McFarlane (2005) is to relativize and keep (i) through (iii) for knowledge ascriptions within contexts of evaluation. Subject Sensitive (Moderate) Invariantists like Hawthorne (2004a) and Stanley (2005) propose that consistency be maintained not relative to contexts of knowledge ascription, but relative to the practical environment of subjects to whom knowledge is being ascribed. From this rough and incomplete characterization of

the current epistemological map, one thing surly stands out; knowledge closure and the avoidance of skepticism are theoretically central.

And yet the kind of skepticism that needs to be avoided according to different accounts of knowledge has not always remained constant across the different suggestions. The skeptical cases above, for instance, are very different. The proposition that if the report in the *Daily* conflicts with the report on *CNN*, it is the *Daily* that is mistaken, does not seem like an alternative to the proposition that it snowed in New York yesterday. In contrast, that I am a handless BIV *does* seem like an alternative to my having hands. Another difference is that some skeptical challenges appeal to systematic error, some don't. The current chapter spells out these challenges and marks some of the differences between them. The importance of this characterization for my own project is to be able to use it later within a wider context and provide a preliminary motivation for knowledge openness.

I do not claim that the proposals advanced by, e.g., the Contextualist or Subject Sensitive Invariantist, do not in fact solve the skeptical problems they set out to respond to. By appeal to the salience of alternative possibilities that make alternatives relevant in a context of ascription or in a subject's practical environment, or some other way, much headway has been made. Nevertheless, I suspect that these solutions will work with regard to cases that somehow make the possibility of *error* relevant. Yet salience of error does not cover the entire spectrum of possible cases that might be used for skeptical purposes. Although some skeptics are bound to appeal to alternatives, to contending incompatible propositions with what one knows that make error salient, they need make such an appeal. For a skeptical closure driven argument all that is needed are propositions that look to be unknown but follow from what is.<sup>65</sup> The *Daily/CNN* case does not fit the standard skeptical mold. Later I will show that given some plausible assumptions about knowledge and its relation to evidence, many such propositions do in fact follow from propositions that are known. This suggests, as I will argue in Chapter 4 and 5, that there are deeper structural properties that need to be understood regarding knowledge and ignorance. I will claim that these structural properties have to do with the relation of known propositions to the

---

<sup>65</sup> For example, Cohen (1988) defines relevance thus: "an alternative (to  $q$ )  $h$  is relevant (for  $S$ ) =df  $S$ 's epistemic position with respect to  $h$  precludes  $S$  from knowing  $q$ ." This definition gives Cohen the resources to solve the skeptical paradox he deals with: Cartesian Skepticism. Once an alternative is relevant in a context what it follows from becomes unknown by definition. But one point that is made in the present work is that we have no reason to suspect that one can know propositions that are not relevant even if they follow by closure from things we know. More specifically, propositions that are irrelevant (and hence known) need not be  $q$ -alternatives ( $q$  = the proposition from which consequences are inferred). So although Cartesian Skepticism is properly answered by the method Cohen suggests, this does not mean that we are out of the woods with respect to similar challenges that invoke propositions that do not straightforwardly seem like relevant alternatives. Similar (perhaps even more straightforward) remarks hold with regard to Lewis' definition of knowledge (1996).

evidence that supports them as well as logical relations that they have to other propositions. Some of the groundwork necessary for appreciating these features is laid out in this chapter in the form of a Cartesian Skeptical challenge. Besides this role, together with other skeptical challenges, Cartesian Skepticism will play a role in making clear the benefit open knowledge enjoys by having a way to respond to them (as well as explaining their force). If open knowledge can respond to skeptical and other central challenges of epistemology, this major benefit can be part of a favorable evaluation of its tenability.<sup>66</sup>

### 2.1.1 Cartesian Skepticism, Moorianism, and what Comes In-Between

Let us start with familiar cases that have been used by skeptics, cases that make salient the possibility of error. We have seen that the principle of epistemic closure serves as premise in arguments for several positions on the epistemological map. At the opposite extremes of this map are Skepticism and Mooreanism: The former will give up (ii) (*Non-Skepticism*) the latter rejects (iii) (*Skepticism*). Both positions seem problematic, but at the same time, the principle on which their arguments are founded, closure, is widely accepted as valid by all the aforementioned positions. It was also claimed that the intuitive plausibility of this principle induces most theorists to seek a way of avoiding the untenable extremes by means other than undermining closure. We have seen, though, some reason in the opening chapter to be cautious with how this intuition should be put to use.

A familiar skeptical argument, Cartesian Skepticism, runs along the rough sketch outlined above: You claim to know that you have hands. But if you are a handless brain in vat, you do not have hands, and you do not know that you're not a handless brain in a vat. Therefore, you do not know that you have hands. The way this argument is regularly represented is as follows:

- (1) I know I have hands.
  - (2) I know that if I have hands then there are external objects.
- Therefore,
- (3) I know there are external objects.

Moore takes (3) to follow from (1) and (2) since closure is valid, while Cartesian skeptics take the falsity of (3) to be a conclusion of a *reductio* argument. As the attempts forged by Contextualists and Subject Sensitive In-

---

<sup>66</sup> As I claimed earlier if knowledge is open it does not follow that Contextualism or Subject Sensitive Invariantism are false.

variantists indicate, both the skeptical and the Mooreian responses do not seem satisfactory. But though the natural reaction would be to deny the principle that both unsatisfactory positions share (see Ramesy's maxim below), closure is not often questioned. And yet, in the context of the skeptical argument this seems surprising. Theoretically it does *not seem* right to say that if I know I have hands and know that if I have hands I am not a handless BIV, then I know I am not a BIV. This, of course, is not to say that closure does not hold nor that it is not intuitive, but only that it is not in the context of the skeptical argument that closure is present in its most favorable intuitive light.

Contrary to what both Moore and some of the early skeptics may have thought, it seems that both drift from commonsense and are motivated by the intuitive pull of epistemic closure.<sup>67</sup> One who denies closure could explain why the two opposing views drift from commonsense and yet enjoy a certain theoretical appeal - an appeal that comes from a shared false assumption. P.F. Ramesy once claimed that regarding longstanding disputes between two opposing views that seem less than satisfactory "it is a heuristic maxim that the truth lies not in one of the two disputed views but in some third possibility which has not yet been thought of, which we can only discover by rejecting something assumed as obvious by both the disputants." (Ramsey, 1960: 115-6)

As the wealth of the theories above indicates, substantial improvements in epistemology have stemmed from attempts to find the underlying assumption as suggested by Ramsey's maxim. Relevant alternatives, contexts, practical environments have been invoked and developed in exciting and sophisticated new directions. Some forms preserve the essential feature<sup>68</sup> of knowledge closure others do not. And yet, despite the theoretical headway that has been gained in developing these theories, the possibility of resolving the skeptical-Moorian bind by denying closure should not be forgotten. Besides the direct considerations supporting knowledge openness, this position will receive more credence the more problems it can simultaneously resolve (or avoid). I believe that the wealth of such problems has been underappreciated.

Let me turn, then, first to consider more thoroughly the nature of Cartesian skepticism and its connection to knowledge closure. Subsequent sections of this chapter will have other forms of skepticism as their main focus.

---

<sup>67</sup> In addition to the reasons we saw in the first chapter there are other reasons to suspect our intuitions regarding closure. *Closure* is a technical term and so is not directly available for inspection with regard to common non-theoretical usage. I also think that principles are not the primary object of intuitions; cases are. And so any intuition regarding closure stems from such cases and not by the general motivations such as the claim that knowledge is expandable by inference. Moreover in chapter four we will see highly intuitive principles that are provably false, principles that have direct relations with knowledge closure.

<sup>68</sup> For the essentials of closure see Chapter 1 p. 29 note 43.

## 2.2. Cartesian Skepticism

Skepticism of a familiar Cartesian type, I argue, relies on two components: the principle of epistemic closure and the underdetermination of knowledge by existing evidence. The second component ultimately depends on an assumption that has been questioned, but which will not be discussed in detail in this chapter,<sup>69</sup> namely, that a subject  $S$  can know that a proposition  $p$  is true, even in cases where her evidence  $e$  does not entail  $p$ . More telegraphically and for further reference the thesis runs as follows:<sup>70</sup>

- (F) It is possible for  $S$  to know  $p$ , even though  $S$ 's evidence and background knowledge,  $e^*$ , does not entail  $p$ .

It is understood, or at least this is the way I want to interpret this claim, that (F) is interpreted probabilistically as follows:<sup>71</sup>

$$\neg[K_S(p) \Rightarrow (\Pr_S(p|e^*)=1)]$$

Where  $e^*$  is all of  $S$ 's evidence and background knowledge and  $\Pr_S(\bullet)$  is her rational credence function mapping any (pertinent) proposition into the interval  $[0,1]$ . I will assume that (F) is true, but this does not mean that I am hereby committed to the idea that there are no sub-domains of knowledge, such as perceptual knowledge, or *a priori* knowledge for which (F) would rightly be considered false. Thus if (F) is true its truth depends on sub-domains of knowledge. For instance, knowledge about the future, knowledge based on induction, and testimony (or "testimony" of instruments such as watches and thermometers) may be underdetermined by the evidence, while knowledge based on memory and perception is knowledge that is based on evidence that entails the known propositions in these restricted sub-domains. Which of the sub-domains we choose to focus on will change the subject matter of a dispute between a skeptic and her opponent, but as long as that sub-domain has underdetermined knowledge, the structure of the argument will remain unscathed.

Let me start by pointing out that if we have a disputed class of propositions  $D$ <sup>72</sup> for which (F) is true, then for a given body of evidence  $e^*$  entailed

---

<sup>69</sup> But see 5.3.3 below, especially 5.3.3.4.

<sup>70</sup> (F) has obvious affinities to the thesis of fallibilism according to, e.g. Stanley (2005: 127); Cohen (1988: 91); Pryor (2000: 518); and Feldman (1981: 266).

<sup>71</sup> Williamson denies this interpretation at least for what he calls evidential or epistemic probability. Here too, 5.3.3. is relevant.

<sup>72</sup>  $D$  I view as a domain containing propositions, hypotheses or theories. As noted in the main text,  $D$  could contain propositions about external world objects, about other minds, about the future, but also it could be characterized by method, i.e. propositions believed on the basis of induction, by perception etcetera. Within  $D$  I might talk about partitions, which, in accordance with the convention used in the literature relating to Bayesianism is not a set of subsets of a

by a proposition  $p$ , there will be another proposition  $p'$  that belongs to  $D$  that is incompatible with  $p$  and that entails  $e^*$ . We can see this if we take  $e^{**}$  to be a conjunction of all the propositions of body of evidence  $e^*$ . Since  $e^{**}$  *a priori* follows from  $(e^{**} \wedge \neg p)$  and is incompatible with  $p$ , we can be confident that for every  $p$  of  $D$  for which (F) holds, there is at least one such proposition. In fact any proposition that entails  $(e^{**} \wedge \neg p)$  will have the same feature and later we will see that this entailment assumption is not required.

Now in probabilistic terms, if both  $p$  and  $p'$  entail  $e^*$ , then their unconditional probability will be lower than their conditional one (no matter how close to 1 they are unconditionally):<sup>73,74</sup>

$$(1) \quad (p \Rightarrow e^*) \rightarrow [\Pr(p) = \Pr(e^* \wedge p)]$$

$$(2) \quad (p \Rightarrow e^*) \rightarrow [\Pr(p \mid e^*) = \frac{\Pr(p \wedge e^*)}{\Pr(e^*)} = \frac{\Pr(p)}{\Pr(e^*)} > \Pr(p)]$$

I am assuming here both that  $0 < \Pr(p) < 1$ ,  $0 < \Pr(e^*) < 1$  as well as that  $p$  is part of the disputed class of propositions  $D$ . Though they need not be cashed out probabilistically, these two features can be used for skeptical purposes.

### 2.2.1 Immodest Skepticism

A skeptic can utilize (2) in two distinct ways. First, she may invoke the following immodest<sup>75</sup> skeptical underdetermination principle:<sup>76</sup>

---

domain but rather a set of propositions that are mutually exclusive that when stated as a disjunction is a tautology, i.e. the probability of a partition is 1. A partition, then, is here viewed as a mutually exclusive and exhaustive set of hypotheses, propositions, theories, that have a common feature upon which we are focusing. For more details see the Appendix.

Relevant here is the fact that if in one sub-domain propositions are not underdetermined, its proposition might follow from propositions of a sub-domain that is. Unless we want to commit ourselves to a view according to which knowledge has different senses, the existence of an undetermined sub-domain has ramifications for all knowledge. (What the ramifications are will depend on how we choose to characterize the domains.)

<sup>73</sup> When the background information plays a distinctive role I will state things more explicitly. It should be understood that  $e^*$  incorporates this background information.

<sup>74</sup> As before I am using the double arrow " $\Rightarrow$ " to denote *a priori* entailment. But the hypothesis need not entail a proposition in order for the probabilistic feature in (2) to hold. What is necessary for current purposes is that  $p$  make  $e^*$  more probable that is unconditionally. See page 53 footnote 97 for proof.

<sup>75</sup> The reason I call it immodest relates to how it compares with a different skeptical tactic that I will turn to shortly, one that relies on a weaker underdetermination principle.

<sup>76</sup> Here and in some of the following, I speak of hypotheses and not of propositions. This is mainly in order to follow common custom. Despite the fact that this switch does make a difference, I do not think it will make much of a difference for issues raised here, and I will say so where it does make a difference.



Immodest Underdetermination (IU):<sup>77</sup> For all  $i, l$  evidence  $e$  and subjects  $S$ , if  $S$  knows that  $h_i$  is true, and  $h_i$  entails  $\neg h_l$ , then  $S$ 's total evidence  $e^*$  favors  $h_i$  over hypothesis  $h_l$ .

The immodest skeptic poses a skeptical hypothesis  $sh$  which together with IU threatens to undermine any claim for knowledge of a hypothesis from the disputed domain  $D$ . This tactic, however, seems less promising than the second way of utilizing (2) (which will be considered in the next section) and hence receives here a somewhat shortened discussion. The basic idea is that a supposed known hypothesis  $h$  will be shown to be unknown by an appeal to  $sh$ , where  $sh$  is a hypothesis that entails the evidence but is incompatible with  $h$ . Since it entails  $e^*$ , by (2) it receives evidential support from  $e^*$ . The argument then continues by claiming that  $e^*$  supports both and so does not favor  $h$  over  $sh$ , and hence, by IU,  $h$  is not known. Since  $h$  is chosen arbitrarily, the claim generalizes to the entire domain of  $D$ .

Notice, however, that what can be derived from (2) (assuming that the raising of probability is sufficient for  $e^*$  counting as evidence for  $sh$ ) is that  $e^*$  is evidence for  $sh$  if  $sh$  *a priori* entails  $e^*$ . What is needed in order to utilize IU for the skeptical purposes, however, is a quantitative claim not merely a qualitative one. It has to be shown that  $h_i$  is not *avored* by  $e^*$  and this involves not merely the that  $e^*$  is evidence for  $h_i$  and  $sh$ , but that the degree of evidential support that  $h_i$  receives from a subject's evidence is no greater than the degree of support that  $sh$  receives.<sup>78</sup> It is by no means an easy task to spell out quantitative evidential support. Indeed in the Bayesian literature on evidential support, there are several measures that have been proposed and little agreement on the correct one.<sup>79</sup> Perhaps here the skeptic will claim that it is the non-skeptic who needs to show that  $h_i$  is favored over  $sh$ . Perhaps she would add that any evidence that can be appealed to has already been accounted for in  $e^*$  and would be compatible with both  $sh$  and  $h_i$  (which we have assumed by letting both entail the total evidence).

I am not convinced that this is how the dialectic is properly characterized and in this respect the immodest skeptics themselves seem to be split.<sup>80</sup> Yet I am not going to try and settle this complicated issue since I think there is an interesting response that a non-skeptic can utilize. The response follows from many of the subjective Bayesian conceptions of evidential support. Given a certain prior subjective probability distribution,  $h_i$  will in fact re-

---

<sup>77</sup> Several epistemologists have advanced this principle (or a similar one) claiming that it is ultimately at the heart of the best version of Cartesian skepticism. Examples are Brueckner (1994), Janvid (2006, MS), Pritchard (2005), Vogel (2004), and Yalçin (1992).

<sup>78</sup> Indeed Vogel claims that this is the way to answer Cartesian skepticism. He claims that the strait hypothesis is supported to a greater extent by the evidence since it better explains the evidence. See Janvid (MS) for an argument that this reply is unsuccessful.

<sup>79</sup> See, for instance, Fitelson (1999), and Eells and Fitelson (2000).

<sup>80</sup> Regarding this claim see Brueckner (2005: 390). For an opposing view see Janvid (MS). Interestingly, Brueckner takes the skeptic to be denying something like the (F) principle.

ceive greater rational credence given  $e^*$  than  $sh$  will if one takes in the total evidential state by a method of Bayesian conditionalization. In other words,  $h_i$  will be supported to a greater degree (for  $S$  given  $e^*$ ) than  $sh$  will, if  $S$  is a rational Bayesian conditionalizer and her prior probability for  $h_i$  is greater than her prior probability for  $sh$ . The Bayesian multiplier (the evidence given the hypothesis divided by the evidence, i.e.,  $\Pr(e^*|h)/\Pr(e^*)$ , equals<sup>81</sup> the skeptical hypothesis' multiplier  $\Pr(e^*|sh)/\Pr(e^*)$ ), is the same and greater than 1, and so conditionalizing will result in greater rational support for the hypotheses that had greater initial subjective credence.<sup>82,83</sup>

Two reactions suggest themselves on behalf of the immodest skeptic. One is that this claim sounds like no more than a shot at pure faith and is not an epistemological answer at all. It sounds as though this non-skeptical Bayesian is answering an epistemological question with a biographical, subjective, (or contingent) answer, based on personal preference or bias. By an appeal to some convergence results this objection can be answered, I think. The details will not concern us here, though the Appendix to this manuscript has some information about convergence that is relevant. The basic idea is that it is rational to stick to your subjective credence since in the long run you are guaranteed (given enough evidence) to arrive at the truth. Part of the project of subjective Bayesianism is to provide a link between truth and agreement on the one hand, and subjective preferences on the other. Conditionalizing on shared evidence provides the link. The skeptic has given the non-skeptic no reason to forfeit her (rational) updating strategy on her subjective credence and this non-skeptic has excellent reason to have great confidence in this conditionalizing strategy. In fact, abandoning her evidence updating strategy is irrational, she would rightly claim.

A second reaction of immodest skepticism is to take the Bayesian multiplier itself as what "favoring" in the (IU) principle refers to. Only if the Bayesian multiplier is greater for  $h_i$  than it is for  $sh$  would a subject count as knowing that  $h_i$  is true. If the evidence itself, the claim would be, gives equal

---

<sup>81</sup> The reason for this is that the likelihoods (the probability of the evidence given the hypotheses) is 1 since the evidence follows from both hypotheses. Hence since the evidence is the same in both the Bayesian multipliers have the same value.

<sup>82</sup> I take this to be a big problem for IU-skepticism, but there are other more minor problems as well. I will mention just one. The immodest skeptic, it seems, needs to provide a clear sense of what it means for a hypothesis to be favored by the evidence. This will presumably depend on some conception or evidential support together with some way of measuring it. It then has to be shown that this conception is better supported than the other (Bayesian) accounts that make other predictions.

<sup>83</sup> I wish to make it clear that the question is not which hypotheses get more or less support from a new item of evidence or from the total evidence taken abstractly. Rather the question is which of the hypotheses is favored by *one's* evidence. What the evidence favors for one subject need not be what it favors for another. For instance, if one is more inclined to think that one is a BIV, then the evidence that does not rule out this hypothesis will favor a skeptical hypothesis over the others. Also, repeated conditionalization will result in greater divergence between one who believes that  $sh$  and one who gives greater credence to  $h_i$ . See Appendix.

support to more than one hypothesis as reflected by this Bayesian multiplier, the hypothesis will not be known. But the non-skeptic has no forthcoming reason to accept this interpretation and if she does, she ought to reject the IU principle. The Bayesian multiplier cannot simply be extracted from Bayes' theorem it reflects only part of what rational conditionalization predicts should be one's posterior probability for a hypothesis. To accept this immodest skeptical reaction is simply to abandon subjective Bayesianism.<sup>84</sup>

Although I will give some reasons for thinking that the Bayesian answers the immodest skeptical challenge, ultimately, it will not be my objective to try to settle this issue. What I aim to do is to provide an outline of an account both of skepticism and of how one can react to it within the open knowledge framework. More accurately, my aim so far was only to raise some doubts about the immodest skeptical argument that later will allow me to compare it with a better skeptical argument that is not susceptible to a Bayesian reply along the lines just mentioned. Once it is clear that the open knowledge view can deal with both modest and immodest Cartesian Skepticism (where other accounts seem to have trouble), this will give at least some support to the idea that knowledge is open.

### 2.2.2. Modest Skepticism

We have been considering a case where both  $h_i$  and  $sh$  entail the evidence, i.e.  $h_i \Rightarrow e^*$ , and  $sh \Rightarrow e^*$ . On a standard Bayesian conception of the *evidence for* relation, there is a qualitative claim and a quantitative claim regarding the relation between an item of evidence and a proposition supported by that item of evidence. (In reality the relation is at least a three-place relation between two propositions and some background knowledge that will here only be presented as a two-place relation for simplicity. The fourth element, a subject, will be represented when necessary as an index. Time can also be added as a fifth relatum when dynamics are pertinent.) The qualitative claim runs as follows and will play a major role in subsequent chapters. Here a slightly stronger principle will be used.<sup>85</sup>

---

<sup>84</sup> There is much more to be said here in defense of Bayesian conditionalization. It solves many epistemological puzzles, there are the convergence theorems that show why one should stick to one's conditionalizing strategy, Dutch book arguments etc. In another context Karl Karlander and I (forthcoming) have proposed a Bayesian reply to the Sleeping Beauty Problem. This is not to say that there are no problems with subjective Bayesianism. For a less but still somewhat enthusiastic approach, see Earman (1993).

<sup>85</sup> This principle does not concern the question of what features  $e$  must have in order to count as evidence. For present purposes one can think of it as certain knowledge.

Qualitative Bayesian Evidence (QBE):  $e$  is evidence for  $h$ , if and only if, the unconditional probability of  $h$  is less than the conditional probability of  $h$  given  $e$ .

$$E(e,h) \leftrightarrow [\Pr(h|e) > \Pr(h)]$$

Assuming (QBE) it is easy to see that  $e^*$  is evidence for both  $h_i$  and for  $sh$  (qualitatively speaking). This is shown by (1) and (2) above. But which is favored? Here, I claimed, prior probabilities are important. If the prior probability of  $h_i$  is greater than the prior probability of  $sh$ , then there is no way that the result of conditioning on  $e^*$ , a proposition that both  $h_i$  and  $sh$  entail, will change this relation. In fact as more and more evidence comes in, the difference between posterior probabilities will grow.<sup>86</sup>

$$(3) \Pr(sh) < \Pr(h_i) \quad \text{[assumption]}$$

$$(4) \Pr(sh | e^*) = \frac{\Pr(sh \wedge e^*)}{\Pr(e^*)} = \frac{\Pr(sh)}{\Pr(e^*)} < \frac{\Pr(h_i)}{\Pr(e^*)} = \frac{\Pr(h_i \wedge e^*)}{\Pr(h_i)} = \Pr(h_i | e^*)$$

This is enough to show that on one measure of evidential support  $e^*$  favors  $h_i$  over  $sh$ . And this accords with the intuitive conviction that it would be irrational to be more confident (or equally confident) given one's evidence in a skeptical hypothesis than one is in one's own favored (straight) hypothesis (if one was not previously inclined to give more credence to the skeptical hypothesis).<sup>87</sup>

Yet, the account of how to measure evidential support can be contested, of course. One could claim that what 'favors' in (IU) means, is not the absolute posterior value of probability that  $e^*$  confers on a proposition, but rather what counts is some other measure - perhaps, as stated above, the Bayesian multiplier measure could serve this end. This reply would have to be backed up by a justification for some conception of quantitative measure of evidential probability that would pull in the skeptical direction. The prospects of such a conception of evidence, I think, are grim. A Bayesian, at least, can fend off the skeptical challenge in a non-question-begging way since the idea of evidential incremental measures such as the one I have used here is not

---

<sup>86</sup> Put in another way, this shows that  $h_i$  is favored at least subjectively. Convergence results, as outlined in the Appendix, show the epistemic connection between this subjective favoring and the objective favoring, i.e. convergence on the truth.

<sup>87</sup> Intuitively, if one initially favors the skeptical hypothesis, one will not know that the straight hypothesis is true. This means that there is place for epistemic luck or subjective bias in coming to know a proposition. Much more can be said in this regard, but I will leave it for another occasion. Some remarks in Appendix are relevant.

developed in response to skepticism and is as well established in its own right as can be expected.<sup>88</sup>

Before moving on to a more promising line of argument on behalf of the skeptic, let it be noted that another way to deal with the skeptical immodest challenge is to simply deny the principle that underlies it – to reject IU. As several authors have claimed, this would entail a denial of epistemic closure and would therefore not be an attractive alternative for many epistemologists.<sup>89</sup> In the present context, however, this is not a deal-breaker. It is an added advantage of open knowledge that it can solve both skeptical challenges.

A second skeptical tactic is to make an initially more modest claim (although the aim is just as ambitious):

Modest Underdetermination (MU): for all  $i$  and subjects  $S$ , if hypothesis  $h_i$  is in the domain of dispute  $D$  and a subject  $S$ 's total evidence  $e^*$  counts in favor  $h_i$ , then  $S$  does not know that  $h_i$  is false.<sup>90</sup>

---

<sup>88</sup> We have been assuming that both hypotheses entail the evidence. A skeptic may challenge this assumption. Strictly speaking there is no entailment relations in some cases. The hypothesis that one is faced with a table does not a priori entail that one has an experience as if there was a table in front of one (and the rest of the subject's evidence is not even mentioned). Or the hypothesis that there is a zebra in the pen does not entail seeing a zebra looking animal in the pen. The simplification I have been using, however, is not what is at issue. There does not seem to be anything stopping one from articulating the cases so that the evidence will follow, e.g. that one is facing a table that causes one to have certain table like experiences (and one can state the other relevant background evidence as well).

I wish to register something that I am neglecting and which does need special care. I have been talking as if  $sh$  is a hypothesis, perhaps a very complex one. This really cannot be so if what is needed is entailment. In order for a hypothesis to entail the evidence the hypothesis would have to change over time and in accordance with the (purported) knowledge under investigation. At this point one might pose a dilemma for the skeptic. If the hypothesis is predictive, then it really does deserve proper attention and should be taken seriously by anyone (at least until it fails repeatedly to predict the evidence). If it is not predictive, then one needs to formulate it as a post factum hypothesis. But if it is a post factum hypothesis it ought to gain no support from the evidence.

Appealing though it may seem, it would be hasty to discard skepticism on the basis of such a dilemma in its present form. The reason being that the common non-skeptical conception of reality is no more predictive than the skeptical one. My view is that they are both models or at the very least propositions that figure in an argument schema (the skeptical argument schema) rather than hypotheses. Only together with some auxiliary hypotheses will these propositions entail the evidence. This is why I earlier noted (p. 44, note 76) that the use of the term "hypothesis" can make a difference.

<sup>89</sup> Brueckner (1994) claims that (IU) and closure are equivalent and Pritchard (2005) argues that (IU) is logically weaker than closure. Unfortunately they are both concerned with justification and not knowledge and considering their arguments would take me too far a field. For a different view on this matter closer to the one I follow here, see Cohen (1998).

<sup>90</sup> If  $S$  antecedently knows that  $h_i$  is false, then the total evidence does not count in favor of  $h_i$  since presumably she would have to have some evidence in its favor that is not overrun by counter evidence. This can be shown by conditionalization. So I need not use the following more detailed principle:

$$\forall(i)[E_S(e^*, h_i) \Rightarrow \neg K_S(\neg h_i)]$$

It should be noted at the outset that (MU) coupled with a principle of knowledge closure will give the same result that (IU) was designed to provide. I will turn to show this after saying some things about (MU).

The (MU) principle can be justified directly. Suppose that one's total evidence raises the probability of a hypothesis  $h_i$  that was not previously known to be false. There will be many cases in which this condition holds without one thereby coming to know that  $h_i$  is true. It seems dubious, then, that one would learn that  $h_i$  is false by means of probability raising evidence to its truth.<sup>91</sup> The modest skeptic I have in mind is appealing to the same kind of challenge that the open knowledge advocate will press in pages 97 and 126.

This direct intuitive basis can be backed up by more fundamental principles. If one accepts (as many do) that the *evidence for* relation is defined by the raising of probabilities, i.e. one accepts (QBE), and furthermore one accepts that in order to know that  $h_i$  is true, one needs evidence for it (the total evidence to count in its favor), (MU) follows.<sup>92</sup>

The point now is, that together with the principle of epistemic closure - (CP), (MU) spells trouble for the non-skeptic (Bayesian or otherwise). To make things a bit easier, let us revert to a simplified formulation of (CP):

If there is a time  $t_0$  when my evidence was not good enough to know  $h_i$  is false, and if my total evidence at a later time  $t_1$  relative to  $t_0$  counts in its favor, then I do not know that  $h_i$  false.

I will later consider a related principle utilized to challenge closure more generally in page 126.

<sup>91</sup> This claim will be explored more extensively in page 126 and in the proceeding sections.

<sup>92</sup> In order to see this, it should be made plain that we are dealing with a domain of dispute  $D$  for which both the skeptic and her opponent accept (F). Next, that  $D$  includes pairwise incompatible hypotheses, i.e. a mutually exclusive and exhaustive partition  $Q$ . Let's assume that we have a proposition  $h$  that antecedently is not known, that we accept (QBE), and that knowledge requires evidence (that one's total state of evidence raises the probability of some hypothesis if it is to count as evidence for it). The idea is now to show that if the total evidence supports a proposition  $h$ ,  $\neg h$  is not known by an appeal to principles we have already accepted. I will assume that the two place predicate  $E(x,y)$  will relate to the total evidence that a subject has ( $e$ ) for a proposition  $y$  and that it supports  $y$ .

- |  |   |
|--|---|
| 1. $E_S(e,h)$  | [Assumption]  |
| 2. $\forall(i)(K_S(h_i) \rightarrow E_S(e,h_i))$                     | [Assumption that knowledge requires S's total evidence needs to count in favor of the known hypothesis] |
| 3. $\forall(i)(E_S(e,h_i) \rightarrow \neg(\Pr(h_i) \geq P(h_i e)))$ | [Follows from (QBE)]  |
| 4. $[\Pr(h e) > P(h)] \leftrightarrow [\Pr(\neg h e) < P(\neg h)]$   | [Follows from the axioms of probability]  |
| 5. $\Pr_S(\neg h e) < \Pr_S(\neg h)$                                 | [1,3,4]   |
| 6. $\neg E_S(e, \neg h)$   | [3,5]   |
| 7. $\neg K_S \neg h$   | [MP, 6 and 2 contraposition]  |

What this argument shows is that if one's total evidence counts in favor of a hypothesis it will not allow one to know that it is false on the assumptions made so far (and total evidence includes antecedently known propositions). Another way of showing this more rigorously would include times but would be more relaxed with respect to what the total evidence is. For current purposes we need not enter into these issues.

$$(CP) \forall(i)\forall(I)[(K_S(h_i \wedge (h_i \Rightarrow h_i)) \rightarrow K_S h_i)]$$

Now the modest skeptic with some logical and imaginative acumen, can grant that for a proposition in the disputed domain  $D$ , the evidence  $e^*$  favors  $h_i$  over its rivals. (In terms of probability, the probability of  $h_i$  on ones total evidence is greater than it is for any incompatible hypothesis.) However, utilizing (CP) she can find another proposition  $\neg sh$ , that *a priori* follows from  $h_i$ , and yet,  $\neg sh$  is not known given the evidence if it is not known antecedently.<sup>93</sup> The reason for this in probabilistic terms is that  $sh$ 's prior unconditional probability (no matter what it is) is not as great as the conditional probability of  $sh$  on the total evidence  $e^*$ . In other words, the probability of  $\neg sh$  goes down. Thus (CP) predicts that it is known yet if one cannot know without evidence we must (by *reductio*) abandon our conviction that  $h_i$  is known whether it is favored or not.

So the general structure of the skeptical argument that is here recommended for the modest skeptic with regard to a given disputed domain  $D$ , runs as follows (where  $Q$  is a partition of the disputed domain):<sup>94</sup>

- i. S knows that  $h_i$  is true, where  $h_i \in Q \subset D$ . [Assumption for skeptic's *reductio* argument]
- ii. S's total evidence  $e^*$  counts in favor of  $h_i$ . [From (i) and assuming that one cannot know a proposition that is not supported or even favored by one's evidence.]
- iii.  $\forall(i)\exists(I)[(h_i, h_i \in Q) \wedge (h_i \Rightarrow \neg h_i) \wedge (h_i \Rightarrow e^*)]$ <sup>95</sup>
- iv. Evidence  $e^*$  supports  $h_i$  (in terms of probability, the conditional probability of  $h_i$  on  $e^*$  is greater than the unconditional probability of  $h_i$ ). [iii, (1)-(2) above]
- v. S does not know  $\neg h_i$  is true. [MU]
- vi.  $\neg h_i$  *a priori* follows from  $h_i$ . [iii]
- vii. S knows  $\neg h_i$  is true. [(i), (PC) and assuming the skeptic runs through an argument that shows this]
- viii. S does not know  $h_i$  is true. [v, vii, by *reductio* and since (i) is the only assumption]

Though the argument might seem complex the idea is basically simple. In contrast to immodest skepticism, modest skepticism relies on two compo-

<sup>93</sup> I do not want to enter the dispute of whether standard skeptical hypotheses are known to be false or not. My point is that whether the standard ones are known to be false or not, for a domain that respects (F), there will be hypotheses that have the features that a modest skeptic can appeal to.

<sup>94</sup> I do not claim that this is an original presentation of the skeptical challenge. Stewart Cohen if not in detail, at least in spirit, is perhaps the first to suggest this argument on behalf of the Cartesian Skeptic. See Cohen (1998). Thanks to Professor Hawthorne for pointing this out.

<sup>95</sup> A proof is on p. 98. Also, the second conjunct is superfluous since we have defined  $Q$  as a partition.

nents. First, that we can prove that for any proposition which does not *a priori* follow from one's evidence, there will be other propositions that entail the evidence and are incompatible with the hypotheses. This will be shown below (p. 98). Second, that one does not know that the skeptical hypotheses are false. Why? Because the hypotheses entail the evidence and so if one did not know antecedently that they are false, getting evidence in their favor is not going to help (MU). This does not connect directly to the skeptics objectives since there is no real threat to the knowledge of the hypotheses for which the subject does gain evidence. This is where (CP) comes in. It ties the lack of knowledge of the falsity of the skeptical hypotheses to the purported knowledge of any hypothesis that is incompatible with them. This tie in the immodest skeptical argument is supplied in the modest one by closure without assuming that the hypothesis that the non-skeptic claims to know is not favored by her evidence. In simple terms it is one thing to claim that one has no evidence against a skeptical (*a posteriori*) hypothesis (and this is what the modest skeptic relies on), it is quite another matter to claim that one does not have any reason to favor one's own initial belief over a rival skeptical hypothesis. The former only claims that the skeptical hypothesis cannot be known to be false on the basis of no evidence. The later needs a much stronger claim, i.e. that one's evidence counts as much in favor of the skeptical hypothesis as it does for one's own favored mundane one. This claim we have seen cause to suspect.

Let me note some interesting features that become clearer with this argument in focus, features that will be central in later discussion. The modest skeptic is in a position to make the skeptical argument only by appealing to a specific hypothesis tailored to entail  $e^*$ . But there are two ways to customize the hypothesis. The first is to forget about general good-for-everything skeptical hypotheses that are good for all knowledge claims of given domain. Thus the modest skeptic will not pose a brain-in-a-vat-skepticism type challenge, for instance, but rather a-brain-in-a-vat-having-experiences-of-a-certain-kind challenge. The reason the modest skeptic needs specific hypothesis is that if  $sh$  is not tailored specifically to *a priori* entail  $e^*$ , nothing would seem to guarantee that S would not know that  $\neg sh$ . She would have no way of showing that  $e^*$  supports  $sh$ , it seems. A consequence of viewing things in this way is that  $sh$  would target only a specific proposition, though repeating this process in a more or less systematic way will make it clear to the non-skeptic that there is no point in going on. Thus this type of argument would succeed only by attrition. But another way to look at modest skepticism is to look at the argument schematically. A hypothesis will only entail the evidence together with auxiliary claims to allow for a proper entailment to hold between it and the evidence. In fact, however, though the modest skeptic may need auxiliary hypotheses, the evidence  $e^*$  need not be entailed. It is enough to show that given the skeptical hypothesis, the probability of



the evidence is raised.<sup>96</sup> So the point is that the modest skeptic can start with weaker assumptions, i.e. with hypothesis that together with auxiliaries will make the evidence more probable than the evidence is unconditionally. That will suffice to bring about the same conclusion the skeptic is after, namely, (viii).

This brings me to the main point. One simple way to avoid skepticism is to deny (CP), i.e. to deny that knowledge is closed under proper inference. The above (rational) reconstruction of Cartesian skepticism, or Modest Skeptics, gives reason to believe, first, that the original item of knowledge is supported by the evidence available to the subject, evidence that has not been undermined by the skeptic's argument. And second, that while the evidence has not been undermined and supports the supposed known hypothesis, the skeptical hypotheses are not known to be false. This is why it is natural to think that knowledge is not closed. In other words, the natural theoretical reaction to Modest Skepticism is that while the skeptical hypotheses are not known to be false, the proposition from which the negation of these hypotheses follow, are known. It should be admitted, I think, that while there may be other ways to respond to this skeptical challenge, the open knowledge advocate does have the tools to respond to it and on the face of it this seems to be the correct reply.

I am well aware that there are several ways one might attempt to answer the modest skeptic (e.g. trying to show that one does know the skeptical hypothesis are false, or claiming that awareness to *sh* in the context of knowledge ascription changes the standards required for a proposition to be known). The point of this chapter (and those that follow) is not to survey or criticize those attempts. I will attempt to make the open knowledge reply to the modest skeptical argument more appealing by placing it within a wider context along with other challenges it can simultaneously stand up to.

## 2.3 Mundane Skepticism

We have seen the workings of Modest Cartesian Skepticism. This section focuses on another type of skepticism. To see its potential to undermine the idea that we have knowledge of everyday empirical claims let me get back to

---

<sup>96</sup> A proof that all that is needed (for a skeptic) is that the probability of the evidence is raised by her skeptical hypothesis, i.e. she need not appeal to hypotheses that entail the evidence:  $[\text{Pr}(e|sh) > \text{Pr}(e)] \rightarrow [\text{Pr}(sh) < \text{Pr}(sh|e)]$ .

- |      |  |                  |
|------|--|------------------|
| i.   | $\text{Pr}(sh e) = (\text{Pr}(e sh)/\text{Pr}(e))\text{Pr}(sh)$  | [Bayes' theorem] |
| ii.  | $[(\text{Pr}(e sh)/\text{Pr}(e))\text{Pr}(sh) > \text{Pr}(sh)] \Leftrightarrow [\text{Pr}(e sh)/\text{Pr}(e) > 1]$ | [i]              |
| iii. | $[(\text{Pr}(e sh)/\text{Pr}(e) > 1) \Leftrightarrow [(\text{Pr}(e sh) > \text{Pr}(e))]$                           | [ii]             |
| iv.  | $[\text{Pr}(e sh) > \text{Pr}(e)] \rightarrow [\text{Pr}(sh) < \text{Pr}(sh e)]$                                   | [i,ii,iii]       |

the broader discussion concerning Moore who argues for (3) from (1) and (2) as opposed to Cartesian Sceptics who claim that  $\neg(1)$  from (2) and  $\neg(3)$ :

- (1) I know I have hands.
  - (2) I know that if I have hands then there are external objects.
- Therefore,
- (3) I know there are external objects.

Besides closure, the Moorean and the Cartesian Skeptical arguments share one other feature - both arguments relate to an entire field of knowledge. Moore's argument is meant to prove that we have knowledge of a wide range of empirical truths, namely, those relating to the existence of objects in the external world. As his Moorean opponent, the Cartesian skeptic also relates to an entire domain of knowledge. Cartesian Skepticism represents a family of skeptical claims aimed at undermining all knowledge of the external world<sup>97</sup> (or of other domains) by raising the possibility of systematic error about it. But error does not have to be systematic. A different type of skepticism is couched in Dretske's example, trading on the possibility of disguised mules. As a skeptical argument Dretske's example will run as follows: Standing in a zoo, looking at a zebra, so the skeptic argues, one does not know that it is a zebra if one cannot tell that it is not a mule disguised as a zebra.

This type of skepticism – call it *Mundane Skepticism* – does not target an entire field of knowledge, but rather depends on existing knowledge and works piecemeal. Let us look at its structure and potential more closely. It has by now become familiar to analyze such skeptical arguments into the following general form: If S knows that  $p$ , she must know all propositions  $q$  that she knows to logically follow from  $p$ . (Since the skeptic has just mentioned it, S knows that  $p$  entails  $q$ .) Furthermore, since she does not know that  $\text{not-}q$  (for some cleverly devised  $q$ ), she consequently does not know  $p$ . This is the argument the Mundane Skeptic rallies against our everyday commitment to particular cases of knowledge: “you must know all that you know to logically follow from your knowledge and you cannot deny that you know that  $\text{not-}q$  follows from what you say you know (I have just shown you that it follows). Since for all you know  $q$  may be the case, you do not know that  $p$ .”

Several writers<sup>98</sup> have noticed the potential of this sort of skepticism separating it from the more famous (or infamous) Cartesian Skeptic.<sup>99</sup> To see the

---

<sup>97</sup> Not all knowledge, since one knows (if Descartes was right) that one exists. If Williamson is right one will know *a priori* that there is at least one false belief (1986).

<sup>98</sup> These include Vogel (1990, 20-3), Cohen (1998: 155), and Hawthorne (2004a: 5). In relation to Mundane Skepticism, Cohen notes (1998: 155) that in contrast, (immodest) underdetermination skepticism only works with the kind of skeptical cases that are strongly underdetermined. I would rather make the distinction by an appeal to anti-realism. The anti-realists

force of the challenge set by this latter type of skepticism let us consider a few more examples. You go to McDonalds and order a Big-Mac hamburger. You pay and receive your meal wrapped in the usual McDonalds wrapping, and the taste is as usual. It is safe to say that you know you are eating a McDonalds Big-Mac. Your only problem is that you have invited your skeptic friend to come with you. She comes up with the following challenge: “For all you know the kitchen might have had a major breakdown. The manager – who is the brother of the manager of the next-door Burger King - might have borrowed 100 hamburgers from his sister, one of which you are now eating.”<sup>100</sup>

Let's consider an example of knowledge of the future: Say Marilyn claims to know that she will go to the movies in an hour. All the mundane skeptic has to do is raise one of the following questions: “Do you know you will not be hit by lightning this afternoon? Do you know your family won't come for a surprise visit?” (No analogy between these types of events is intended.) Naturally, Marilyn cannot pretend to know that these possibilities will not materialize. The skeptic can thus remain confident in her ability to come up with these scenarios time and time again (unless, as we will see later, the evidence guaranties the truth of the know proposition). Granted, some cases are harder for the skeptic than others and some may even be impossible. Nevertheless, by committing ourselves to closure, it seems that the mundane skeptic can force us to forfeit a large bulk of our empirical knowledge. In fact, since she does not rely on the plausibility or even the mere possibility of what most would consider outlandish and improbable scenarios such as the ones Cartesian skeptics have been using, mundane skepticism is in some

---

have no answer to Mundane Skepticism even though they may have a ready reply for Cartesian Skepticism. Vogel and Hawthorne take Mundane Skepticism to be primarily an issue regarding lottery-like propositions.

<sup>99</sup> Perhaps the first to consider Mundane Skepticism closely as apposed to Cartesian Skepticism is J.L. Austin. A defense of this interpretation of Austin can be found in Kaplan M. (forthcoming).

<sup>100</sup> This and like cases are important for several reasons only some of which will be considered in the main text. One important consideration which I will not enter into in detail has to do with the claim made by e.g. Cohen (1988) and Vogel (1990), that skepticism of this sort turns on the statistical or lottery-like nature of the background information that allows us to know propositions such as “I'm eating a Big-Mac hamburger”. In this case, this suggestion seems implausible unless “statistical” just means that we attach some credence or subjective probability to the eventuality of a major breakdown in the kitchen and the manager borrowing hamburgers from the next-door Burger King. Many cases do have the feature of being based on probability estimates of the lottery kind. I also agree that one does not know propositions are true on mere statistical background information. My contention is that these things need not be connected to one another. The example suggests (and Cohen (1988) agrees) that some consequences of things we know are such that would we consider them on the basis of our evidence directly, we would not know them to be true. My claim is that the class of such propositions is wider than what has been suggested. Once it is admitted that there is a gap between knowledge and evidence, there seems to be no reason to think that all propositions that may be used by a mundane skeptic will be statistical. My view is that the class is determined by the evidence.

respects more threatening than the Cartesian skepticism. All the MS skeptic needs is an ability to imagine scenarios incompatible with something that is entailed by what we claim to know and which is not ruled out by our evidence. This skepticism, therefore, depends on an ability rather than a type of underdetermined single (schematic) scenario.

Consider the case from the last chapter from Vogel (1990) – knowledge gained by memory: Say you parked your car in a parking space directly in front of your house. 10 minutes later a visiting skeptic asks you whether you know where your car is. ‘Sure’, you unsuspectingly reply, ‘it is right in front of my house!’ The skeptic’s reply is by now almost too familiar: ‘You say you know where your car is. Does that mean you know it hasn’t been stolen?’ Assuming that your car is still in front of your house and that it has not been stolen and returned in the expanse of 10 minutes, your knowing that the car is in front of your house entails that the car has not been stolen. But obviously you do not know that your car has not been stolen if all you have to go on is your memory of where you parked it. Then again, as long as you are not a skeptic yourself, you do not want to say that under normal conditions when asked whether you know where your car is, you shouldn’t honestly say ‘yes’.<sup>101</sup>

As the above cases illustrate, Mundane Skepticism is generated by appeal to the ability to devise scenarios the negations of which follow from the purportedly known propositions, but which are not supported by the evidence one has. Mundane Skepticism is threatening since it relies on very little and turns on assumptions (perhaps unavoidably) endorsed by its opponents, namely, that knowledge is often based on non-conclusive evidence, that the scenarios appealed to (whose truth is entailed by the (purportedly) known proposition  $p$ ) are not supported by the evidence, and finally, that the principle of epistemic closure is valid.

But issues are not altogether clear-cut. From the vantage point of the skeptic, Mundane Skepticism is both more and less advantageous than Cartesian Skepticism. On the one hand, its relative weakness lies in the fact that since it works piecemeal, it does not undermine entire fields of knowledge at once. Also, Mundane Skepticism depends on the evidence available to the subject. If you are sitting in your car, the Mundane Skeptic will not be able to undermine your knowledge that it has not been stolen (at least not by the same argument used before). Mundane Skepticism may not work when the non-skeptical opponent has second order knowledge.<sup>102</sup> On the other hand,

---

<sup>101</sup> If one is concerned that other evidence in some or all of the above cases is needed in order to know in the first place, the cases should be viewed in a way that takes all of the evidence under consideration. If this evidence does not a priori entail the known propositions, the (MS) skeptic has good reason to think that an (MS) scenario can be found. For a proof of the existence of such propositions see page 98.

<sup>102</sup> Depending on what it is that one sees as required for knowledge of knowledge, this in itself does not diminish the threat Mundane Skepticism poses. Say you know that  $p$  and have

MS relies not on what is conceivable or possible, but rather on known realities. While it may be possible that we are all brains in vats (or at least that I am), we definitely do not know that we sometimes are. We do, however, know that cars get stolen, lightning strikes, people suffer heart attacks, families make surprise visits, people die and hamburgers might get switched.

Placing higher demands on knowledge would not help keep the Mundane Skeptic at bay unless we demand that knowledge be based on conclusive evidence, i.e. that knowledge would be based on evidence that make the proposition certain. This will keep the skeptic at bay only at the cost of our becoming skeptics ourselves. Settling for higher demands that fall short of conclusively supported knowledge would bring us in one respect closer to the skeptic (since we would then agree at the outset that the set of known propositions is smaller than usually conceived) yet it would be harder for the Mundane Skeptic to come up with challenging cases. Nevertheless nothing essentially is changed if we just make things harder (for ourselves and for the Mundane Skeptic) in this way, and so I find little promise in looking for a solution in this direction.<sup>103</sup>

It seems that Mundane Skepticism resists at least some of the strategies developed against Cartesian skepticism. Idealist or verificationist positions might be immune to the challenge set by Cartesian skepticism. The verificationist can respond to a Cartesian skeptic by claiming that a skeptical hypothesis of the sort she invokes is not true. Depending on the specifics of the account, the reply would either be that what cannot be verified by all possible data is simply false, or meaningless, or that it is not true. In any case, the point is that sentences that are by their nature underdetermined are not true according to anti-realists. (Though admittedly this is a reply that works only against skeptics who invoke hypotheses that are underdetermined by all possible data).

But Mundane Skepticism's propositions do not allow the verificationist to take advantage of the under-determination feature of such scenarios or hypotheses (e.g. the world was created 5 minutes ago with dinosaur bones and all the required "historical" furniture). The MS propositions are not underdetermined by all possible data, they are underdetermined merely by the available evidence. Hence these proposition (or sentences) are simply too robust for a verificationist tactic.<sup>104</sup> Take the following dialog as an example:

S (Skeptic): Do you know what time it is?  
V (Verificationist): Yes, its 5:00

---

extra evidence regarding your evidence so that you qualify as knowing that you know *p*. The Mundane Skeptic would first target the evidence that you know that you know, and if successful, in a further stage one's first order knowledge would be challenged.

<sup>103</sup> See section 5.3.3 regarding Williamson's suggestion to define epistemic probability so that all known propositions have probability 1 on the knower's evidence.

<sup>104</sup> I am indebted to Professor Williamson for this point.

S: How do you know this?

V: Look at my watch - it says "5:00."

S: But how do you know your family has not played a trick on you by secretly setting your watch one hour back. Surely you can't rule this out by looking at your watch.

The verificationist can hardly appeal to her semantics to save her from such mundane scenarios.

## 2.4 Live Skepticism

A third kind of skepticism different from Mundane Skepticism and Cartesian skepticism has recently been suggested by Bryan Frances (2004). The type of skepticism he suggests - Live Skepticism (LS, for short) - is generated by the following principle.

The Live Hypothesis Principle: ... If S is as aware as just about anyone that [*h*] is live... and that [*p*] entails  $\neg$ [*h*], then if [*h*] isn't ruled out with respect to S, S doesn't know [*p*].

After introducing this principle and some surrounding (plausible) assumptions, Frances argues that some very central claims can fit the role of *h* in the Live Hypothesis Principle (or at least that there are close possible worlds in which *h* is a live hypothesis). For instance, I think that I know that I believe that Obama is the US president. But if I know that I so believe and know that if I have this belief then those who claim that there are no beliefs (some eliminativists)<sup>105</sup> are mistaken, then I know that eliminativists are mistaken. Now although it is probably true that I do have such beliefs, it is also true that I don't know that these theorists are mistaken since they are just as informed as I am (probably more informed about this matter than I am) and I have nothing special to say that would make other theorists convinced that I know that they are wrong. These theorists disagree with belief-eliminativists but they would also admit that eliminativism is a live option (or that it could easily become a live option in the sense that it is an improbable yet real contender for replacing the leading accounts). They might even be able to mention future results in psychology that would convince them (though they would be surprised if things turned out that way). Now if the theorists who believe in beliefs do not think they know that the eliminativists are wrong, how could I claim I know them to be wrong?

---

<sup>105</sup> Frances suggests Stich, the Chirchlands, Rorty, Feyerabend, Quine and Dennett as examples. See Frances (2004: 38).

Frances, for different reasons from those given by Dretske and Nozick, believes that closure is invalid.<sup>106</sup> Yet he claims that what Dretske and Nozick have to “say against the [closure] principle does not, in my opinion, give any reason at all to reject the Live Hypothesis Principle.” (2004: 26) I disagree, at least partially. Although I agree that Drestke's relevant alternative theory does not have the required resources for rejecting the Live Hypothesis Principle, and although I am not sure about Nozick's sensitivity requirement<sup>107</sup> (the open knowledge condition), I think a proper open knowledge account would find it faulty. The basic reason is simple (though the full details will only be given in Chapter 4). Assuming that the eliminativists have a way to accommodate all the (good) reasons I have for believing that I believe Obama is the US president, I do not have any evidence for the proposition *the eliminativists are wrong*. Assuming I cannot know without evidence, I do not know them to be mistaken. Of course, this is only a preliminary sketch, but it shows how ultimately LS depends on knowledge closure and that open knowledge seems to have the resources to keep LS contained.

## Chapter Summary

We have seen three skeptical challenges. Cartesian Skepticism, the first of the skepticisms we have seen, uses epistemic closure as premise together with the Modest Underdetermination principle (p. 49). The scenarios that are typically used are not accidentally far-fetched. In order for them to threaten fields of supposed knowledge (domains that are defined as disputed by the type of proposition or method of knowing), these scenarios must be highly underdetermined (though they need not be absolutely underdetermined). Also, Sceptics use cases that are incompatible with what is known so that they can be used in a *reductio* argument that uses closure as premise. The feeling of unease that one might feel when presented with Moore's argument is evidence (at least some evidence) that these scenarios are not known to be false. And this unease gains support from the plausibility of the modest underdetermination principle.

In contrast with the Cartesian Skeptic, the Mundane Skeptic is initially more reserved than the Cartesian Skeptic. The order of quantification is what is at issue: The Cartesian Skeptic claims that there is a scenario that is not known to be false such that for any proposition in the domain of dispute, e.g. empirical propositions that are not *a priori* certain (*I am thinking*, for instance), the negation of this scenario follows from any one of these propositions. In contrast, the Mundane Skeptic relies on an ability to conjure up a

---

<sup>106</sup> See Frances (1999).

<sup>107</sup> See 1.5.2 (above) for the details regarding Nozick's open knowledge tracking account.

scenario for any proposition in the domain that is not *a priori* certain (or follows from one's evidence). That the Mundane Skeptic has a right to be optimistic regarding the existence of such scenarios is shown by the proofs in pages 98 below and 53 above. (Further analysis of this feature of underdetermined knowledge will be given in due course.)

It was also claimed that both the Cartesian and Mundane Sceptics (as well as the (LS) skeptics), share a tacit assumption that hinders the scope of their suggested cases. They share the idea that they need to find *alternatives*, i.e. cases that are incompatible with what is purportedly known that compete with the known proposition for being the best explanation of the evidence, favored by the evidence, or receive the greatest probability from the evidence. But with closure they need not do this, all that needs to be assumed is that propositions that follow from what is known are themselves not known to be true under the circumstances. One way of showing this is by showing that these alternatives have a greater conditional probability on the evidence than they do unconditionally. Another way is to employ propositions that do not contend for being explanations of the evidence at all – that is, propositions that are not in any real sense alternatives (besides the fact that their negation is incompatible with the known proposition). The potential for this kind of skepticism will be the focus of the next chapter with regard to the implausibility of epistemic ascent: Dogmatism (Kripke, Forthcoming), Easy Knowledge (Cohen 2002, 2005) and Bootstrapping (Vogel, 2000, 2007)).

Frances' Live Skepticism, a skepticism that relies on the ability to use knowledge of everyday propositions to infer the falsity of philosophical and scientific theories and thereby come to know them to be false (assuming that they are), was briefly considered. I claimed, contrary to what Frances says regarding this matter, that if one denies closure (at least for the right reasons), one can avoid the threat posed by this type of skepticism.

The centrality of closure is hard to dispute, yet the considerations regarding closure have not usually been of a direct nature. The way the dispute has been conducted is mostly by way of asking whether there is a cheaper way to get out of the skeptical bind. The open knowledge advocates though they are quick to mention their ability to respond to the skeptic have nonetheless argued for knowledge openness by appealing to their respective theories. Not much has been done to explain how or why knowledge is closed or open regardless of any knowledge account. The way I propose to think about the issue is as directly as possible. Nevertheless I do not think the benefit of open knowledge that can be drawn from having a ready non-skeptical argument ought to be ignored in thinking about the issue.

The challenge that the closed knowledge advocate faces is to account for one of two alternatives. First, where does the added property originate from that allows the negation of skeptical claims to be known by inference? Since evidence is not available, something else needs to fill its place. Deduction itself does not give us the extra property, and so it seems like either we sys-



tematically underestimate the kind of evidence we have, or knowledge is not closed. Alternately, second, one needs to account for why it is that when one deduces a conclusion that is not known from a premise that is, knowledge of this premise is lost. In other words one needs to have an account of why coming to believe that a skeptical hypothesis is false destroys the knowledge of the premise from which it was inferred. The evidence for the premise has not been lost so something else needs to be appealed to as an explanation of this phenomenon. While in some cases it seems plausible that awareness of skeptical claims raises doubt as to whether a proposition is known, other cases do not lend themselves as easily to this tactic. Such is the case with regard to the scenarios that will be the main concern of the next chapter.

# Chapter 4: Evidence and Open Knowledge

## Chapter outline

The current and next are the central chapters of this monograph. Though the argument might at times be involved, it is driven by a simple idea: in the analysis of knowledge, the logic of evidence should play a pivotal role. A proper account of knowledge, in other words, must be constrained by facts about the relation of evidential support.<sup>162</sup> Appealing as this idea may seem, even among contemporary epistemologists who address evidence in their theories, little attention has been given to the actual workings of evidence. A proper understanding of constraints on the relation between knowledge and evidence, I will argue, has ramifications for epistemology that are as wide-ranging as they are fundamental. Specifically, since the relation of evidential support is not closed under known entailment, there is good reason to believe that knowledge is also open. Another way of stating the objective of these next two chapters is as setting a challenge for epistemic closure: if, as I argue, the openness of evidence can be established, how can knowledge be closed? The current chapter argues that evidence is open and the next ties evidence openness to knowledge openness.

In this chapter I first - section 4.1 - present a simple case that seems to pose an intuitive challenge to knowledge closure. I then argue that regardless of intuitions (which actually seem to accord with the analysis), the case is one where one lacks evidence for what follows from what is known. Using probabilistic reasoning, I show that evidence is not closed, i.e. that one may have evidence for  $p$ , know that  $q$  follows from  $p$  *a priori*, and yet have no evidence for  $q$ . I turn then - section 4.2 - sketch a preliminary argument based on evidence openness for the conclusion that knowledge is open. In section 4.3 I show that rejecting the idea that the “evidence for” relation is open cannot be maintained on the basis of refusing to view this relation as one that entails that probabilities have not been lowered, or refusing to view the relation probabilistically. These arguments will also make plain the way

---

<sup>162</sup> There are two central questions that one might focus on in relation to evidence. First, one might ask; what is evidence? What does something have to be in order to be counted as evidence? Second; what is the “evidence for” relation? What kind of relation is it? My main focus is on the second question, specifically, if the relation is closed under known inference. Some of the comments in this chapter, however, have more to do with the first question.

in which Hawthorne's arguments (of Chapter 1.5.1) can be answered by an open knowledge advocate.

## 4.1 The Watch Case and its Probabilistic Analysis.

You look at your watch and see that it reads "3:00". Assuming that the time actually is 3:00 o'clock and that all other things are normal, you know that the time is 3:00. By trivial reflection you also know that *if the time is 3:00 o'clock, then if your watch reads "3:00", it is showing the correct time*. Do you know that *if your watch reads "3:00", it is showing the correct time*? Do you know, just by looking at it, that *even if the watch has stopped, it is showing the correct time*?

Intuitively, it does not seem that you do. Perhaps you already know beforehand - relying on other sources - that your watch is working properly, or at least, that it is accurate. But if you don't, it does not seem like the kind of thing that can be known on the basis of the fact that the watch shows "3:00".

And yet, epistemological orthodoxy says that you do (or, at least, that you often can) know this. Since knowledge is closed under known entailment, the claim goes, a belief properly derived from a known proposition is itself known. Having derived the belief that *if my watch reads "3:00", it is showing the correct time*, from your knowledge that *the time is 3:00 o'clock*, you know this conditional is true. Knowing that your watch shows "3:00", you can derive the consequent of the conditional and hence know that *your watch is showing the correct time*.

Why hold fast to this counter-intuitive conclusion? The answer, as in many similar cases, is the principle of epistemic closure. The watch reading example, however, brings out not only the counter-intuitive consequences of closure, but also theoretical reasons for thinking that it fails in cases of this sort. To put it succinctly, the reason we tend to deny the status of knowledge to the conclusions of such inferences is that they lack evidential support. In what follows I wish to elaborate and support this claim by analyzing this example in detail. This analysis will, in turn, serve my more ambitious attempt to motivate knowledge openness and lay bare its benefits.

### 4.1.1. Evidence and Probabilities

Your reading of the watch provides you with the evidence in virtue of which your belief that it is 3:00 o'clock gets to count as knowledge. But what does it mean that reading "3:00" off the watch is evidence for your belief? One

common rendering of this relation is in terms of conditional probability.<sup>163</sup> Let “ $p$ ” denote the proposition that *the time is 3:00*, and the proposition that *the watch reads “3:00”* will be “ $e$ ”. We can then see that the probability that *the time is 3:00* given the evidence (which in this case is  $e$ ), is greater than the probability of  $p$  without this evidence:  $\Pr(p|e) > \Pr(p)$ .

Even if it cannot be accepted as a definition of evidential support, it seems that any account of evidence should grant the following criterion:<sup>164,165</sup>

---

<sup>163</sup> The notion of probability that I am appealing to here is a more or less a standard construal of subjective Bayesian probability. The prior probabilities reflect, either possible betting behavior or are read off of other sorts of behavior with no small bit of idealization. To dodge the well-known Bertrand Paradox as well as for other reasons, I am not very optimistic about objective readings of prior probabilities. Let me just say here (though these issues will be considered more carefully in this and the next chapter), that I take evidence to either be a special kind of knowledge, roughly, certain knowledge, known knowledge or less than certain knowledge using Jeffrey Conditionalization, to represent the proper change from prior to posterior probability. Let me also note that I sometimes treat probabilities as if they amount to 1 for the purposes of ease and clarity. I don't see any special problem with doing this for the purposes of certain examples. As far as I can tell all the cases will yield the same results without this simplification.

Central here is the assumption that there is a way to define prior probabilities that can then be compared with conditional probabilities in accordance with the axioms of probability. This allows a comparison between an event's likelihood in relation to a body of evidence prior to the availability of the item of evidence under consideration, with its probability after it becomes available. Surely the issue is vexed, yet by leaving the method of setting the priors pretty much open, I hope I can get things started without, so to speak, loading the dice in anyone's favor.

Together with *a priori* propositions, I do treat necessary *a posteriori* propositions, e.g. “I exist” and, if Williamson is right, “there is at least one believer” (Williamson 1986) etc, as having a prior probability of 1, but not much else. In particular, although in some cases it seems plausible to associate probability 1 to known propositions, I do not accept that this is the case across the board. Some of the arguments proposed below can be reformulated as a challenge to those who, following Williamson (2000: 184-237), view knowledge as always having probability 1 and 5.3.3 spells some of these challenges out more directly.

<sup>164</sup> Note that this criterion is weaker than the identification of the *evidence for* relation with the raising of probabilities.

<sup>165</sup> Some may be worried that not all evidence is propositional, that experiences, for instance, such as the experience of a blue patch in one's visual field, may be evidence for one that there is something blue in the vicinity. If you have such worries, take as the relata figuring in (EC) (and the other evidence principles below) the proposition that *S is experiencing a blue patch in his field of vision*. I propose this measure only in order to sidestep this thorny issue.

Let me just note in passing that I view Williamson's argument (Williamson 2000: 194-200) from the fact that we use evidence in, e.g. inferences to the best explanation, to the conclusion that evidence is propositional, as somewhat too quick. One might view the evidence as represented in such inferences as I do here and still insist that the evidence itself, e.g. experiences, or perceptions, is non-propositional. Moreover, one might also view the probabilities as representing e.g. the non-propositional brain states and the causal relations these states have with behavior. That I am in a certain brain state, for instance, might be correlated with behavior probabilistically. For helpful conversation regarding this issue, thanks to Professor Kathrin Glüer-Pagin.

(EC) Necessarily, if  $e$  evidentially supports  $p$ , then the probability of  $p$  given  $e$  is not lower than the prior unconditional probability of  $p$ .

$$\Box(E(e,p) \rightarrow (\Pr(p|e) \geq \Pr(p)))$$

“ $E(e,p)$ ” means that  $e$  evidentially supports  $p$ .<sup>166</sup> How does your situation vis-à-vis the accuracy of your watch fare with respect to this principle? From  $p$  it follows that *if the watch shows “3:00”, then the watch is showing the correct time*. The antecedent of this conditional is just  $e$ , materially implies the consequent of this conditional (call it “ $q$ ”)<sup>167</sup> assuming  $p$ ; that the time is indeed 3:00. Hence one can know by mere reflection that:<sup>168</sup>

$$(1) p \Rightarrow (e \rightarrow q)$$

It follows by closure that if you know  $p$ , you know:

$$(2) e \rightarrow q$$

But do you have evidence for (2)? Presumably, if you do, it must be the evidence that facilitated knowledge of  $p$  in the first place (or some other aspect of your evidential situation), namely,  $e$ . But if it is a necessary condition on evidence that it not decrease the probability of that for which it is evidence (EC), then  $e$  does not provide (2) with evidential support. This is because the conditional probability of (2) on  $e$  is not greater than the probability of (2). In fact, since  $e$  verifies the antecedent of (2), given  $e$ , the probability of the conditional is *lowered*. The reason for this is that the truth of  $e$  excludes all the cases in which (2) is true in virtue of the falsity of its antecedent. So in fact,

$$(3) \Pr(e \rightarrow q|e) < \Pr(e \rightarrow q)$$

This can be proven assuming that  $\Pr(e \wedge \neg q) > 0$  (that there is some, perhaps even minute, probability that the watch shows “3:00” and it is not 3:00):

$$i. \Pr(\neg(e \rightarrow q|e)) = \frac{\Pr((e \wedge \neg q) \wedge e)}{\Pr(e)} = \frac{\Pr(e \wedge \neg q)}{\Pr(e)}$$

<sup>166</sup> I will mostly suppress representing the necessity operator and quantification.

<sup>167</sup>  $q$  can be viewed as a conjunction of two claims, that your watch reads “3:00” ( $e$ ) and that it is 3:00 ( $p$ ).

<sup>168</sup> Since  $q$  is just  $p \wedge e$ , we have:  $[p \Rightarrow (e \rightarrow (e \wedge p))] \Leftrightarrow [p \Rightarrow (e \rightarrow q)]$ , and since the left hand side of the a priori equivalence is a tautology, the right side is a tautology (and knowable as such).

- ii.  $\frac{\Pr(e \wedge \neg q)}{\Pr(e)} > \Pr(e \wedge \neg q) = \Pr(\neg(e \rightarrow q))$  [assuming that  $0 < \Pr(e) < 1$ ]
- iii.  $\Pr(x|y) > \Pr(x) \leftrightarrow \Pr(\neg x|y) < \Pr(\neg x)$  [follows from the axioms]
- iv.  $\Pr(e \rightarrow q|e) < \Pr(e \rightarrow q)$  [i, ii, iii]

Assuming in (ii), as in this case one should, that the unconditional probability of  $e$  is not 1 or 0, (3) results.<sup>169</sup> And from (EC) and (3) it follows that  $e$  is not evidence for  $(e \rightarrow q)$ . It is this lack of evidence that explains why, although properly derived from a known premise, (2) is not known.<sup>170,171</sup>

Similar considerations explain a host of other examples some of which are proposed in the literature as challenges to the validity of epistemic closure. We have already seen some in the previous chapters: It follows from something's being a zebra that it is not a mule disguised to look like a zebra. And yet, seeing a zebra-looking animal in the pen, although providing one with evidence that there is a zebra in the pen, does not provide any evidence that the animal is not a disguised mule. In fact, that there is a zebra-looking animal in the vicinity is, at least to some extent, an indication that there is a zebra-looking disguised mule in the area. Memory of having parked one's car in the driveway ten minutes ago evidentially supports the belief that one's car is in the driveway. It provides no evidential support for the entailed belief that one's car has not been stolen in the last ten minutes. Perceiving one's hands is evidence that one has hands, it is not evidence that one is not a bodiless brain in a vat. The same holds for the Dogmatism problem, the Easy Knowledge problem, Bootstrapping, other epistemic ascent cases, Live skepticism, Mundane skepticism and many more that may not seem intuitively problematical yet under analysis will be revealed as unsupported by one's evidence.<sup>172</sup> Examples of all these sorts abound. What is common to

---

<sup>169</sup> Much thanks here to Peter Pagin for first proving the above claim. Also, many thanks to Karl Karlander for making the necessary assumptions clear. The proof I give here connects easiest to other arguments in this and the next chapter.

<sup>170</sup> If one is worried that I have not taken all the evidence under consideration, perhaps not taking into account background knowledge of the reliability of one's watch, one can simply interpret  $e$  as including this background knowledge nothing short of entailment will matter.

<sup>171</sup> The problem discussed here is akin to the Easy Knowledge problem discussed by Hawthorne (2004a: 73-8) introduced in Cohen (2002) and further developed in his (2005). As mentioned in the previous chapter the issue discussed here does not relate to the question of whether the method for knowledge gain is right, whether the knowledge is basic, or easy. As will become apparent, the core issue, I believe, relates to the failure of evidence closure and differs significantly from Cohen's and Hawthorne's discussion with respect to the analysis and solution of these problems. Regardless of the differences I am indebted to them for their groundbreaking work on these as well as related issues.

<sup>172</sup> This last claim is important to appreciate since the fact that one has not evidence for a proposition is a good reason to think that she does not know it to be true. And yet, since these

all, I claim, is the failure of the evidence for the originally known proposition to carry over and support the inferred proposition. Lacking evidential support, it seems, empirical beliefs of this sort do not qualify as knowledge (though I will claim later that they do qualify as justified beliefs). The same idea, then, accounts for a large number of unhappy consequences of epistemic closure. For instance, having proper evidence that  $p$  is true can allow one to know  $p$ , but not that the means by which the evidence was acquired are reliable (the bootstrapping problem), or that evidence against  $p$  is misleading (the dogmatism problem).<sup>173</sup> Proper evidence can provide warrant in believing that  $p$ , but does not supply one with reasons for believing that this evidence is not misleading.

Admittedly, denying the status of knowledge to properly inferred beliefs exemplified in these cases has its cost, namely, the rejection of the intuitive and extremely popular principle of epistemic closure. Yet in view of a theoretical explanation of why and under what conditions it might fail, and in light of the vast epistemic challenges it overcomes, it would be irresponsible not to give it more serious attention. This chapter aims to explain why evidence is not closed and why the connection evidence has with knowledge, together with the plausible (perhaps even inevitable) assumption of knowledge fallibilism, open knowledge seems to result.

I shall also propose an open knowledge reply to the Hawthorne's arguments of the first chapter, an argument we saw Dretske's and Nozick's accounts do not have the resources to handle. I will also try and show why some ways to circumvent knowledge openness in light of evidence openness are ultimately not workable, and this will be further developed in the next chapter. The challenge, I will claim, remains: if knowledge is not always based on conclusive evidence, and if evidence is not closed, how can knowledge be closed?

---

cases are not readily available to our attention, we will not tend to deny knowledge of the premise from which the unknown proposition was inferred.

<sup>173</sup> Notice that my formulation of the problem is more general in that it does not rely on the intuition that knowledge is gained too easily nor on the intuitive oddity of bootstrapping oneself into knowledge of the reliability of one's sources. As will become evident in what follows I rely on structural features of evidence and the principles governing the relation of evidential support. Contrary to what some have alleged, the denial of closure is motivated not merely by the desire to avoid Cartesian skepticism, or by specific accounts of knowledge (like Nozick's or Dretske's). Besides the argument here relating to the structural features of evidence, since epistemic closure is implicated in a wealth of epistemic puzzles, including, in addition to those already mentioned, the lottery paradox (see Vogel 1990, Hawthorne 2004a), some of the semantic self-knowledge puzzles, and more (see Chapter 1), open knowledge enjoys the benefit of solving or defusing many, perhaps even all, of these puzzles.

## 4.2. Probabilistic Argument for Evidence and Knowledge Openness

In the course of analyzing the watch case, I have claimed on probabilistic grounds and the (EC) principle, that (a) evidence is open, i.e. one can have evidence for a proposition  $p$ , know that  $q$  follows from  $p$ , yet the evidence for  $p$  cannot count in favor of  $q$ . That is, we have seen a proof that shows that the evidence that raises the probability of a proposition  $p$  can lower the probability of a proposition that *a priori* follows from it.<sup>174</sup> (b) that this phenomena accounts for many cases that seem intuitively to be problematical.

A tacit appeal was made to the thesis of fallibilism that has already played a role in previous chapters, the idea that one can know that  $p$  on the basis of evidence and background knowledge that does not entail  $p$ . This assumption was explicit in the proof with regard to the watch case when it was assumed that  $\Pr(e \wedge \neg q) > 0$ . This assumption need not be true in many cases, for example, it is not true with regard to mathematics when one's evidence on the basis of which one comes to know a mathematical truth is a proof. But given that the probability of a known proposition is less than 1, the argument for knowledge openness is *prima facie* compelling.

To make the argument more explicit, let me state it schematically.

- (i) *Fallibilism* - One can know that  $p$  even though one's total evidence does not guarantee that  $p$ . That is, one can know that  $p$  even though the probability of  $p$  given one's total evidence is less than 1.
- (ii) (ii) (EC) - If the probability of a proposition  $p$  is lower given  $e$  then it is un-conditionally,  $e$  is not evidence for  $p$ .
- (iii) (KE) - If one knows that the  $p$  is true, one has evidence for  $p$ .
- (iv) So, Knowledge is not closed.

(iv), I will argue, follows from (i), (ii) and (iii). But this argument is in need of elaboration. Specifically, not only do the premises need support, there are further claims that need to be established before it is made clear why (iv) follows from the premises. I will repeat the argument of chapter two that for all propositions  $p$ , if all available evidence  $e$  (which could be a long conjunction) does not conclusively establish that  $p$ , then there is a proposition  $q$  such that  $q$  follows from  $p$  and  $e$  lowers the probability that  $q$ . With this claim established, for all subjects  $S$  and every known proposition  $p$  that does not follow from  $S$ 's total evidence, there will be a proposition, in fact many

---

<sup>174</sup> Unless we are dealing with Multi Premise Closure, the probability of a proposition that follows from  $p$  cannot be any lower than the conditional probability of  $p$  on the evidence. Nevertheless, it can lower its prior probability; cases where its unconditional probability is greater than the conditional probability of  $p$  on  $e$ . I will have more to say about this later in this chapter.



propositions, which are not supported by the totality of S's evidence which can be properly inferred from  $p$ . Hence, given that in order to know that a proposition is true, S needs to have evidence for it, knowledge is open. So the basic idea is that since evidence is open, so is knowledge. The more detailed argument, then, is as follows (*Open Knowledge Argument*):

- I. *Fallibilism* - One can know that  $p$  even though one's total evidence does not guarantee that  $p$ . That is, one can know that  $p$  even though the probability of  $p$  given one's total evidence is less than 1. [assumption]
- II. For any fallibly known proposition  $p$  (in the sense of (I)) there is a proposition  $q$  that *a priori* follows from  $p$  such that the conditional probability of  $q$  given all the available evidence is lower than the unconditional probability of  $q$ . [independent proof]
- III. (*EC*) - If  $e$  is evidence for  $q$  then the conditional probability of  $q$  given  $e$  is not lower than the unconditional probability of  $q$ . [assumption]
- IV. For any fallibly known proposition  $p$ , there are propositions that *a priori* follow from it which the available evidence does not support (the probability of the proposition is lowered). [I, II, III]
- V. (*KE*) For all subjects S and propositions  $p$ , if S knows that  $p$ , then S has evidence for  $p$ .<sup>175</sup> [assumption]
- VI. *Knowledge is open* - For all subjects S and propositions  $p$  and  $q$ , if a proposition  $p$  is fallibly known by S to be true, there is a proposition  $q$  such that even if S knows that it follows from  $p$ , S is in no position to know  $q$  by competent inference (keeping the evidence fixed).

Several of the premises of this argument can be challenged, of course, namely, (I), (III) and (V). Also the argument can be taken to be a *reductio* for one of the premises. I take Cartesian skepticism to be making such an argument to the conclusion that no proposition is fallibly known.<sup>176</sup> Yet (II), I now argue, cannot be challenged.

No matter how high the probability of  $p$  given  $e$  is, since the probability of  $e \wedge \neg p$  need not be 0 (given fallibilism), the probability of  $\neg(e \wedge \neg p)$  will be

---

<sup>175</sup> One might want to restrict these claims to empirical propositions. The restriction depends on one's conception of evidence. One might take proofs and *a priori* reasoning as a type of evidence, or not. If you believe that there are propositions that need no evidence to be known, restrict the claim in any way you see fit. Surely there are many types of propositions that have to be supported by available evidence in order to be considered as known - it is those that we need to quantify over.

<sup>176</sup> Since a Cartesian Skeptic wants to challenge almost all empirical knowledge, she will have to argue that all empirical knowledge is fallibly known. The dream or demon hypothesis is supposed to supply a motivation (among other things) for accepting this premise. Chapter 2 focuses on this matter. At this stage let me just note that a more promising line for the skeptic to take for substantiating the premise is to claim that if evidence does not entail a proposition, that proposition cannot be known. This would place the non-skeptic in the uncomfortable position (if she is interested in showing the skeptic to be wrong) to either jettison closure, or argue that empirical knowledge is not fallible.

lower given  $e$  than it is unconditionally (assuming that  $e$  is S's total evidence).

- (1)  $K_S(p) \wedge (\Pr(p|e) < 1)$  [Fallibilist assumption]
- (2)  $\Pr(e \wedge \neg p | e) = \frac{\Pr((e \wedge \neg p) \wedge e)}{\Pr(e)} > \Pr(e \wedge \neg p)$
- (3)  $\Pr(\neg(e \wedge \neg p) | e) < \Pr(\neg(e \wedge \neg p))$  [(2), Kolmogorov axioms]
- (4)  $K_S(p \Rightarrow \neg(e \wedge \neg p))$  [assuming that S so reasons]
- (5)  $\neg E_S(e, \neg(e \wedge \neg p))$  [(3), (EC)]
- (6)  $\neg K_S \neg(e \wedge \neg p)$  [(5), (KE)]

Assuming, that is, that  $e$ 's unconditional probability is greater than 0 less than 1 ( $0 < \Pr(e) < 1$ ). (3) has a feature that is not essential to the argument. For the argument to go through suffice it that the probability of the evidence is lowered by the proposition that follows from the known proposition. Also, the argument could have terminated at (3), the point of continuing to (6) was to place (3) in the context of the open knowledge argument: Since S knows that  $p$  and does not know a proposition S knows logically follows from  $p$ , knowledge is open.

As I just mentioned, any proposition that follows from  $p$  and lowers the probability of  $e$ , will have the same features, i.e., it follows from  $p$  and yet the evidence lowers its probability.<sup>177</sup> For every fallibly known proposition there are, then, a host of propositions that follow from one's knowledge that are not supported by the totality of one's evidence, and are hence (given KE) not known.

Rejecting (II) as a way to resist the open knowledge argument does not, then, seem promising. Another way to respond to the argument is to reject (I). But I think it cannot simply be rejected, one would need to come up with some reason to think that given one's evidence, the probability of a known proposition will always be 1. One reason for rejecting (I) is equivalent to skepticism<sup>178</sup> and can be improved on by accepting the conclusion of open knowledge, i.e. open knowledge is better than skepticism (other things being equal). Another way has been proposed by Williamson and will be considered in detail in the next chapter. So let us see how well one might do by questioning the other premises. Here are the options (at least those I can think of) for one who wants to resist the open knowledge conclusion (VI):

<sup>177</sup> See proof in page 53, note 96, Chapter 2.

<sup>178</sup> Other ways of rejecting the first premise will be considered in Chapter 5.

- (a) Evidence can lower probabilities: It is not the case that whenever the probability of  $p$  given  $e$  is not as great as the unconditional probability of  $p$ ,  $e$  cannot be evidence for  $p$ .
- (b) Evidence is not probabilistic: No account of the relation of “evidence for”, or more generally perhaps, of the relation of “reasons to justifiably believe” can be given in terms of probabilities.
- (c) Knowledge is not dependent on evidence: There are some cases where one knows that a proposition  $p$  is true, even though one has no evidence for  $p$ . Perhaps all that is needed, in some cases, are general reasons (or suitable initial probabilities) with no special evidence concerning  $p$ .

The next section contains an argument aimed at defusing (a) and (b). No appeal to probabilities is needed, I will argue, in order to reach the conclusion that evidence is open. (c) will be one of the subjects of the next chapter (5.2). To be sure, (c) (and what it is aimed at) is in need of refinement. The purpose here was to display in broad strokes the central ways to avoid the conclusion of the argument for knowledge openness.

It should be noted that responding to the argument by seeking a refined notion of evidence does not seem promising. I have not assumed any special conception of evidence nor does the argument presuppose that all knowledge requires evidence. Suffice it that some knowledge requires evidence (of whatever non conclusive kind) and one may fill out the details regarding the nature of this evidence in many ways.

In sum, we have seen so far a case – the watch case – that intuitively poses a challenge for knowledge closure. This intuition was reinforced by considerations concerning probabilities and the (EC) principle. The argument first established evidence openness since the watch case itself is a clear probabilistic counterexample to evidence closure, and the argument continued by claiming that this openness of evidence can be used as an argument for open knowledge. Since evidence is open, it was argued, assuming the possibility of fallible knowledge, there is a compelling argument to the conclusion that knowledge is open as well. Three major directions of responding to this argument were sketched and rejoinders to these responses will be argued for in further sections.

### 4.3 Evidence Openness from Principles

The previous section included an argument for knowledge openness that appealed to evidence openness assuming that evidence conforms both to the axioms of probability and to (EC). Two reactions to the argument were men-

tioned that I will respond to here. One is that (EC) fails miserably, evidence can lower the probability of a proposition and count as evidence for it, and this happens systematically with regard to all the cases where an entailed proposition has its probability lowered by the evidence that supports one's knowledge. This was option (a) above. Option (b) states that evidence is simply a non-probabilistic notion. The argument below defuses both these responses by an appeal to non-probabilistic reasoning showing that on very weak assumptions, it is more or less an established fact that if one's evidence for a proposition  $p$  is non-conclusive, then one does not necessarily have evidence for  $q$  even if one knows that  $q$  follows from  $p$  *a priori*.

A second objective of this section is to defend open knowledge from John Hawthorne's argument (of Chapter 1) aimed at refuting it. Hawthorne's argument, as we saw in Chapter 1, shows that in accepting weaker principles than the closure principle, leading figures of open knowledge (Nozick and Dretske) have tacitly or explicitly accepted inconsistent commitments. The current section will show that not only does Hawthorne's argument not refute a correct open knowledge account; his argument helps reveal the underlying structure of open evidence which can be used as an argument to support it.

Imagine you are looking for zebra-look-alike mules. Where would you look? It would be natural for you to do so among zebra-looking animals. Admittedly, zebra-looking mules would be hard to find, but if there is any chance of finding some (at least the ones which are well disguised) you had better search among zebra-looking animals. Seeking them among the elephant-looking animals, or the banana-looking objects holds little promise of success. Your chances of encountering a zebra-looking mule are slim. But the probability that an object encountered is *not* a zebra-looking mule is even lower when a zebra-looking animal (say, a zebra) is visually observed. Although it does not constitute strong evidence, a zebra-looking animal gives some support to the proposition that a given object is a zebra-looking mule. In probabilistic terms, the presence of a zebra-looking animal raises the probability that a mule disguised to look like a zebra is present.<sup>179</sup>

Now in the normal case, when one sees a zebra-looking animal, one has evidence that the animal is a zebra. But anyone who knows that a zebra is not a mule, must realize that at the same time that one gains evidence for  $Z$  (= there is a zebra in the pen) in this way, one *loses* evidence for  $\neg DM$  (= it is not the case that there is a disguised mule in the pen), or, in other words, one *gains* evidence for  $DM$ . Denying this quickly gets one into serious trouble in trying to provide a plausible account of evidence. The argument below shows why (in non-probabilistic terms).

---

<sup>179</sup> There is no essential probabilistic point here. All that I am claiming is that some measure of support is given to  $DM$  by a visual observation of a zebra looking animal. This will be argued for more directly below.

The following principle, it seems, must be a part of any plausible theory of evidential support:

*Consistency of Evidence* (CS): For all evidence  $e$  and propositions  $p$ , if  $e$  evidentially supports  $p$ ,  $e$  does not evidentially support the negation of  $p$ .

*Evidence addition* (EAD): For all evidence  $e$  and propositions  $p_1$  and  $p_2$ , if  $e$  evidentially supports  $p_1$ ,  $e$  evidentially supports  $p_1$  or  $p_2$ .

*Evidence equivalence* (EEQ): For all evidence  $e$  and propositions  $p_1$  and  $p_2$ , if  $e$  evidentially supports  $p_1$ , and  $p_2$  is logically (or *a priori*) equivalent to  $p_1$ ,  $e$  evidentially supports  $p_2$ .

The latter two principles are designed to emulate principles concerning knowledge that Hawthorne employed in the service of his argument against knowledge openness, but other principles could be used as well. (EAD) is an evidential analog of Hawthorne's (AD) which states (roughly) that adding disjuncts to a known proposition results in knowing the disjunction. (EEQ) emulates the (EQ) principle stating that if one knows that  $p$  is *a priori* equivalent to  $q$ , then one knows  $q$  if one knows that  $p$ . We will have occasion to come back and consider Hawthorne's arguments, shortly. For the time being, note that the three principles enjoy a high degree of intuitive appeal. The first principle, the consistency of evidence, expresses the simple idea that if something is to count as evidence for some theory, hypothesis, proposition or what have you, it cannot also support its negation. Or, in other words, that a proposition supporting both a hypothesis and its negation, does not constitute evidence for either. The principle (EAD) stems from the idea that evidential support is a relation that is closed under certain logical operations. Addition captures the idea that adding disjuncts to a supported hypothesis does not undermine the degree of support for a disjunction.

Equivalence - (EEQ) - expresses the idea that "confirmation of a hypothesis is independent of the way in which it is formulated" (Hempel 1965: 13). The truth-values of logically equivalent hypotheses stand or fall together, so equivalent hypotheses must also be supported together.<sup>180</sup>

Although the principles appear plausible, their conjunction with (even a particularly weak version of) the thesis of underdetermination of theory by

---

<sup>180</sup> In terms of a coarse-grained possible world semantics, we might say that any evidence that the actual world is one of the  $p_1$ -worlds (the possible worlds in which  $p_1$  is true) is also evidence that the actual world is a  $p_2$ -world, since in those terms the sentences express the same proposition. The (EEQ) principle is justified by the claim that any evidence that the actual world is one of the  $p_1$ -worlds, is also evidence that the world is an  $p_2$ -world if " $p_1$ " and " $p_2$ " are true in the same worlds. I might add that I don't want to lay too much weight on this kind of reasoning in support of (EEQ) since possible world semantics faces trouble of its own when it comes to closure principles.

evidence, leads to a contradiction. Strong underdetermination is the contentious claim that all possible evidence cannot fully determine the choice between (some) mutually incompatible theories. Weak underdetermination (henceforth: UD), however, states that, at least insofar as actual evidence goes, there can be two theories (or more) that are supported by the evidence but incompatible with each other.<sup>181</sup>

(UD) A certain body of evidence may support two or more theories, hypotheses or propositions that are incompatible with one another.<sup>182</sup>

Here are some examples. Consider first the case of competing interpretations of formulas of quantum mechanics. The two leading interpretations of quantum theory are, apparently, compatible with all (possible) observations. And yet, since one, the Copenhagen interpretation, entails that every particle has a momentum and the other, the Bohmian interpretation, implies that particles have no momentum,<sup>183</sup> the two are mutually incompatible. Presumably, the evidence we have supports both interpretations. Thus, there is evidence supporting the Copenhagen interpretation (CQM). By (EAD) it follows that this evidence also supports: *the Copenhagen interpretation is true or the Bohmian interpretation is false* ( $CQM \vee \neg BQM$ ). But this is equivalent to: *it is not the case that the Copenhagen interpretation is false and the Bohmian is true* ( $\neg(\neg CQM \wedge BQM)$ ), and so the evidence supports this latter proposition as well (by EEQ). Now since the truth of one interpretation entails the falsity of the other, *the Bohmian interpretation is true* (BQM) is equivalent to *the Bohmian interpretation is true and the Copenhagen interpretation is false* ( $\neg CQM \wedge BQM$ ). Thus by evidence equivalence (EEQ) the evidence supports the claim that *it is not the case that the Bohmian interpretation is true*. It follows from (CS) that the evidence supports neither the Bohmian interpretation, nor its negation – in contradiction to what we have assumed.

Or take a more mundane, non-scientific, example. Say David is slightly color-blind so he cannot always tell blue from green, although usually he gets such things right. Seeing what appears to him to be a blue car David presumably has evidence – albeit inconclusive – for there being a blue car before him. But the evidence is compatible with there being a green car before him. It follows by (EAD) from this latter claim that the evidence sup-

---

<sup>181</sup> For the argument below all that is needed is that there are some cases of underdetermination, for instance, induction to a claim that might even turn out to be false after many confirming instances.

<sup>182</sup> A probabilistic proof would be easy enough, but here I am limiting myself to non probabilistic arguments.

<sup>183</sup> I have been told that the difference between the Copenhagen interpretation and the Bohmian interpretation is not quite as I present it here. Since quantum physics is not the issue, let us pretend that I have presented the issue correctly. In any case, the interpretations are apparently incompatible.

ports: *there is a green car or it is not the case that there is a blue car*. Since the second disjunct is implied by the first, the disjunction is equivalent to: *it is not the case that there is a blue car*. By equivalence, then, the evidence supports the claim that *it is not the case that there is a blue car*. Once again, assuming (CS), the argument shows that David has no evidence that there is a blue car before him.

This may not seem too problematic if you take David's case to be unique in some sense or are willing to disqualify his perception as evidence for there being a blue car in front of him (despite the fact that it clearly raises the probability that this is the case).<sup>184</sup> But the problem generalizes. The above examples both have the same form and follow from the apparently undeniable reality of underdetermining evidence. (UD) entails that given a finite set of evidence propositions  $e$ , this evidence can equally support two incompatible theories,  $T_1$  and  $T_2$ . Thus,  $T_1$  implies  $\text{not-}T_2$ , and  $T_2$  implies  $\text{not-}T_1$ . Let me state this more formally as follows:

$$(1) \quad E(e, T_1) \wedge E(e, T_2) \wedge (T_1 \Rightarrow \neg T_2) \quad [\text{UD}]$$

(EAD) entails the following:

$$(2) \quad E(e, T_1) \Rightarrow E(e, T_1 \vee \neg T_2) \quad [\text{EAD}]$$

which entails:

$$(3) \quad E(e, T_1 \vee \neg T_2) \quad [\text{MP, 1, 2}]$$

Now, since  $T_1$  entails the negation of  $T_2$ , the following *a priori* equivalence holds:<sup>185</sup>

---

<sup>184</sup> Note that this may not be a happy position even with regard to this case. Imagine that David identifies these colors correctly only 50 percent of the time. Now if David believes that there was a blue car in front of him, I take it that that would constitute *some* evidence for there being a blue car there. Imagine that in a court case, a defendant's car is green and David is a witness. Would not his belief that it was blue count as some evidence for the car being blue? Is David's testimony just as relevant to the case as someone who did not see the getaway car? Suppose someone says "yes, it is only evidence for the following: *Either the car was green or it was blue.*" By the Kolmogorov axioms (since the events are mutually exclusive), the probability of this proposition is equal to the sum of the probabilities of both disjuncts ( $\text{Pr}(Bc \vee Gc) = \text{Pr}(Gc) + \text{Pr}(Bc)$ ). Now if you believe that the raising of probabilities is evidence, then (assuming that the probability goes up to 0.9)  $\text{Pr}(Gc) + \text{Pr}(Bc) = 0.9$  and assuming that the prior probability of  $Gc$  was low, e.g.  $\text{Pr}(Gc) = 0.05$  (since it could be any one of the different car colors), then the probability of each of the disjuncts has gone up considerably. It therefore seems that on this construal, it would count as evidence for both disjuncts. If you still find this case unconvincing just focus on the others. My purpose is not to convince anyone on issues related to color perception. As noted, any underdetermined case will do.

<sup>185</sup> The proof is straightforward. First, from left to right: Given that  $\vdash T_1 \rightarrow \neg T_2$  and assuming  $T_1 \vee \neg T_2$ , either  $T_1$  is true, in which case so is  $\neg T_2$  (by the implication), or  $\neg T_2$  is true. So

$$(4) \quad (T_1 \vee \neg T_2) \Leftrightarrow \neg T_2$$

It thus follows from (3) that:

$$(5) \quad E(e, \neg T_2) \quad [(\text{EEQ}), 4, 3]$$

But given the principle of consistency (CS), this entails that  $e$  does not evidentially support  $T_2$ .

$$(6) \quad \neg E(e, T_2) \quad [\text{CS}, 5, 1]$$

and (6) contradicts (1).<sup>186</sup>

Call this argument the Underdetermination Argument (UA for short) for the evidence openness. Notice that in the course of UA there has been no essential appeal to probabilities (although it follows on probabilistic grounds as well). The question now will be which of the principles should be jettisoned.

#### 4.3.1. Equivalence, Consistency and Addition

The equivalence of evidence is involved in one of philosophy's notorious paradoxes, namely Hempel's paradox of confirmation. In Hempel's argument the standard conception of evidential confirmation leads to apparently unreasonable results. Specifically, Hempel showed that coupled with the Nicod principle,<sup>187</sup> the principle of equivalence leads to the conclusion that pink stockings evidentially confirm the claim that all ravens are black. The present argument shows that a black raven equally supports the claim that not all ravens are black and is thus no evidence at all. Given even weak (UD), no matter what conception of confirmation it is coupled with, whether Nicod's or some other conception, the (EEQ) and (EAD) principles lead to paradox. These considerations seem to point to the rejection of (EEQ).<sup>188</sup>

---

$T_1 \vee \neg T_2$  implies  $\neg T_2$ . Now from right to left,  $\neg T_2$  clearly implies the disjunction  $T_1 \vee \neg T_2$ . Hence,  $\vdash (T_1 \vee \neg T_2) \Leftrightarrow \neg T_2$ .

<sup>186</sup> In terms of Modal Logic, what we have here is a failure of Monotonicity+Equivalence. Monotonicity is the condition that:  $\vdash \Box(p) \rightarrow \Box(p \vee q)$  and replacement of Modal Equivalence is  $\vdash \Box(p \leftrightarrow q) \rightarrow (\Box p \leftrightarrow \Box q)$ . Together these two entail closure. Thanks to Johan van Benthem for pointing out the analogy.

<sup>187</sup> In simple terms the principle states that universal generalizations of the form  $\forall(x)(Rx \rightarrow Bx)$  are supported by instances of the form:  $Rx \wedge Bx$ .

<sup>188</sup> Scheffler and Goodman reject equivalence as a reaction to the raven paradox. See, e.g. Scheffler and Goodman (1972: 78), Scheffler (1963: 289), Goodman (1955: 71-2). A more recent attempt is made in Yablo (MS).



It may be suggested at this point that neither (EEQ) nor (EAD) are to be identified as the culprit. It is rather (UD) that is incompatible with (CS). To be sure, Hempel, who did not want to give up on evidence closure realized the challenge that simple cases of evidence underdetermination may pose for his account:

A finite set of measurements concerning the changes of one physical magnitude,  $x$ , associated with those of another,  $y$ , may conform to and thus be said to confirm, several different hypotheses as to the particular mathematical function in terms of which the relationship of  $x$  and  $y$  can be expressed; but such hypotheses are incompatible because to at least one value of  $x$ , they will assign different values for  $y$ .

Hempel (1965: 33)

Unlike Hempel, Carnap was less reluctant to endorse the full consequences of (UD) – the rejection of consistency (1950: 474-6). Thus, it may be suggested, it is (UD) that is incompatible with consistency, not (EEQ).

But the consistency principle that Hempel and Carnap had in mind is significantly stronger than (CS):

(CS\*) If  $e$  evidentially confirms  $p_1$ ,  $e$  does not evidentially confirm a *contradicting* hypothesis  $p_2$  (to an equal degree).

Surely, this principle is not compatible with (UD), and given the pervasiveness of (UD) it must be rejected. (CS), however, is not as disposable as (CS\*). How can a piece of evidence support some hypothesis if it supports its negation? (CS), it seems, is a principle that no plausible theory of confirmation can deny, for otherwise what is left of empirical refutation of a theory? Indeed Hempel and Carnap both embrace the following definition of “disconfirmation”:<sup>189</sup>

(DC)  $e$  disconfirms a proposition (or hypothesis)  $p$  if it confirms  $\neg p$ .

Thus, if  $e$  confirms both  $p$  and its negation, it both confirms and disconfirms  $p$  and is thus evidence neither for  $p$  nor for *not- $p$* . Insisting that (DC) is true and (CS) false would lead to theoretical nihilism with regard to evidence.

The distinction between (CS) and (CS\*) may seem contrived. Nevertheless, a closer look at the substance of (UD) supports such a distinction. Presumably, when a proposition is underdetermined by evidence, there are two incompatible propositions  $p_1$  and  $p_2$  supported by the available evidence  $e$ . A set of measurements, to take Hempel’s example, is entailed by both functions  $f_1$  and  $f_2$  and thus supports both. Yet it does not seem to support the

---

<sup>189</sup> Hempel (1965: 37). Cf. Carnap (1950: 479).

negations of either function. Observations that some animal looks like a rabbit can be taken to support both the hypothesis that it is a rabbit and the hypothesis that it is a hare (by someone who cannot tell a rabbit from a hare), since this observation is entailed by both hypotheses.<sup>190</sup> Yet ‘it is not a rabbit’ and ‘it is not a hare’ do not imply the observation of a rabbit-looking animal (quite the contrary). Thus neither of these hypotheses should gain any support from the observation of a rabbit-looking animal.

So, although (UD) arguably entails that (CS\*) is to be rejected, (CS) is not expendable for any proper theory of the relation between evidence and hypotheses. And yet, as the argument above shows, assuming so much as the weak (UD) and (EAD), the (EEQ) principle conflicts with (CS). Taken together these principles lead to the conclusion that *e* evidentially supports *p* and evidentially supports *not-p*. The implausibility of this conclusion is even more striking in light of the following consideration. It may be suggested that evidential support collects over conjunction:<sup>191</sup>

(ECJ) if *e* evidentially supports *p* and *e* evidentially supports *q*, then *e* evidentially supports *p and q*.

Given this principle, the Underdetermination Argument (p. 104) entails that any underdetermining evidence *e* supports a contradiction. This seems like a conclusion no proper theory of confirmation can endorse.

But even if (ECJ) is rejected (as I think it must be - see below) our problem is still very pressing. Since (UD) appears to be an undeniable reality, and since (CS) must be regarded non-negotiable, it seems that we must give up either (EAD) or (EEQ).

If you are still not convinced, here is an argument to an implausible conclusion that does not assume (CS).<sup>192</sup> For simplicity let's say that we have evidence consisting in a conjunction of two atomic propositions *a* and *b*. Two hypotheses result from operations on the three atomic propositions *a*, *b* and *c* ( $a \wedge b \wedge c$  and  $a \wedge b \wedge \neg c$ ). Both *c* and  $\neg c$  are consistent with *a* and *b*. Now, the evidence raises the probability of  $a \wedge b \wedge c$ , as well the probability of  $a \wedge b \wedge \neg c$  (this can be proven trivially). But we need not use probabilities and the assumption that the raising of probabilities is what the “evidence for”

---

<sup>190</sup> Using probabilities, a proof would be straightforward.  $\Pr(p|e)$  when *p* entails *e* will be greater than the prior probability of *p*. The reason is:  $\Pr(p|e) = \Pr(p \wedge e) / \Pr(e) = \Pr(p) / \Pr(e) > \Pr(p)$  (whenever the prior probability of *e* is less than 1). But here I am restricting myself to non-probabilistic reasoning.

<sup>191</sup> The principle follows from multi premise closure for evidence, but, in fact, evidence does not collect over conjunction. See Carnap's example below. Moreover, the principle leads to paradox - the lottery paradox and the preface paradox.

<sup>192</sup> In conversation Ofra Magidor suggested that one who goes as far as to claim that evidence is non-probabilistic and that (CS\*) is to be rejected, might very well (given the options) opt for rejecting (CS). I see her point but I think that the argument in the main text will deter one from going in such a direction.

relation consists in, to be convinced that  $a \wedge b$  is evidence for both  $a \wedge b \wedge \neg c$  and  $a \wedge b \wedge c$  if everything else is equal. So the following holds:

$$(7) \quad E((a \wedge b), (a \wedge b \wedge c)) \wedge E((a \wedge b), (a \wedge b \wedge \neg c))$$

by (EAD) and the first conjunct of (7), we have,

$$(8) \quad E(a \wedge b, ((a \wedge b \wedge c) \vee \neg(a \wedge b \wedge \neg c)))$$

from (8) we derive by (EEQ):

$$(9) \quad E((a \wedge b), \neg(a \wedge b \wedge \neg c)).$$

Another application of equivalence will give us:

$$(10) \quad E((a \wedge b), (\neg a \vee \neg b \vee c))$$

(10) is already implausible (and adding (CS) or using probabilities would make for incoherence). But things get worse for evidence-closure given the assumption that  $a \wedge b$  is evidence neither for  $c$ , nor for  $\neg c$ . A plausible claim (at least in this context) is that if  $e$  is evidence for  $p$ -or- $q$ , and  $e$  is not evidence for  $q$ , then  $e$  is evidence for  $p$  (at least for atomic independent propositions), that is:

$$(11) \quad [E(e, p \vee q) \wedge \neg E(e, q)] \rightarrow E(e, p).$$

(11), however, would lead us to the impossible situation that  $a \wedge b$  is evidence for either  $a$ -or- $b$  being false (remember that these are atomic propositions):

$$(12) \quad E((a \wedge b), (\neg a \vee \neg b))$$

Equivalence would give us:

$$(13) \quad E((a \wedge b), \neg(a \wedge b))$$

I take it that a proposition cannot be evidence against itself (unless it is paradoxical, self undermining or contradictory perhaps, which presumably atomic independent propositions are not). And so we have an argument with no appeal to (CS) showing that evidence is not closed. Since we must accept (UD), either (EEQ) or (EAD) must be rejected. But if this is the case, then we have shown that evidence is not closed under known entailment. To state things clearly, let us represent evidence closure more explicitly as a principle:

Evidence Closure:<sup>193</sup> For all propositions  $p$  and  $q$  and evidence  $e$ , if  $e$  is evidence for  $p$ , and  $q$  follows *a priori* from  $p$ ,  $e$  is evidence for  $q$ .

Since  $p \vee q$  follows *a priori* from  $p$ , and assuming that we can show all the relevant equivalences hold *a priori*, we can state the logical relations between the different principles as follows:

- (14) Evidence Closure  $\vdash$  (EEQ)
- (15) Evidence Closure  $\vdash$  (EAD)

The following is also true, and plays a pivotal role in responding to Hawthorne's argument against knowledge openness:

- (16) (EEQ), (EAD)  $\vdash$  Evidence Closure

We have seen an argument for the following:

- (17) (EEQ), (EAD), (CS), (UD)  $\vdash \perp$

The last argument made no use of (CS):

- (18) (EEQ), (EAD), (UD)  $\vdash \perp$

That is, that even without (CS) the conjunction of (EEQ) and (EAD) in light of the fact that evidence in at least some cases underdetermines propositions it supports, leads to absurdities. Since (14) and (15) clearly hold, I think we can safely conclude that Evidence Closure is invalid. No matter which of (EAD) or (EEQ) will turn out to be responsible for the contradiction, Evidence Closure must be relinquished. And since no special appeal has been made to probabilities or the lowering of probabilities, the attempt to circumvent the *Open Knowledge Argument* (argument (I-VI) p. 97), is wrongheaded. The only other option is to claim that there are no cases of (UD) in conjunction with a claim that evidence is not probabilistic, and this must be regarded as a very desperate measure indeed.

Essentially, the argument for evidence openness is complete, but it will be important to try to determine whether it is (EEQ) or (EAD) which is to be dispensed with. To this end it might be important to be clear that (EEQ) is involved in the same kind of argument together with another principle. The

---

<sup>193</sup> As with the other principles, I suppress mention, for simplicity, of subjects and other fine-tuning. Nothing will turn on this simplification.

next section contains an argument showing that (UD), (EEQ) and the distribution of evidence principle are incompatible:

$$(19) \quad (\text{EEQ}), (\text{EDIS}), (\text{UD}) \vdash \perp$$

### 4.3.2. Equivalence and Distribution

As in the previous case, let us start with two underdetermined theories  $T_1$  and  $T_2$ :

$$(20) \quad E(e, T_1) \wedge E(e, T_2) \wedge (T_1 \Rightarrow \neg T_2) \quad [\text{assumption}]$$

By (EEQ), we have:

$$(21) \quad E(e, T_1 \wedge \neg T_2) \quad [(\text{EEQ}), (20)]$$

The following would seem to be a highly plausible principle and is an evidential analog of Hawthorne's (DIS) principle for knowledge:<sup>194</sup>

Evidence Distribution (EDIS): If  $e$  is evidence for  $p$  and  $q$ , then  $e$  is evidence for  $p$  and  $e$  is evidence for  $q$ .

(EDIS) and (21) gives us:

$$(22) \quad E(e, \neg T_2). \quad [\text{EDIS}, (21)]$$

Yet having assumed that  $e$  supports  $T_2$ , the result is the same contradiction we had before. So we need to give up either (EDIS), (EEQ), (CS) or (20) (a particular instance of (UD)). The last two I have claimed are virtually undeniable, and I have argued that even without (CS) an absurdity follows, i.e. that the conjunction of any two atomic propositions is evidence that one of them is false (p. 107). It seems clear, then, that either (EEQ) or (EDIS) must be given up. And we have also seen that we need to decide whether to give up (EEQ) or (EAD). So the choice is really between (EEQ) on the one hand, and both (EDIS) and (EAD) on the other. In other words, the Underdetermination Arguments show that equivalence cannot be maintained along with evidence addition or with evidence distribution. The next section is an attempt to say which is the culprit.

---

<sup>194</sup> See p. 30 above.

### 4.3.3. EAD, EDIS and the Logic of Evidence

Which of (EEQ) or (EAD) and (EDIS) is to be rejected? The case against (EEQ) has been mounting. Yet, despite the appeal there is in rejecting (EEQ); despite the fact that rejecting (EEQ) provides a quick way out of the paradox of the ravens, despite the fact that we would be giving up one principle rather than two, I will argue that (EEQ) should be maintained and (EAD) and (EDIS) are to be rejected.<sup>195</sup> Here my reasons for rejecting these latter principles will be partly probabilistic.<sup>196</sup>

If  $T_1$  and  $T_2$  are incompatible theories and have the same initial probability, then this probability must be equal to or less than 0.5 (thus:  $(T_1 \Rightarrow \neg T_2) \rightarrow [\Pr(T_1) + \Pr(T_2) \leq 1]$ ). Let us suppose that each theory has an initially probability of 0.2. The probability of  $\neg T_2$  is therefore 0.8. Now say we receive evidence  $e$  that supports both  $T_1$  and  $T_2$  to an equal degree (for simplicity). Assume that  $\Pr(T_2|e) = 0.4$  and likewise  $\Pr(T_1|e) = 0.4$ . The initial probability of  $(T_1 \vee \neg T_2) = \Pr(T_1) + \Pr(\neg T_2) - \Pr(T_1 \wedge \neg T_2)$ , which equals 0.8.<sup>197</sup> Now if  $e$  supports  $T_2$  and  $T_1$  equally (as we have assumed), then given  $e$  the probability of  $T_1 \vee \neg T_2$  decreases to 0.6 ( $\Pr((T_1 \vee \neg T_2)|e) = 0.6$ ). The reason is simple, since the probability of  $T_2$  rises, the probability of  $\neg T_2$  decreases and since from  $T_1$  it follows that  $\neg T_2$ , the probability of the disjunction  $T_1 \vee \neg T_2$  drops (equivalently  $\Pr(\neg T_2|e) = 0.6$ ). Hence,  $e$  is not evidence for  $T_1 \vee \neg T_2$  even though it is evidence for  $T_1$  (assuming, that is, that if  $e$  lowers the probability of a given theory  $T$ , it does not count as evidence in its favor).<sup>198</sup>

A similar argument holds for (EDIS). If we treat the *evidence for* relation as a conditional probability relation such that it raises it relative to the unconditional probability, the probability of  $T_1$  given  $e$  is the probability of  $e$  and  $T_1$  divided by the probability of  $e$  ( $\Pr(T_1|e) = \Pr(T_1 \wedge e) / \Pr(e)$ ). Assume that the probability of  $T_1$  is 0.2 and that of  $e$  is less than 1 (and greater than 0). If  $T_1$  entails  $e$ , then  $\Pr(T_1|e) > 0.2$ , and likewise for  $T_2$ .<sup>199</sup> This means that the probability of  $\neg T_2$  given  $e$ , is less than 0.8. Now since  $T_1$  entails  $\neg T_2$ , the prior probability of  $T_1 \wedge \neg T_2$  is just the probability of  $T_1$ . And, as you may have figure out already,  $e$  must increase the probability of  $T_1$  to no lesser a

<sup>195</sup> Harman and Sherman (2004) reject (EQ) as a reaction to Hawthorne's arguments (that are presented in Chapter 1.5.1). (EQ) is the epistemic counterpart of (EEQ) and so it may be that they would deny the later principle. Goodman, Scheffler and Yablo reject (EEQ) (see note 189) as a way to resolve the raven paradox. I think that in light of the considerations in the main text, this direction is wrongheaded.

<sup>196</sup> I have already completed my argument on non-probabilistic grounds that evidence is not closed under known a priori entailment since whether (EEQ) or (EAD) and (EDIS) are to be rejected Evidence Closure must be rejected. So there is no apparent reason for me to continue restricting myself to non-probabilistic reasoning.

<sup>197</sup> Since the probability of  $T_2$  was stipulated to be 0.2, the probability of  $\neg T_2$  is 0.8, and the probability of  $T_1 \wedge \neg T_2$  is just the probability of  $T_1$ .

<sup>198</sup> This assumption is not to be confused with the stronger claim that evidence just is increasing the probability of a hypothesis. Here I am merely assuming that evidence cannot lower the probability that the proposition it supports is true.

<sup>199</sup> This is shown by simple application of Bayes' theorem.

degree than it does that of  $T_1 \wedge \neg T_2$ , so all is well for the first equivalence step of the argument that utilizes (EDIS) (assuming that  $e$  supports both  $T_1$  and  $T_2$  to an equal degree). But what about the last step that uses (EDIS)? Clearly, although  $e$  raises the probability of  $T_1 \wedge \neg T_2$ ,  $e$  lowers the probability of  $\neg T_2$  (since it raises the probability of  $T_2$ ). And so, unless we want to claim that  $e$  provides evidential support for a theory, proposition, or hypothesis, even though it lowers its probability, we must give up (EDIS).

This explanation appealed to probabilities, but this is not the only way to determine which of the principles are to be rejected. Let me briefly outline an explanation of what I take to be going on in these cases without employing probabilities. This explanation will serve the open knowledge proponent's arguments below since it will explain (if she is correct), why closure seems compelling in the first place.

The basic idea is that before evidence is received, a disjunctive proposition,  $p$  or  $q$ , can already be well supported (the notion of support need not be construed probabilistically). The evidence, although lending support to one of the disjuncts, can count against the disjunction. Since most objects are not disguised mules, the assumption that some (yet unperceived) object *is a zebra or not a disguised mule* is highly plausible. But the fact that the object looks like a zebra makes it more likely (again – not necessarily in term of probability) that the object is a mule disguised to look like a zebra. It is the neglect of such possibilities that inclines us to accept (EAD) and (EDIS) and may explain the appeal that knowledge closure enjoys, as well.

#### 4.3.4. Carnap's Matrix

The formal considerations of the previous subsections are exemplified in the following scenario devised by Carnap (1950: 382-5).<sup>200</sup> The table below rep-

---

<sup>200</sup> Hempel (1965: 31-33) argues that if evidence is closed under strict implication, every proposition is evidence for any other. I present Carnap's example since, as will become evident, it relates directly to the principles that we have been concerned with. Namely, (EAD) and (EDIS). Hempel's argument proceeds via what he labels the "converse consequence condition." But here is another more general way to proceed: Assuming that a proposition  $e$  is evidence for a proposition  $h$  iff the probability of  $h$  given  $e$  is higher than the probability of  $h$  ( $E(e,h) =_{\text{def}} \Pr(h|e) > \Pr(h)$ ), let me first establish a lemma:

Lemma. For all empirical propositions  $p$  and  $q$ , if  $p$  entails  $q$ , then  $q$  is evidence for  $p$ .

Recall the definition of conditional probability:  $\Pr(p|q) = \Pr(p \wedge q) / \Pr(q)$ . Now let us assume (as is plausible if we are considering empirical matters) that  $0 < \Pr(q) < 1$  and that  $p$  strictly implies  $q$ . It then follows that:

- (1)  $(p \Rightarrow q) \Rightarrow [\Pr(p) = \Pr(p \wedge q)]$  [Since  $p$  and  $p \wedge q$  are equivalent, Kolmogorov axioms]
- (2)  $\Pr(p \wedge q) / \Pr(q) > \Pr(p)$  [1, since  $0 < \Pr(q) < 1$  and  $\Pr(p \wedge q) = \Pr(p)$ ]
- (3)  $\Pr(p|q) > \Pr(p)$  [2, conditional probability]
- (4)  $E(q,p)$  [3, Evidence def.]

Now, the lemma entails:

resents players in a game all of whom have equal chances of winning. The ‘M’s represent male contestants and the ‘F’s denote female contestants.

	Local	Out-of-towner
Junior	F F M	M M
Senior	M M	F F F

Recall the evidence criterion (EC) regarding the relation of evidential support (p. 93 above):

(EC) Necessarily, if  $e$  evidentially supports  $p$ , then the probability of  $p$  given  $e$  is not lower than the prior unconditional probability of  $p$ .  
 $\square(E(e,p) \rightarrow (\Pr(p|e) \geq \Pr(p)))$

Let  $j$  be the proposition that the winner is junior,  $s$  that the winner is a senior,  $l$  that the winner is local and  $o$  that the winner is from out of town. Let  $f$  represent the proposition that the winner is a female contestant and  $m$  that the winner is a male contestant. Now as can be seen in the table above:

$$\Pr(s)=0.5, \Pr(s|f)=0.6$$

$$\Pr(o)=0.5, \Pr(o|f)=0.6$$

So the probabilities of both  $s$  and  $o$  are raised given the information that the winner is female. Now let us look at the probability of the disjunction  $o \vee s$ :

---


$$(5) \quad \forall p \forall q E(p, p \wedge q) \quad [\text{Lemma}]$$

Assuming for *reductio* that evidence is closed under (known) entailment we have:

$$(6) \quad \forall p \forall q \forall r (E(p, q) \wedge (q \Rightarrow r)) \Rightarrow E(p, r)$$

But then since  $q$  follows from  $p \wedge q$ , we have the triviality result:

$$(7) \quad \forall p \forall q E(p, q) \quad [5,6]$$

(7) is surely unacceptable, so one must either reject evidence closure or the proposed definition of evidence (or the Kolmogorov axioms). By relying on the weaker criterion EC rather than on the definition of evidence as the raising of probabilities, we avoid the rejection of the proposed definition of evidence as a reply to the argument against evidence closure. Another version of Hempel’s argument – similar to the one presented here – can be found in Kaplan (1996: 45-56).



Initially,  $\Pr(o \vee s) = 0.7$ . Given  $f$  it becomes 0.6.

Thus, the probability that the winner is *either a senior or from out-of-town* decreases given that the winner is a female contestant. So by (ES), although  $f$  is evidence for  $s$  and evidence for  $o$ ,<sup>201</sup> it is not evidence for  $s$ -or- $o$ . If anything, the fact that the winner is a female is counter-evidence to the claim that the winner is either a foreigner or a senior.<sup>202</sup>

Now to (EDIS). We have the following initial probability assignments:

$\Pr(l) = 0.5$ ;  $\Pr(j) = 0.5$ ; and  $\Pr(j \wedge l) = 0.3$ .

Given the evidence that the winner is a female the probabilities are:

$\Pr(l|f) = 0.4$  ;  $\Pr(j|f) = 0.4$ ; and  $\Pr(j \wedge l|f) = 0.4$ .

Thus while  $f$  raises the probability of the conjunction it lowers the probability of each conjunct. This is a stronger result than the one argued for above (though the assumptions are stronger).<sup>203</sup>

---

<sup>201</sup> This can be denied of course. But in light of the arguments above, this seems to be the wrong thing to say.

<sup>202</sup> Notice that Carnap's argument can be blocked by claiming that although there is a raising of probabilities, that does not mean that  $f$  is evidence for either disjunct. To prove the point in the way Carnap envisaged one would have to appeal to a stronger principle than (EC) (roughly, that raising probability is a *sufficient* condition of evidence). Notice, however, that no such principle is needed for the argument from (UD), nor for the watch case argument. Carnap's point, however, is entirely in line with what I have been arguing for.

Notice also how strange it is that one could have evidence for  $p$  and evidence for  $q$ , yet lack evidence for  $p$  or  $q$ . Asserting as much in ordinary conversation would seem barely intelligible.

<sup>203</sup> The table also provides a counter-example to collection of evidence. Consider the probability distribution for the following proposition:  $l$ ,  $j$  and  $l$ -and- $j$ .  $\Pr(l) = 0.5$ ,  $\Pr(j) = 0.5$ ,  $\Pr(l \wedge j) = 0.3$ . Now consider the information that the winner is a male contestant ( $m$ ). The probabilities given the truth of  $m$  are:  $\Pr(l|m) = 0.6$ ,  $\Pr(j|m) = 0.6$ ,  $\Pr(l \wedge j|m) = 0.2$ . So the probability of each of the conjuncts goes up, yet the probability of the conjunction goes down. If we take any of the contradictory underdetermined cases and a place them as a conjunction, we will get a similar result.

There is an amusing case in which the probability of the conjunction goes down to zero. Say  $n$  people come to a party each wearing a hat. We want to evaluate the probability that each of the people will go home with someone else's hat (we can represent this as a long conjunction: "the first guest to leave takes a hat that does not belong to him and the second guest to leave. ...") Suppose we know the first  $k$  people go home with other people's hats when  $k = n - 3$ . Now focus on the last three people left at the party, assuming the three remaining hats belong to them. If we get the evidence that the third before last took the hat of the second to last, and the second to last took the hat of the third to last, the probability that the hypothesis is true drops to zero. The last guest at the party must go home with her own hat (or with no hat). In some respects this case is similar to the lottery cases, however, it is different in the sense that with every departure of a guests, the probability of the conjunction goes up while the actions of the second to last guest brings it down to zero (from 0.5). This case is given (with some changes) by Rosenkrantz (1981).

The conclusion, I maintain, is that (EEQ) is valid, and that both (EAD) and (EDIS) are not. Granted, denial of two plausible principles is more costly (other things being equal) than denial of one, and although denial of (EEQ) dissolves the raven paradox, we have seen an example - Carnap's example - that shows (EAD) and (EDIS) to be false, and we have a probabilistic and non probabilistic explanation of their failure. Moreover, probabilistically, (EEQ) is easy to prove:

Let us assume that  $p$  and  $q$  are *a priori* equivalent. Then  $p$  iff  $q$  and  $p \vee \neg q$  is a tautology, and thus,  $\Pr(p) + \Pr(\neg q) = 1$ . Hence  $\Pr(p) = 1 - \Pr(\neg q) = 1 - \Pr(\neg p)$ , so the initial probabilities of  $q$  and  $p$  are the same. Now by the same reasoning if  $e$  raises the probability of  $q$  it will likewise do so for  $p$  (the equivalence is *a priori*). Hence, if  $p$  and  $q$  are equivalent, then if  $e$  is evidence for  $p$ , it is evidence for  $q$ . Another way to go is simply to say that if  $q$  follows from  $p$  its probability can be no lower than that of  $p$ , and likewise for  $p$  with regard to  $q$  (since  $p$  follows from  $q$ ). Hence, before and after evidence  $e$  comes to light and is used as evidence, their probability will be the same.

So the conclusion that (EEQ) is to be accepted and (EAD) and (EDIS) rejected, is quite conclusive. It is the failure of evidence to be closed under these, and other modes of valid deductive inference, that (as I will argue here and in the next chapter) explains many of the cases that trouble contemporary epistemologists. Whether or not epistemic closure will survive, is another matter.

## Chapter Summary

Before returning to Hawthorne's arguments against knowledge openness as well as some other arguments, let me quickly take stock of what has been argued for in the preceding sections and add a few comments that will motivate some of the discussion of the next chapter. First, the watch case was presented. I argued that propositions that follow from the knowledge that is gained by looking at one's watch is not supported by the evidence that allows the knowledge of what time it is. This I take to be intuitively right, but the argument did not rely on this intuition.<sup>204</sup> The argument depends on the idea that evidence does not lower probabilities, and supposing that one did not know beforehand that one's watch is showing the correct time, if it reads "3:00", it does not seem possible that one will learn this from looking at one's watch. (I am supposing that one does not have other watch independ-

---

<sup>204</sup> Other cases have the same structure but are less intuitive. For instance, it seems that if looking at one's watch gives one evidence for the time being 3:00, it should also count in favor of the proposition that *It is 3:00 or it is not the case that the time is not-3:00 and it seems like it is*.

ent information that comes to light together with reading the watch, e.g., the sound of three clock-tolls.)

One might complain that my argument was too quick since if one did not know this beforehand one would not know that the time is 3:00 by looking at one's watch. This reply is problematical. It is hard to uphold such a requirement on knowledge unless one supposes that one also knows, for instance, that one's ticket will be a loser before buying a ticket. It is hard to see how someone can know that one's watch is accurate at a specific occasion on mere probabilistic grounds but not that one will lose a lottery (even though the probabilities can be much higher that one will lose a certain lottery). One might be highly justified in believing that one's watch is accurate now, but this will not be changed by looking at one's watch and gaining evidence that, if anything, counts against the proposition that *if the watch reads "3:00", it is showing the correct time*. It seems inevitable, for instance, that in order to know the following proposition one needs evidence: *If your watch reads "3:15," your watch and not my watch that reads "3:00", needs to be corrected*. I doubt someone would claim that such a proposition can be known without evidence, and yet, we have seen that the evidence for knowing that the time is 3:00, need not transfer to all propositions that follow from what is known. So even if one would be tempted to claim that the examples I used are known somehow antecedently without empirical evidence, we know that there are *some* propositions that follow from the known proposition and are not supported by the available evidence and background knowledge. This is because we can prove that if the evidence together with background knowledge does not entail the known proposition, there are such propositions. So unless one would want to claim that all the unsupported propositions are known antecedently, I don't see how this reply could be made to work. In particular I have claimed that looking at your watch and noting that it reads "3:00" gives you no evidence for the proposition that *if your watch reads "3:00," it is showing the correct time*. Although this was a mere example, it shows that as good as your evidence is for the latter proposition, it seems doubtful that you can go from not knowing that it is true to knowing that it is true by looking at your watch. And indeed if we assume that you did not know it antecedently, how could you come to know it by inferring it indirectly from a different proposition that is supported by your new evidence, namely, the new knowledge that the time is 3:00? To put it another way, suppose you do not know that the time is 3:00 by looking at your watch and noting that it shows "3:00". It is clear now, I think, that you have no new evidence for the proposition that *if your watch reads "3:00," it is showing the correct time*. But whether your evidence is good enough for knowledge or not is irrelevant for the case at hand. Unless you antecedently know all propositions that are not supported by your evidence, this suggestion will

only offer a reply to regarding certain cases. It will not offer a more general reply.<sup>205</sup>

Second, on the basis of the probabilistic argument for the openness of evidence, I presented a generalized argument against knowledge closure. The basic idea was that since evidence is not closed, the evidence that allows one to know a proposition will not put one in a position to know the logical consequences of this new knowledge since this new evidence need not give one evidence for those consequences (even in cases when one is aware that they are consequences). I also argued, that for any proposition knowledge of which is based on evidence that does not raise the probability of this proposition to 1, there will be propositions that follow from this proposition that are not supported by the evidence. Since this has been established, I think it is safe to say that if knowledge depends at least sometimes on evidence (and perhaps in certain domains it always does), then for every body of knowledge that is based on non-conclusive evidence, there will be cases that pose a formidable challenge for closed knowledge whether or not we can recognize them as such.

So, third, a straightforward way to respond to the openness of knowledge argument seems to be to either question the idea that knowledge does not have to have probability 1, or to question the idea that evidence is open.<sup>206</sup> The arguments from the logic of evidence are meant to serve as a response to the latter way of trying to respond to the knowledge openness argument.

A response that I will have less to say about, is the claim that, e.g., empirical knowledge does not require evidence. This suggestion seems wrong-headed. It seems that although we do not have a good grip on all the features of knowledge, of evidence, and how they relate, we can feel secure that some domains of knowledge require evidence. But even if knowledge does not always require evidence in a certain domain, there is still no reason to think that the propositions that follow from known ones that are not supported by the available evidence are precisely those that do not require evidence. Absent any reason to think that this is wrong, I think it is pretty safe to say that this type of defense of closure is useless. There are further suggestions of how to respond to the Open Knowledge Argument that I have not yet dealt with. One such is that the central cases, such as the Dreske's zebra case, are one's for which the probability is, in fact, raised. This last claim

---

<sup>205</sup> Roger White (2006) advances an argument on probabilistic grounds against Mooreianism (or Dogmatism) similar to the one advanced here. As I understand his argument it is not directed against knowledge closure and assumes, in contrast to the Underdetermination Argument, that evidence is probabilistic (an assumption that can be denied by a Mooreian). Unfortunately his argument is focused on justification closure, a principle I will later argue is compatible with open knowledge and so it does not directly relate to the issues I discuss here.

<sup>206</sup> Other less obvious ways will also be considered in Chapter 5. Let me just say that I will not be able to consider all replies but merely those that I think might seem plausible. Ultimately the argument will rely on the idea that open knowledge is the simple and elegant solution to the challenge and brings with it many theoretical benefits.

does not target the Open Knowledge Argument as such, but will nevertheless undermine much of its attraction if it turned out that central cases do not lend support to the open knowledge account. I will critically consider one such argument in the next chapter (5.1.1.1.) and will also respond to further ways that might be considered as a proper reply to the Open Knowledge Argument.

These last remarks concern the overall strategy and direction of this and the next chapter. But it is also important to take notice of more specific issues. First, given the high plausibility (or inevitability) of (UD) and (EEQ), (EAD) and (EDIS) must be rejected in order to avoid contradictions. We have also seen independent reasons for rejecting (EAD) and (EDIS) having to do with the logic of evidential support. Now it is easy to see that if evidence is to play an important role in characterizing central features of a logic of knowledge, these conclusions put considerable pressure on the epistemic counterparts of these principles, namely (AD) (roughly, that if one knows that  $p$  one can know that  $p$ -or- $q$ ) and (DIS) (that if one knows that  $p$ -and- $q$  one can know that  $q$ ). By ‘an important role’ I mean roughly, conceptions of knowledge according to which if one lacks evidence for a claim, one does not know it. (I provide a more detailed articulation of possible roles evidence might play in the next chapter.) Second, although denial of (EEQ) may have some advantages even in relation to evidence (i.e. avoiding the ravens paradox), the other principles must in any case be rejected as they have definitive counter-examples and a clear explanation of their failure. Third, proponents of knowledge openness (at least those of the sort I have in mind) would be well motivated to deny the epistemic counterparts of the principles (EAD) and (EDIS) on the same grounds that give rise to the denial of closure. The reason to deny closure and these weaker principles as well is that evidence supporting a known proposition need not carry over through these modes of inference to the propositions inferred. One knows that there is a zebra in the pen even though one does not know that there is no painted mule in the pen, because, while one has evidence for the former proposition, one has no evidence for the latter relative to a fixed body of background information. In fact, as I have claimed earlier, one’s evidence tells *against* the hypothesis that there is no disguised mule in the pen.

A fourth point should be stressed in this context. I have demonstrated that in order to keep fundamental features of evidence one must give up other principles that may at first seem undeniable. For instance, it seems highly implausible that one can have evidence for  $p$ -and- $q$  and yet have no evidence for  $p$  and no evidence for  $q$ . It is conceivable, therefore, that much of the current distaste with closure-denial stems from convictions about evidence that are in any case misguided. I will have more to say about this as well in the next chapter. Finally, we have seen that whether or not the relation of evidential support can be construed in terms of probabilities, there is good reason to reject the principles; (EAD), (EDIS). In both probabilistic terms –

that evidence does not reduce probability (EC), and non-probabilistic terms – that evidence cannot support both a proposition and its negation (CS), I have argued that these principles are invalid (and even this last principle has been shown not to be essential). As we will see, there is reason to suspect that the appeal of Hawthorne’s arguments against epistemic openness turns on the same feature of knowledge that, as we have seen in the case of evidence, give rise to principles that might seem plausible, yet must ultimately be rejected. To give substance to this claim, I now return in the next chapter to first consider Hawthorne’s argument and then consider other challenges to the Open Knowledge Argument.

# Chapter 5: Open Knowledge – Costs and Benefits

## 5.1. From Evidence to Knowledge

To defend closure against examples advertised by its deniers, John Hawthorne (2004a, 2005) argues (as presented in Chapter 1.5.1) that interpreted in the way closure deniers would have us interpret them, these examples conflict with other, more basic and weaker epistemic principles. The advocate of knowledge openness, he claims, is forced to reject these highly compelling principles. In other words, to deny closure on the basis of these examples is tantamount to denying a number of weaker principles as well. Thus, if his arguments are cogent, Hawthorne manages to significantly raise the cost of knowledge openness.<sup>207</sup>

---

<sup>207</sup> Hawthorne's argument is, of course, not the only argument against knowledge openness. The argument that open knowledge has unintuitive consequences is perhaps the most prominent one employed among contemporary epistemologists. I do not think it to be a very convincing argument, however.

On any open knowledge account, conjunctions of the form: 'he knows it's a zebra, but he does not know it's not a disguised mule' which DeRose calls "abominable conjunctions," - come out true (1995). Nevertheless, not all closure advocates agree that these conjunctions are as bad as DeRose makes them out to be. There is much to be said about these conjunctions and surely I cannot address all aspects of this vexed issue here. I shall note just a few points oriented towards explicating the relation between this issue and evidence openness, and explain why I do not take these consequences as seriously as many epistemologist have.

First, it is not clear that proponents of epistemic openness should even attempt to alleviate the unintuitive consequences of non-closure. The reason is simple. Presumably, they have at least one eye focused on providing a reply to skepticism. Regardless of whether their reply is successful or not, it should include an explanation of the force of skepticism – why the skeptic's challenge is so gripping. The oddity of abominable conjunctions can be one source of this force. If their abominability is alleviated, this avenue for explaining the force of skepticism is no longer available to the proponents of epistemic openness. There is something to be said, therefore, for just accepting a surface unintuitiveness involved in epistemic openness, particularly if – as the account of knowledge openness aims to do – it is explained by the underlying evidential structure.

Second, to say that the open knowledge account is committed to abominable conjunctions is not quite to say that it is committed to their assertibility or believability. The fact that they come out true according to the theory does not mean that there are no independent constraints on belief and assertion that prevent them from fitting either of these categories. Pragmatic considerations, for one, may explain this. Thus, for example, if beliefs are rationally closed under logical inference, then one is not necessarily committed to first-person abominable conjunctions such as "I know that these are my hands, yet I do not know that I am not a handless brain in a vat." Although I don't know that I am not a brain in a vat, since I believe it,

After looking closely at principles and probabilistic properties of the *evidence for* relation in Chapter 4, we are now in a position to revisit Hawthorne's arguments. The purpose of the section 4.3 was, among other things, to see whether those arguments raise the cost of the open knowledge account that is motivated by the logic of the *evidence for* relation. I argued in Chapter 1 that these arguments are in fact successful against Dretske's and Nozick's open knowledge accounts. Yet the way I have presented the argument for open evidence makes it manifest, I hope, that they do not impede on open knowledge that relies on open evidence.

Hawthorne assumes in the course of his arguments that the following principles would be accepted, at least initially, by open knowledge advocates:

---

there is a perfectly good pragmatic explanation for why I will not ordinarily assert the conjunction, and why, if I were to assert it, it would appear odd. In general, all the proponent of epistemic openness needs is a way of explaining why forming the relevant beliefs is blocked. And this leads to the final comment.

Most importantly, third, I have argued that any plausible account of evidence is committed to the truth of what we may call abominable evidence constructions. For instance, that S may have an item of evidence for  $p$  and for  $q$  but not for  $p$ -and- $q$ , or have evidence for  $p$  and evidence for  $q$ , but not evidence for  $p$ -or- $q$ , or evidence for  $p$ -and- $q$  but no evidence for either  $p$  or evidence for  $q$ . So defenders of closure are in the same boat as their opponents when it comes to evidence. They too are committed to the possible truth of the unintuitive conjunction "I have evidence that the winner will be a local junior player but not that it will be a local player." Moreover, it is the closure advocate that is committed to the truth of claim such as "she knows there's no colored water in the bottle although her evidence tells against it." Hence it is far from clear that the closure of knowledge liberates from costs at the level of intuition. Williamson, as I will argue below, is committed to justification openness, and most closure advocates believe that there are constraints on closure such that it is possible to know that one's car is in the driveway but not that it has not been stolen (e.g. since one has not formed a belief on the basis of proper inference). So conjunctions of the abominable type are a problem for more or less everyone. Moreover, the open knowledge advocate relies on evidence openness to be unintuitive, if where intuitive, one could use evidence to define closed context sets. Thus since evidence is not intuitively open, salience of propositions that are known on the basis of evidence which is in fact open will provide contextual predictions of knowledge where evidence is absent.

Still, you might think, better to have unintuitive consequences in one realm (i.e. evidence) than in two (knowledge and evidence). But if the open knowledge account is cogent these unintuitive results are not really distinct. The truth of unintuitive claims about knowledge stems directly from the truth of the analogous claims about evidence to which all are committed. But more importantly, given these undeniable features of evidence, the open knowledge account is better equipped to provide an *explanation* of the unintuitive conjunctions regarding knowledge. Knowledge has unintuitive consequences because evidence (on which knowledge is based) has unintuitive logic. Moreover, since justification can be closed on an open knowledge account (see below), at least one type of unintuitive conjunction might be avoidable. Regarding the role of intuitions in philosophical methodology my view is that cases should have primacy, not principles. Since principles are valid only if they hold in all cases, I am suspicious of arguments that proceed solely on their surface intuitiveness. This is not to say that we always have a choice. We might have to start from such premises but I think we ought to make the best efforts in making sure that there are no further theoretical considerations at our disposal.



- (AD) Necessarily, if S knows that  $p$  and infers  $p \vee q$  from  $p$ , then S knows that  $p \vee q$ .
- (DIS) Necessarily, if S knows that  $p \wedge q$  and infers  $q$  from  $p \wedge q$ , then S knows that  $q$ .
- (EQ) Necessarily, if S knows that  $p$  and knows that  $q$  is *a priori* equivalent to  $p$ , then S knows that  $q$ .

We saw that the pairs (AD)-(EQ) and (DIS)-(EQ) give the same results as (CP) contrary to what both Dreske and Nozick seem to have realized (though Nozick anticipates the equivalence of (CP) and the combination (DIS)-(EQ)). Thus, e.g. Nozick argued inconsistently that (AD) and (EQ) are valid principles while denying that (CP) is. Thus, Hawthorne shows that these open knowledge proponents need to reevaluate their accounts and concede that what they took as valid principles in their theories must be denied.

But this is not the case with open knowledge that is motivated by open evidence. We saw that the evidential counterparts of (AD) and (DIS), that is, (EAD) and (EDIS), are not valid principles, thus one who accepts that knowledge is open because evidence is, will deny the principles that are necessary for Hawthorne's arguments. So, for instance, an open knowledge proponent of this variety would be well advised to admit that if one knows that there is a zebra in the pen, one can know by competent inference utilizing (EQ) that *there is a zebra in the pen and that there is no disguised mule in the pen*.<sup>208</sup> But in some cases, one will not thereby be in a position to come to know by proper inference that there is no disguised mule in the pen. So although Hawthorne's arguments are valid, what I am recommending for the open knowledge proponent is that she claim that they are not sound. Although (EQ) is a valid principle, since (EAD) and (EDIS) are not, (AD) and (DIS) are not valid from an open knowledge perspective and so the costs of open knowledge are not expected to rise by Hawthorne's arguments.

I have proposed, then, in light of Hawthorne's argument against open knowledge a unified rationale for the rejection of epistemic closure, addition and distribution, namely, that evidence is not closed under these operations. Still it is important to note why this account does not provide a defense for Dreske's and Nozick's knowledge openness. Roughly on their accounts knowledge is not closed due to its underlying subjunctive structure. On the account proposed here, in contrast, it is not the subjunctive structure of knowledge to which its openness is indebted, but rather the evidential nature of knowledge and the logic of evidential relations.

---

<sup>208</sup> Harman and Sherman (2004) respond to Hawthorne's argument by denying that this is right. That is, they deny (EQ). I think that this is the wrong response but I cannot go into this issue here. Though my own reasons are slightly different, I think Hawthorne (2004b) adequately addresses their concerns.

Having a response to Hawthorne's arguments does not yet justify thinking that knowledge is open. That justification was supplied in the previous chapter by the Open Knowledge Argument. It is time to examine closely its premises to see whether they can be questioned. The current section is aimed at looking more closely at the link between evidence and knowledge that the Open Knowledge Argument (p. 97) appeals to. After making clear what is needed for the Open Knowledge Argument as far as the evidence-knowledge link is concerned, I will employ an objection to correct a lacuna in this argument (5.1.1). I will then look briefly at a suggestion that knowledge is closed even though one lacks evidence for the consequences of knowledge since proper inference itself provides the added required justification (5.2.). I then consider the suggestion that gaining knowledge of a proposition changes the evidential situation since all knowledge can count as new evidence. I conclude that this suggestion on one variant (Jeffrey Conditionalization) does not challenge the Open Knowledge Argument (5.3.1) while using knowledge as evidence by Standard Conditionalization does (5.3.2). I then offer general reasons for why this second suggestion is problematic, and then more specific arguments criticizing Williamsonian epistemology that takes this route in defending (multi-premise) closure (5.3.3). I conclude section (5.3) with some benefits the open knowledge account has and then go on to defend justification closure within this framework ((5.3.4) and (5.4.2), respectively). To make this route somewhat clearer, let me get back to the main features of the Open Knowledge Argument.

The probabilistic Open Knowledge Argument is as follows:

- I. *Fallibilism* - One can know that  $p$  even though one's total evidence does not guarantee that  $p$ . That is, one can know that  $p$  even though the probability of  $p$  given one's total evidence is less than 1. [assumption]
- II. For any fallibly known proposition  $p$  (in the sense of (I)) there is a proposition  $q$  that *a priori* follows from  $p$  such that the conditional probability of  $q$  given all the available evidence is lower than the unconditional probability of  $q$ . [independent proof]
- III. (EC) - If  $e$  is evidence for  $q$  then the conditional probability of  $q$  given  $e$  is not lower than the unconditional probability of  $q$ . [assumption]
- IV. For any fallibly known proposition  $p$ , there are propositions that *a priori* follow from it which the available evidence does not support (the probability of the proposition is lowered). [I, II, III]
- V. (KE) - For all subjects  $S$  and propositions  $p$ , if  $S$  knows that  $p$ , then  $S$  has evidence for  $p$ . [assumption]
- VI. So, *Knowledge is open* - For all subjects  $S$  and propositions  $p$  and  $q$ , if a proposition  $p$  is fallibly known by  $S$  to be true, there is a proposition  $q$  such that even if  $S$  knows that it follows from  $p$ ,  $S$

is in no position to know  $q$  by competent inference (keeping the evidence fixed).

We have seen a proof of (II), and (I) will be considered more closely later in this chapter. Regarding (III), the basic idea behind it was to argue that the following principles are incompatible:

(EC) For all evidence  $e$  and propositions  $p$ , if  $e$  is evidence for  $p$ , then  $e$  does not lower the probability that  $p$  is true.

and

*Closure of Evidence* (CE): For all subjects  $S$ , evidence  $e$  and propositions  $p$  and  $q$ , if

- (i)  $S$  has evidence  $e$ ,
- (ii)  $S$  knows that  $S$  has  $e$ ,
- (iii)  $S$  knows  $e$  evidentially supports  $p$ ,
- (iv)  $S$  knows that  $p$  (logically or conceptually) entails  $q$ ,

then,  $e$  evidentially supports  $q$  for  $S$ .

It is of course possible to retain (CE) at the expense of (EC). But first, this must be regarded as a significant cost. It is hard to imagine a theory that captures a workable notion of evidence while violating (EC). Second, the examples I have been considering all invoke a strong intuition that, regardless of (EC), there is reason to doubt (CE). Although one does have evidence that the time is three o'clock (i.e. the watch showing "3:00"), one does not have evidence for the truth of: *even if the watch has stopped, it is showing the correct time*, or that *if the watch shows "3:00", it is showing the correct time*. While one's memory of having parked the car is evidence that the car is in the driveway, one does not have evidence that the car has not been stolen. And likewise for many other cases. (Yet, other cases do not arouse the same confidence.) Finally, third, to be in a position to claim that (CE) is valid one would need to argue either that there are no cases of underdetermination as I have described them something which follows from premise (I) (reminder: all that is needed is inductive underdetermination perhaps even to a false conclusion), or, one would have to find some problem with the (UA) for open evidence (p. 104).

Moreover, we have now independent arguments and an explanation of what the mistake is in endorsing (CE). This explanation was given both in probabilistic and non-probabilistic terms. The following is another way of explaining what is transpiring in these cases in terms of possible worlds. Having evidence supporting a proposition  $p$  may be explicated as having

reason to believe that the actual world is one of the  $p$ -worlds. Or if you prefer, that evidence for  $p$  raises the probability that the actual world is one of the worlds in which  $p$  is the case. Now if  $q$  (*a priori*) follows from  $p$ , then any world that is a  $p$ -world is also a  $q$ -world. So you might think that evidence that the actual world is a  $p$ -world must also be evidence that the actual world is a  $q$ -world. But, intuitive as it may be, this last step is incorrect. If  $p$  implies  $q$ , then surely, a world that is a  $p$ -world is also a  $q$ -world (and this might also explain why knowledge closure seems so intuitive). However, whether the evidence supporting the claim that a world is a  $p$ -world also supports the claim that it is a  $q$ -world depends on the relation between the purported evidence and  $q$ . Specifically, it depends on whether the evidence raises or lowers the probability that the world is a  $q$ -world (or in other words, whether it counts in favor or against the world being a  $q$ -world). Now, although the probability that the world is a  $q$ -world cannot be lower than the probability that it is a  $p$ -world, if the initial probability that the world is a  $q$ -world is higher than the posterior probability of  $p$ , evidence that it is a  $p$ -world might lower it. This is why, as we have seen, a proposition that increases the probability of  $p$  can lower the probability of a proposition  $q$  implied by  $p$ .<sup>209</sup> This suggests that if, as urged, the idea that evidence must not lower the probability of the proposition which it supports is to be preserved, (CE) must be renounced. Evidence is not closed under known entailment.

Now since items of knowledge that are underdetermined (or fallibly known) entail propositions for which there will be no evidence, or alternatively, since any item of knowledge of a disputed domain<sup>210</sup> needs evidence, epistemic openness is a simple and elegant upshot of the openness of evidence.

But as far as the argument goes, premise (V) (or (KE)) regardless of its plausibility, has not been supported by argument. The next sections fill this lacuna.

### 5.1.1. Premise (V) of the Open Knowledge Argument

It will be shown in this section that in a certain sense, (KE) might not serve the purpose of the Open Knowledge Argument, and yet a more careful and modest knowledge-evidence link, will.

It may be that not all knowledge is based on evidence. A person may know that she intends to  $\phi$ , without having any evidence (other than her very

---

<sup>209</sup> The same explanation with minor modification can account for cases involving disjunction and conjunction. The explanation cannot account for cases involving equivalence, which is one reason to think this is the correct response to the cases above.

<sup>210</sup> A disputed domain of knowledge in the present context is just one for which there is agreement between the open knowledge proponent and the closure proponent that it is a domain for which there is no requirement that knowledge is based on conclusive evidence.

intention, perhaps).<sup>211</sup> Perhaps one can know some things *a priori* having no empirical evidence at all, one might know that one does not have one's hand in a fist behind one's back without looking at it or feeling it, or one might have basic perceptual knowledge that is non-evidential. I don't pretend to have any new considerations about this complex issue. I want to focus attention on propositions knowledge of which uncontroversially depends on evidence. In what follows, then, (unless I indicate otherwise explicitly) my focus will be on the type of knowledge that plausibly depends on evidence.

Chapter 4 opened with such a case. One cannot know, it seems, that a watch, even if it has stopped, is showing the correct time unless one has independent evidence to that effect. In such cases, I have argued, it is sometimes the case that one has evidence allowing knowledge of a proposition *p*, which nonetheless (and taken as a whole) *counts against* what one knows to follow from *p*. Thus evidential closure conflicts even with the rather weak idea that evidence cannot decrease the probability that the proposition it supports is true. Given this modest assumption, the relation of evidential support is not closed under known implication and the argument from (UD) proceeds with an even weaker non-probabilistic assumption, namely, on (CS) (and even (CS) is not required in the last analysis - see p. 107).

But the open knowledge argument leaves the nature of the dependence of knowledge on evidence unspecified. As I hope to show, the open knowledge advocate need not appeal to an especially strong form of dependence. The dependence I will appeal to on behalf of the open knowledge proponent is modest in both scope and strength. Its scope does not range over all justified belief, but rather encompasses only knowledge in a domain of dispute.<sup>212</sup> In fact, I want to remain more or less neutral on the question of *what* kind of knowledge depends on evidence. I do maintain, however, that at least some knowledge both depends on evidence and is underdetermined by it (weakly underdetermined). Also, the *strength* of the dependence claim here is minimal. Strong dependence claims that enough evidence is both necessary and sufficient for believed true propositions in a domain of dispute to count as knowledge. A milder form of dependence is the claim that evidence is only a necessary condition for knowledge in such a domain. While I find this posi-

---

<sup>211</sup> If evidence cannot support itself, it seems that knowledge one has of one's intention is knowledge without evidence. That claim that evidence does not support itself has been challenged by Williamson (2000). Other examples suggest themselves that do not depend on this assumption. It seems that I know that my hand behind my back is not in a fist without evidence. Perhaps here the evidence is that I lack a certain tension in my muscles. The issue surely deserves more inquiry.

<sup>212</sup> Contrast Conee and Feldman: "evidentialism is a supervenience thesis according to which facts about whether or not a person is *justified in believing* a proposition supervene on facts about the evidence that the person has" (2004, 1, emphasis added). The difference in scope is substantive for, as will become clearer below, I maintain a distinction between justified belief and knowledge-promoting justification.

tion agreeable, the argument for epistemic openness does not even require this much. A more modest dependence constraint is the following:

*Negative Evidence Dependence (NED)*: For a domain of dispute  $D$ , if an agent  $S$  does not know  $p$  at time  $t_o$ , between which time and some later time  $t_1$  the only change in  $S$ 's evidential state is the addition of information  $i$  counting as a whole (together with one's former evidence) against  $p$ , then  $S$  does not know  $p$  at  $t_1$ .<sup>213</sup>

So long as we accept this rather weak constraint, there are good grounds (in light of the Open Knowledge Argument) for denying epistemic closure. Since evidence supporting one proposition can count against another proposition entailed by it, and since knowledge is prevented when the change in evidential status counts against the latter proposition, there is reason to conclude that knowledge is not closed under known entailment.

An example might help. Take the proposition that *your car has not been stolen from 87<sup>th</sup> and Broadway*. Having parked your car a while ago (you don't know where in Manhattan you parked the car), presumably you don't know whether this proposition is true or false. At some time, you suddenly remember parking the car there. Since the probability that it was stolen from 87<sup>th</sup> and Broadway goes up, if you did not know that your car has not been

---

<sup>213</sup> Further qualifications are needed, since one might correct one's reasoning, find a mistake, etc, between  $t_o$  and  $t_1$  and thereby come to know that  $p$ , without any new evidence. So here one might add that other things are equal, or that  $S$  has not reconsidered during this time anything with regard to  $p$ , simply treat the correction as new evidence, and so forth. Also, there are cases where the antecedent of the claim is satisfied, such as when one does not believe the proposition yet the additional evidence that counts against the proposition cause  $S$  to believe it (perhaps for the first time). I hope that the spirit of (NED) is clear enough so that I can use it without adding these (and perhaps other) qualifications. Thanks here to Professor Pagin.

There is another kind of concern that I am ignoring. Suppose part of the evidence that is gained counts in favor and part of it against  $p$ . It might be the case that even though as a whole this new evidence counts against  $p$ , it still allows for knowledge that  $p$ . I don't know myself if a plausible example of this kind can be constructed, but I see no reason why this would be impossible. Nevertheless, since any underdetermined knowledge will have propositions that follow from the known proposition but are not supported by the evidence only one case where (NED) holds is needed. And, since the examples are clearly not of the type that would cause trouble for (NED), I will not explore this possibility further.

To exclude cases where some of the evidence counts against  $p$  and some for it as cases posing an obstacle to knowledge, I interpret "additional information" as counting against  $p$  only if the information gained as a whole counts against  $p$ . When we view the evidence as a whole, I will assume that differences in the order of the receiving of information will not matter. To this end it might be required to instate a threshold or to be able to appeal to times before  $t_0$  for comparison. Again, since only one case of (NED) is needed, I will not spend time refining it. Here too, thanks to Professor Pagin.

In light of these worries, then, I am not advancing (NED) as a principle but merely as a claim about how knowledge normally relates to evidence. To advance it as a principle one would have to deal not only with the worries just outlined, but also with cases where one moves from a Gettier case to a non-Gettier situation, or cases where fake barns are removed from the vicinity. Thanks here to Professor Hawthorne and Maria Lasonen-Aarnio.

stolen from 87<sup>th</sup> and Broadway before remembering where you parked your car, surely you do not know it now that you remember having parked it there. If anything, your memory counts against the truth of this proposition (i.e. it makes the possibility that the car was stolen from this location more likely). Nonetheless assuming memory can sometimes facilitate knowledge of the location of your car and that there are no defeaters in this case, you know that your car is on 87<sup>th</sup> and Broadway. Since you know that that the car is on 87<sup>th</sup> and Broadway entails that it was not stolen, closure entails - contra to what we have just seen - that you know this as well. So if (NED) is not absolutely misguided, i.e. if it does not fail in *all cases*<sup>214</sup> of underdetermined knowledge, and if you do not know a priori all the probability lowered propositions that follow from your knowledge, we have found good reason to question the validity of closure.

#### 5.1.1.1. An Objection to the Use of (V)

Quite surprisingly, perhaps, premise (V) – the idea (roughly) that by necessity if one knows, one has evidence – is problematical if used in an argument concluding with open knowledge (IV). The best way to appreciate the problem is to consider an objection to (V), one that I will claim can be met by using (NED) in its place. The problem with premise (V) is that it will not, in the last analysis, serve the open knowledge proponent’s purpose.

To see the objection consider the probabilistic argument for knowledge openness. It may be objected that even if all I have said is true and the evidence in these cases lowers the probability of the derived propositions, in cases of knowledge other evidence always creeps in which raises it. E.g. the absence of signs indicating that the animal in the pen is a disguised mule raises the probability that it is not a disguised mule, given that it is more likely than not that if there were a disguised mule in the pen, the zoo keepers would tip visitors off using sign posts.<sup>215</sup>

The objection more specifically is as follows. From the axioms of probability it follows that  $\Pr(p|e) > \Pr(p)$  iff  $\Pr(\neg p|\neg e) > \Pr(\neg p)$ . Now, since this is the case, the absence of evidence that does not come to light, counts against any given proposition if in fact this evidence has not come to light (assuming, perhaps, that a subject can reflect and realize that such evidence has not materialized). But since any fallibly known proposition has some evidence that would make the proposition more likely, there is always evidence

---

<sup>214</sup> I say “all cases,” since we have seen that for any underdetermined proposition, there are propositions that follow from it that are not supported by the evidence taken as a whole (p. 98). In other words, we need not treat (NED) as a principle or necessary requirement on the evidence knowledge relation. Suffice it that we treat (NED) as a claim about how knowledge regularly relates to evidence. See previous note.

<sup>215</sup> This objection was raised by anonymous referee to a similar argument to the one proposed here in a coauthored manuscript (Sharon and Spectre, unpublished manuscript).

(which is constituted by this evidence not materializing) against any proposition that is fallibly known. No wonder, then, that for any fallibly known proposition  $p$  there are propositions that follow from it that the available evidence counts against. There is evidence, after all, that counts against the proposition  $p$  itself. Yet this in and of itself is no cause for worry, since presumably, whenever we know a proposition fallibly, the overall evidence will count in its favor. We can therefore rest assured that although some evidence will always count against a fallibly known proposition, this evidence will be overrun by the evidence for it, taken as a whole. What matters, then, with regard to knowledge is what a body of evidence counts for as a whole.

But now a new thought might be of concern to the open knowledge proponent. In the zebra case, for example, there will always be evidence (the type of evidence that materializes in the form of lack of counter-evidence to a fallibly known proposition) that counts in favor of there not being a disguised mule in the pen. For instance, as noted above, the fact that the zookeepers did not put up a sign warning visitors of disguised mules, can count (together with background knowledge that zookeepers are usually reliable) in favor of the animal not being a disguised mule. But looking back at (KE) it is now apparent, that at least in some cases (perhaps even in all cases), for a derived proposition  $q$  from a known proposition  $p$ , there will be *some* evidence counting in favor of  $q$ . Hence the conclusion that knowledge is open would not follow from the premises since (KE) might be true, yet one might always have *some* evidence for any proposition inferred from a known proposition. Contra-posed, (KE) is the claim that if one has no evidence, one does not know, but here, and perhaps in all cases, one does have some evidence in favor of the derived proposition.

The open knowledge proponent needs to do more work, then, in showing that properly inferred propositions from fallibly known propositions could be unknowable relative to a fixed body of evidence. She can either try to show that not all derived propositions are of the kind that have some evidence in their favor, or alternatively she can appeal to a different principle. So I think the open knowledge advocate should accept that as things stand, the Open Knowledge Argument is not sound. Yet a closer look at the substance of (KE) will reveal that there is firmer ground on which to place the argument.

This realization is what motivates the use of (NED) in place of (KE) (the second way in which an open knowledge proponent may respond to this worry). In countering this objection in light of (NED) what has to be done is to show that the evidence gained *as a whole* counts, in the relevant cases, against the proposition that is entailed by the known proposition. So in the zebra case, if we call the absence of warning signs etc  $\neg e^*$ , what the open knowledge proponent needs to show is that  $e$  counts in favor of  $DM$  *more than*  $\neg e^*$  counts against it. Specifically, the evidence as a whole, including the visual appearance of a zebra looking animal needs to lower the probability of  $\neg DM$  more than  $\neg e^*$  raises it.



I will soon return to just such an argument, but first allow me to note that the objection hardly seems applicable to all the examples we have seen (this was the first way an open knowledge advocate may respond to the objection). Consider, for instance, of the watch<sup>216</sup> or the stolen car cases. Although with relation to the watch case the point was formal, intuitively, it is hard to think of convincing cases where evidence can play the same role as  $\neg e^*$  does in the zebra case. Similarly with regard to the car case.

Moreover, there is always a way to formulate the deduced proposition so as to either eliminate the possibility of such evidence as  $\neg e^*$  (“the animal in the pen was deviously swapped for a disguised mule to intentionally deceive visitors”) or to include one’s total evidence and generate the same argument I presented above with relation to the watch case (p. 93). The idea is that when the known proposition is  $p$  and the added (or total) evidence that allows one to know it is  $e$ , the directive is to infer  $\neg(e \wedge \neg p)$  from  $p$ . If one does not have evidence for this claim beforehand, one will not have it after receiving  $e$ . From the fallibilism assumption (I), the subject will count as knowing  $p$  on the basis of  $e$  even though the conditional probability of  $p$  given  $e$  is less than 1. One way to understand this assumption is that one does not know before  $e$  becomes one’s total evidence, that  $e \wedge \neg p$  is false. But the probability of the falsity of  $e \wedge \neg p$  only decreases when  $e$  is one’s (new) total evidence, hence, one has no evidence for it. More generally, any proposition that both makes the total evidence more probable (even if it does not strictly entail it) and is incompatible with the known proposition, will have the effect of gaining support from the evidence that allows knowledge of  $p$  such that its negation can be inferred. These propositions will have the features that the objection neglects.

To make things yet more explicit I will argue that even if I stick with the original inferred proposition (that *the animal in the pen is not a disguised mule*), the kind of evidence such as  $\neg e^*$  does not pose a problem for the Open Knowledge Argument. To pose a problem it is not enough that whenever it is known that an animal is a zebra there will be some evidence that raises the probability that it is not a disguised mule. As noted,  $\neg e^*$  must raise the probability that the animal is not a disguised mule more than the fact that it looks like a zebra lowers it. If this is not the case, then (NED) is applicable and the Open Knowledge Argument is reinstated.

Let us consider what the objection requires. Take a concrete case. In visiting a zoo one can see animals of various types. When looking at an animal from a far, it could be any one of the 1,000,000 animals populating zoos (call them  $o$ ). Suppose there are 10,000 zebra-looking animals ( $=z$ ), out of which

---

<sup>216</sup> In the watch case we have a proof that generalizes as follows: For all propositions  $p$  and all evidence  $e$ ,  $\Pr(e \rightarrow p | e) < \Pr(e \rightarrow p)$ . [assuming that  $\Pr(\neg p \wedge e) > 0$ ]. Since  $(e \rightarrow p)$  follows from  $p$ , we know that as a whole the evidence  $e$  (which will include that fact that  $\neg e^*$  has not materialized). Hence, as a whole, the evidence will count against previously unknown propositions that follow from one’s knowledge.

10 are **disguised mules** ( $=dm$ ). Nine of the disguised **mules** have some indication of the **disguise** ( $=idm$ ) that would tip off zoo visitors that there is a disguised mule present, and one of them is an unmarked perfectly **disguised mule** ( $=dm^*$ ). The rest of the zebra looking animals are **zebras** ( $=z$ ). Upon having the visual input caused in the right way by having a zebra present, one will presumably know that it is a zebra. Now let us see whether the fact that there is no sign indicating a disguised mule ( $=\neg e^*$ ) raises the probability that the animal is not a disguised mule more than the **appearance** of a zebra ( $=a$ ) lowers it. Assuming that it does, the following must be true (I am treating the predicates, e.g.  $z$ ,  $dm$ , as propositions in the formulation such that ' $dm$ ' is short for 'that is a  $dm$ '):

$$\Pr(dm|a \wedge \neg e^*) < \Pr(dm)$$

that is, that the appearance of a zebra-looking animal *and* the absence of indications that there is a disguised mule in the pen, the conditional probability, is lower than the unconditional probability of the presence of a disguised mule. However, the left side equals (slightly more than)  $10^{-4}$ . The right side equals:  $10/10^6=10^{-5}$ . Starting with plausible numerical assignments, then, no problem arises for the original analysis of the example. Seeing a zebra looking animal in the zoo raises the probability that the animal is a disguised mule more than it lowers it even though the zookeepers are more likely than not to put signs up that warn visitors of the disguise (9 out of 10 times, in this example).

But the above is not specific to this assignment. In fact for the conditions to be as the objection would require, it must be the case that the frequency of disguised mules is greater than the total number of objects divided by the number of zebra-looking animals. This can be argued for in the following way: The objection requires that the probability that the animal is a disguised mule given the evidence be greater than its probability before the evidence is in. The initial probability that something is a disguised mule is just the number of disguised mules divided by the number of animals ( $\underline{DM}/\underline{O}$ ).<sup>217</sup> The probability given the evidence is the delimited  $\underline{DM}^*/(\underline{ZL}-(\underline{DM}-\underline{DM}^*))$ . Since in normal cases the total number of animals ( $\underline{O}$ ) is much greater than that of the zebra-looking animals, for  $\underline{DM}^*/(\underline{ZL}-(\underline{DM}-\underline{DM}^*))$  to be less than  $\underline{DM}/\underline{O}$ , the number of disguised mules times the number of zebra looking animals must be very large as well (given that  $\underline{DM}^*=1$ , roughly,  $\underline{ZL} \times \underline{DM} > \underline{O}$ ).<sup>218</sup>

<sup>217</sup> Underlined capital letters stand for the number of objects the propositions are about.

<sup>218</sup> Even though this is very implausible, it is even less plausible if we drop two assumptions. First, that we restrict ourselves to animals and not objects more generally. Second, we have assumed that the number of well-disguised mules ( $\underline{DM}^*$ ) is 1. If the number is greater than one, circumstances must be even more extreme for the objection to apply. To see this, consider a more general version of the objection. The objection depends on the following claim

Such circumstances are not impossible, but are certainly quite unusual and they do not look like the kind of conditions necessarily imposed on knowledge. In normal cases, the presence of peripheral information of the sort invoked in the objection can raise the probability of the entailed proposition, but usually this will be by a minute degree, a degree that will not overcome the decrease in probability brought about by the whole of the evidence supporting the known proposition. Moreover, the open knowledge proponent needs only one possible case where the decrease of probability will not in fact be overcome by the peripheral evidence that the objection appeals to.<sup>219</sup>

Actually, there is one way of taking the case and turning it in favor of the Open Knowledge Argument. Let us restrict the case so that a visitor to the zoo knows that the zoo has a designated “zebra” pen, but doesn’t know that it contains a zebra (z), nor does he know that it does not contain a disguised mule (whether with some indication (idm) or without (dm\*)). Let the number of zebra-looking animals be N. Before coming to the pen area the probability of dm\* is  $\underline{DM}^*/N$ , but after looking to see what the animal and the surroundings look like – seeing that the animal looks like a zebra and finding no indication of it being otherwise – the probability of dm\* rises. It is now  $\underline{DM}^*/(N-\underline{IDM})$ . That is, the probability of there being a hoax or some other form of deception rises.

One could claim in reaction that I have changed the proposition from *the animal is not a painted mule*, to something else (perhaps: *the animal is not a mule deceptively made to appear as a zebra*). This is true, but first, the previous argument makes it implausible that the probability of painted mules given the total evidence drops. And second, the proposition, i.e. the animal in the pen is not a cleverly disguised mule, can be taken without too much of

---

$\Pr(\neg dm|e \wedge a) > \Pr(\neg dm)$  (where *e* is the total evidence including the absence of telling marks). This claim is true iff  $\Pr(e \wedge a|\neg dm) > \Pr(e \wedge a|dm)$ . Proof:

1.  $\Pr(\neg dm|e \wedge a) > \Pr(\neg dm)$
2.  $\Pr(e \wedge a|\neg dm)\Pr(\neg dm)/(\Pr(e \wedge a|\neg dm)\Pr(\neg dm)+\Pr(e \wedge a|dm)\Pr(dm)) > \Pr(\neg dm)$  [Bayes' Theorem and law of total probability]
3.  $\Pr(e \wedge a|\neg dm) > \Pr(e \wedge a|\neg dm)\Pr(\neg dm)+\Pr(e \wedge a|dm)\Pr(dm)$
4.  $1 > \Pr(\neg dm)+\Pr(e \wedge a|dm)\Pr(dm)/\Pr(e \wedge a|\neg dm)$
5.  $\Pr(dm) > \Pr(e \wedge a|dm)\Pr(dm)/\Pr(e \wedge a|\neg dm)$  [1 - Pr(-A) = Pr(A)]
6.  $\Pr(e \wedge a|\neg dm) > \Pr(e \wedge a|dm)$  [ $\div \Pr(dm)$ ,  $\times \Pr(e \wedge a|\neg dm)$ ]

$$\frac{\Pr(e \wedge a \wedge \neg dm)}{\Pr(\neg dm)} > \frac{\Pr(e \wedge a \wedge dm)}{\Pr(dm)}$$

which by the definition of conditional probability translates into: Using the notation I adopted, the proportions need to be as follows:  $\underline{Z}/\underline{O-DM} > \underline{DM}^*/\underline{DM}$  iff  $\underline{DM}^* \times (\underline{O-DM}) < \underline{DM} \times \underline{Z}$ . As I have claimed, although this is not impossible, it’s unlikely and doesn’t seem to accord with any intuitive knowledge criterion. In fact it seems to me to be a *reductio* of any conception of knowledge that entails that these proportions must hold *a priori*. Perhaps one would be inclined to restrict the calculation in some way. Yet to impose further restrictions is simply to claim that the visual evidence always has a very limited impact for the purposes of identifying and thereby coming to know what object one is facing.

<sup>219</sup> Here too thanks to Professor Pagin.

an interpretive stretch, to include absence of indications to the contrary. But even waving this claim about the original proposition, we now have the present case as a new and more properly founded example supporting the objectives of the Open Knowledge Argument.

So to conclude the rejoinder to the objection, I think both paths of responding to the objection are viable. First, not all propositions will have some evidence counting in their favor. Second, in place of (KE) one can safely use (NED) for the purposes of the Open Knowledge Argument. I have found no reason to think that anything special is going on in cases of knowledge that is not the case with respect to evidence more generally. In fact, given the proof that any underdetermined proposition has propositions that follow from it that are not supported by the evidence, all that might be shown by the objection is that some cases are not of this type. Yet even the case that might seem to accord with the objection, namely, the zebra case, has been found to rely on some very peculiar initial conditions. Thus, given the openness of evidential support and given the plausible assumption that some knowledge is underdetermined (i.e. falls within the domain of dispute) and accords with (NED), the Open Knowledge Argument seems to be on firm ground.

#### 5.1.1.2. An Objection to (NED)<sup>220</sup>

What is knowledge? What is evidence? What is the relation between knowledge and evidence? These are notoriously difficult questions and I cannot pretend to be able to answer them. But nevertheless we can gain a better understanding of these notions. One way to approach this, the way in which I have so far tried to guide the discussion, is by exploring the principles that we think pertain to these notions and see if there is cohesion among them. Often they come into conflict and we need to make a choice. The attempt here will be to make a choice between which three unappealing claims about knowledge and evidence we need to accept given some basic assumptions. These are the more basic assumptions (I repeat them here for convenience):

Fallibilism: One can know that  $p$  even though one's evidence (including background information) does not conclusively establish that  $p$ . This means that one can know that  $p$  even though one's epistemic posterior probability for  $p$  (however we want to spell it out) is less than 1.

(EC): The *evidence for* relation is constrained by the probabilistic evidence criteria: If evidence  $e$  supports the truth of proposition  $p$ , then the conditional probability of  $p$  given  $e$  –  $\Pr(p|e)$  – is greater than or at least equal to the unconditional probability of  $p$  –  $\Pr(p)$ .

---

<sup>220</sup> The discussion of this section reacts to problems for (NED) posed by Professor Pagin.

Now the problem is that these assumptions lead to the following choice between three claims none of which seem like anything we would want a theory to include (at least initially). We need, then, to accept one of the following:

- a One can know a proposition even though relative to a time where one did not know this proposition to be true, one merely gains information that *lowers* its probability. Given (EC) this means that one could come to know that  $p$  without additional evidence.
- b The order in which one receives evidence makes a crucial difference. Two subjects may have the same prior probability same evidence and same posterior probability yet one knows and one does not.
- c subject may lose knowledge even though all she has is slightly undermining evidence (and her posterior probability surpasses any plausible knowledge threshold).

To see that we are faced with this choice (given our more basic assumptions) let us consider the following stories.

**Case A:** Subject A goes to a zoo not knowing whether there is a zebra, a disguised mule or whether there is neither, at the zoo. He has a high prior (subjective) probability given his background information that there is no painted mule at the zoo he is visiting, yet he does not know it. At the entrance, he is warned by the ticket salesperson that there will be no disguised mule in the pen unless there are signs posted at the pen that say that there is a painted mule in the pen. A proceeds and sees what looks to him to be a zebra at the pen. A believes it is a zebra, but since he cannot tell a zebra from a painted mule his total evidence slightly lowers the probability of there not being a disguised mule in the pen. Had he not been tipped-off at the entrance, his evidence would further lower the (rational) probability of there not being a disguised mule at the pen.

**Case B:** Subject B has the same prior probability that A had for there not being a disguised mule in the pen upon coming into the zoo. She is not tipped-off by the ticket salesperson as A was, but goes directly to the pen. She too has her probability for there not being a mule in the pen lowered since she (like A) cannot tell a disguised mule from a zebra. Her probability for there not being a disguised mule in the pen is lower than A's at the pen. However, when she leaves the zoo, the salesperson tells her what he told A. Thus her probability for there not being a painted mule in the pen goes up and she leaves the zoo with the same posterior probability as A.

We will soon compare these cases to two further stories, but for now let us see how a-c relate to the A and B stories. To make things a bit easier let us adopt the conventions above:  $Z$  = the animal in the pen is a zebra;  $DM$  =

the animal in the pen is a mule made to look to a non expert as a zebra; ZL = the animal in the pen looks like a zebra (or a well disguised mule); SP = the salesperson says that if there is no indication at the pen, there are no painted mules there; NS = there are no signs at the pen indicating that there is a disguised mule in the pen.  $\Pr_A(\bullet)$  = A's rational probability function. (Similarly with B when the subscript changed).

We have four times: before going to the zoo =  $t$ , going into the zoo =  $t'$ , at the pen =  $t''$ , and coming out =  $t'''$ . These will be indicated along side the probability function thus:  $\Pr_{A,t}(\bullet)$  (= the unconditional probability of  $\bullet$ , for A, at  $t'$ ).

Before going into the zoo A and B have the same probability for  $\neg DM$ ,  $\Pr_{A,t}(\neg DM) = \Pr_{B,t}(\neg DM)$ , and they persist with this same probability at the zoo entrance (at  $t'$ ). Whether A's conditional probability function changes at  $t'$ , is an issue we will come back to. For now, let's assume not. At the pen, the following holds:

$$\Pr_{A,t'}(\neg DM) > \Pr_{A,t''}(\neg DM) > \Pr_{B,t''}(\neg DM)$$

The reason for this is that  $\Pr_{A,B}(\neg DM|LZ) = \Pr_{A,B}(\neg DM|LZ \wedge SP) < \Pr_{A,B}(\neg DM|LZ \wedge SP \wedge NS) < \Pr_{A,B}(\neg DM)$ . The lack of signs, let us suppose, without the information from the salesperson adds nothing to the evidence that supports the hypothesis that there is a mule in the pen, i.e. the animal looking like a zebra or a disguised mule. We are assuming that the evidence for the presence of a disguised mule, the evidence of an animal which looks like a disguised mule, counts more for there being a disguised mule in the pen than the lack of signs and the salespersons warning combined. If this sounds wrong we can imagine that A and B merely hear some of what the salesperson says such that they are not very confident that he actually was asserting SP. He may have been saying that this is how it is in other zoos. Also, in accordance with standard Bayesian scruples, at  $t'''$  A and B adjust their probabilities to their evidence and since they started off with the same priors, we have;  $\Pr_{A,B}(\neg DM|LZ \wedge SP \wedge NS) = \Pr_{A,B,t'''}(\neg DM)$ .

Now given (EC), A has no new evidence at the pen for  $\neg PM$ , or any time after and hence if we want to reject a, we will have to agree that he does not come to know  $\neg DM$ . However, B does have evidence since  $\Pr_{B,t''}(\neg DM) < \Pr_{B,t'''}(\neg DM)$  (even though it is lower than  $\Pr_{B,t}(\neg DM)$ ). Hence, if we reject a, A does not know that  $\neg DM$  when leaving the zoo while B might.

It seems doubtful that A can come to know  $\neg DM$  by probability lowering evidence. But if A does not, neither should B and yet, B does have evidence when coming out of the zoo relative to the low probability B assigns to  $\neg DM$  at the pen. But if we let B know that  $\neg DM$ , we will have to accept option b. The order of the evidence will matter for what one comes to know.

The best solution, it seems, is to reject both a and b by quantifying over times in the following way (which is a version of (NED)):

For all times  $t'$  and  $t''$ , if S does not know that  $p$  at  $t'$ , and gains no evidence in favor of the truth of  $p$  between  $t'$  and (the later)  $t''$ , then S does not know that  $p$  at  $t''$ .

Regarding our two subjects A and B, the principle seems to give the right result. Both A and B came into the zoo without knowing that  $\neg DM$  and they leave without knowing that  $\neg DM$  simply since the evidence raised the probability of DM more than it lowered it. It seems right, then, to say that both did not know that  $\neg DM$  before going to the zoo (as we have assumed), and they leave the zoo ignorant. Their posterior and prior probabilities are the same and the order of receiving the evidence should not matter. Since they both have a greater posterior probability relative to the time they were ignorant for DM, they leave the zoo not knowing that  $\neg DM$ .

But this response will leave us with the odd position of having to accept c. The following problem regarding a third case C will bring this out more clearly.

**Case C:** Subject C goes into the same zoo with the same probability as A and B had for  $\neg DM$ . Our trusty salesperson tells C that there is no disguised mule in the pen. This raises the probability of  $\neg DM$  and we can suppose that in this way (even though C has no idea about how reliable our salesperson is) he comes to know that  $\neg DM$ . We are to imagine that this item of evidence is not so strong since seeing the zebra in the pen, S's evidence of a disguised mule decreases the probability of  $\neg DM$  in such a way that it equals A and B's probability to what it was leaving the zoo (at  $t''$ ).

If we accept b, we must say that C does not know that there is no painted mule in the pen. The reason for this is that if someone were to first go to the pen, the probability would first be lowered and then raised upon hearing that there is no mule in the pen to a level beneath the one this person came into the zoo with. Since we have already seen that this cannot be a case of knowledge (if we want to accept a and b), we must say that C loses her knowledge upon receiving slight disconfirming evidence. Hence, we must accept c.

I think that this is the best option. Yet I admit that as it stands it seems to be an unattractive one. Before trying to explain why I take this to be the best option, let us see how the probability changes for C with regard to  $\neg DM$ .

$$\Pr_{C,t''}(\neg DM) = \Pr_{C,t'}(\neg DM) < \Pr_{A,B,C,t}(\neg DM) < \Pr_{C,t}(\neg DM)$$

We are also imagining that  $K_{B,t}(\neg DM)$  and the implausible result is that  $\neg K_{C,t''}(\neg DM)$  and  $\neg K_{C,t'}(\neg DM)$ . In words this is a case where upon coming to the zoo C does not know that there is no disguised mule in the pen, then learns that that this is the case, and upon receiving slight disconfirming evidence, C loses this knowledge.

There are two things we need to think of at this point. One is whether we actually need to agree that  $K_{C,t}(\neg DM)$ . The second thing to think of is whether assuming that  $K_{C,t}(\neg DM)$ , there is no way to make  $\neg K_{C,t}(\neg DM)$  more plausible. I think that whether or not we answer the first question affirmatively the answer to the second question is affirmative.

Why is it doubtful that  $K_{C,t}(\neg DM)$ ? After all,  $\Pr_{C,t}(\neg DM)$  is greater than the initial probability, and if there is a knowledge threshold, it surely surpasses it (it must be the case that  $\Pr_{A,B,C,t}(\neg DM)$  is above the threshold if it is a constant threshold for both  $Z$  and  $\neg DM$ ). One reason to doubt that  $C$  has this knowledge is that the evidence would not be strong enough for a subject who had a lower prior probability. Imagine we had a subject  $D$  who has a lower prior probability for  $\neg DM$  than  $C$  has. The question we need to ask ourselves is whether the evidence  $C$  got in favor of  $\neg PM$  is good enough for  $D$  to know that  $\neg PM$ . If it is, the next question we need to ask ourselves is whether the evidence  $ZL$  that lowered the probability of  $\neg DM$  for  $C$  beneath what it was prior to entering the zoo, would do the same for subject  $D$ . If the conditional probabilities are objective (as most subjective Bayesians believe),<sup>221</sup> then it must. And this means that the counter evidence is not as weak as we have imagined, or alternatively it would not be strong enough for  $D$  to come to know  $\neg PM$ . Either way, we have a motivation for saying that  $C$  loses her knowledge at  $t'$  (which is the answer to our second query).

It might seem, however, that this suggestion is *ad hoc*. The point is, however, more general and when looked at carefully makes good sense. Let us imagine that the evidence  $C$  receives would have been good enough for  $D$  (with a lower prior probability for  $\neg DM$ ) to know  $\neg DM$ . No matter how strong this evidence is, we know that the counter evidence is stronger since  $C$ 's probability for  $\neg DM$  upon leaving the zoo is lower (if this case is to pose any problem in the first place) than it was when  $C$  came into the zoo. This is the case no matter how strong or weak we believe  $C$ 's evidence must be for him to know that  $\neg DM$ . Since the counter evidence will be stronger than whatever we require, accepting  $c$  no longer seems objectionable. First, if we think that in order to know a proposition one needs to have strong evidence, the kind of evidence that would raise the probability of a proposition for a large range of prior probabilities<sup>222</sup>, we need not accept  $a$ ,  $b$  or  $c$ . The reason

---

<sup>221</sup> Here what I mean is that the likelihoods are objective, i.e. the probability of the hypothesis given the evidence. The likelihood divided by the prior probability of the evidence is the Bayesian multiplier that will determine whether the evidence counts for or against a proposition. On an objective reading of the likelihoods any body of evidence will count in favor or against the same propositions for any pair of subjects. For more see Appendix specifically with regard to convergence theorems.

<sup>222</sup> We cannot just allow any range of prior probabilities. Say we do, say for instance, that we require that the body of evidence will raise the probability no matter how low for a proposition  $p$ , such that in order to know that  $p$  one needs to have evidence that raises the probability of  $p$ . Since for any given body of evidence no matter how strong we can find a prior probability distribution that will be so low that one who had this probability would rationally remain



for this is that in order to lose knowledge it is now clear that the counter evidence must be strong as well. Second, we will need to accept  $c$  if we are relaxed with respect to how strong evidence needs to be in order for a subject to come to know that a proposition is true. Nevertheless this will not be an intuitively bad position since accepting that slight disconfirming evidence can destroy knowledge will be explained by the fact that this knowledge was acquired through even slighter evidence. In other words, a subject may lose knowledge by slight disconfirming evidence (above the knowledge threshold if there is one) if the way in which she came to this knowledge is by evidence that is even slighter.

The moral of our stories is that all subjects came in and went out of the zoo ignorant of whether or not there is a disguised mule in the zoo. Our problem with subject  $C$  was only an apparent problem.  $C$ 's prior probably upon coming to the zoo is greater than the one he walked out of the zoo with, hence the counter evidence he got is stronger than whatever evidence for  $\neg DM$  he encountered. If so, it seems plausible to say that he came in and went out not knowing that  $\neg DM$  and we can leave the question of whether the weak evidence ever allowed him to know it for further investigation. The apparent problems we encountered are the result of the possibility of dividing a body of evidence in different ways such that they either first count against and then for a proposition or vice-versa. The principles we accepted seem resilient to these kinds of worries upon closer scrutiny.

## 5.2. Knowledge Without Evidence

There are other more general ways that a proponent of knowledge closure may respond to the Open Knowledge Argument. One suggestion is that knowledge need not depend on evidence. The propositions that closure entails to be known are properly deduced from true, justified premises, and proper deduction is as good a justification as anyone can ask for. Alternatively one may accept the dependence of knowledge on evidence, yet claim that the original known proposition ( $p$ ) is one's evidence.<sup>223</sup> Since  $p$  implies

---

extremely confident that it is false, being free with the prior probabilities does not seem like a good idea. Nevertheless this objection to allowing prior probabilities to be as one chooses neglects the fact that evidence is not the only necessary requirement on knowledge. We might want to place rational confidence at some threshold as such a requirement.

<sup>223</sup> These two replies are not meant to be exhaustive, but seem to be the most pertinent. I cannot hope to survey all the possible reactions. I do claim however, that there are only two general possibilities here. Either reject (NED) or claim that somehow it is not violated. This first option seems highly questionable though this rejection might be given an explanation I am unaware of. The second seems to be more promising and can take on several forms only some of which have been mentioned. Hawthorne and Cohen raise a similar worry not with regard to evidence in general, but to cases of easy knowledge that turn, as I claimed in the main text, on failure of evidence closure. Hawthorne mentions three replies: 1. A priori contingent knowledge is much more widespread than we may have imagined. 2. Skepticism. 3.

$q$ , if  $p$  is one's evidence one can know  $q$ . I will consider these options in turn, the first in this section and the second in the next.

Knowledge indeed requires justification and a priori inference from a single premise is usually good justification for belief. Thus it may be claimed that a belief that is the fruit of valid inference from knowledge is justified, and is hence known, or that it puts the subject in a position to know. As it is undisputed that knowledge of the sort under consideration is normally based on evidence, however, proponents of this view must provide the source of justification in cases where evidence is lacking. Specifically, since the justificatory force of inference depends on the justification of the premises, when this justification does not transmit (as is the case in the examples above) something else must be proposed as the source of justification for the conclusion.

Here it may be suggested that the dependence claim (NED) should be denied.<sup>224</sup> The inference from  $p$  is itself a source of epistemic justification for  $q$ ,

---

When a proposition is known, its probability is raised to 1, with a widespread violation of van Fraassen's Reflection Principle (Hawthorne 2004a: 75-7).

I will have more to say some of these options in the main text, it might be worth, however to say something that directly addresses them here. (1), I think, would quite surprising if it turned out to be true. Even if it could be argued for with respect to cases of easy knowledge, it is hard to see how in every instance of failure of evidential closure a priori knowledge would be present. It seems highly unlikely that one can have a priori knowledge that cars are not frequently stolen or that zoos do not paint mules to look like zebras, etc. Even granting a *priori* contingent knowledge of this sort, or denying that it must be strictly a *priori* (perhaps it is inductive background knowledge), will not relieve any worries since this knowledge pertains to types not to tokens. In other words, although I may know that cars are not often stolen, I certainly don't know by memory that *my car today* has not been stolen. If memory and induction won't deliver such knowledge, claiming to know a *priori* that my car is not stolen seems questionable. (I mention in passing that there seems to be a regress argument if it is claimed that one will (always?) know by induction that one's car today has not been stolen and the like.) For the suggestion to work, an explanation should be offered for the systematic divorce between evidence and knowledge it entails. In other words, the challenge for those who want to advance such a position is to spell it out not only with respect to certain cases but systematically. This is because we have a proof that generalizes underdetermined knowledge to unsupported propositions that can be validly inferred from knowledge. Even if such an account can be spelled out I doubt it will be as simple as the present open knowledge suggestion. Moreover, it might not solve all the problems for knowledge outlined in Chapters 2 and 3 as well as the other challenges for closed knowledge.

The second response (2) can be improved on by rejection of closure. Thus one can be skeptical of knowledge (or even the possibility of knowledge) of problematic propositions such as that one is not a brain in a vat, that there are mind-independent objects etc., without denying knowledge of ordinary empirical truths.

(3) might be the best track for the defender of closure to take, but this entails, as I will argue below 5.3, several unfavorable consequences. The most developed account along these lines is Williamson's.

Cohen's favored reaction (at least at one time it was) is to claim that closure is inapplicable to what Sosa has called "animal knowledge." It seems that even Cohen is not entirely happy with this response, perhaps because he recognizes that it amounts to a significant limitation of epistemic closure (see Cohen 2002).

<sup>224</sup> In effect, denial of dependence is not enough for this line of argument to work. Since closure must be valid in all cases, what must be accepted is the claim that in all cases where

not supervening on the evidential justification for  $p$ .<sup>225</sup> But this seems to conflict with the general worry: Presumably the inferred proposition,  $q$ , is not known prior to its inference from  $p$ . So prior to acquiring evidence  $e$ ,  $q$  was not known. But since  $e$  reduces the probability of  $q$ , it is apparently not in virtue of the acquisition of  $e$  that  $q$  came to be known. If  $e$  does not provide the justification enabling knowledge of  $q$ , its role must be in facilitating the inference. Indeed without  $e$ ,  $p$  would not have been known and  $q$  could not have been inferred from knowledge. But if  $e$  cannot justify  $q$  and if the inference from knowledge of  $q$  is not available without  $e$ , how can the inference provide more than  $e$  itself could?

In effect what I am recommending on behalf of the open knowledge account, is a step back to enable a view of the entire process of acquiring (the contested) knowledge that  $q$ . From this vantage point, knowledge that  $q$  seems doubtful: Taking the entire process of (would be) knowledge that  $q$ , if it cannot be gained directly through  $e$ , finding a proposition that can be known from  $e$  and then inferring  $q$  from that proposition, would seem in general to be a highly questionable method for gaining new knowledge. Yet this in effect is what this suggestion amounts to. From the perspective of the whole process, this seems like not much better than wishful thinking. And this, I think, is what (NED) brings out quite straightforwardly; aside from peculiar cases, a proposition cannot come to be known on the basis of evidence that lowers its probability, directly or indirectly. The present proposal amounts to claim that probability-lowering evidence is a general method for knowledge gain.

In addition, even if knowledge does not rely on evidence, counter-evidence to one's belief can defeat the belief's status as knowledge. Now in the examples we have been considering it is not just that one lacks evidence for the  $q$ -propositions, rather one actually has evidence against them (on the assumption that decrease of probability constitutes counter-evidence). In the aforementioned cases one has evidence that tells against the truth of the  $q$ -propositions. Given the evidence, that is, the probability of  $q$  decreases. Thus the divorce from evidence required of this defender of epistemic closure is more radical than merely claiming that knowledge can be had without evidence. Knowledge, she must claim, can be had in the face of counter-evidence. If epistemic closure is to hold, that is, not only is the connection between knowledge and possession of evidence to be severed, it must also be

---

the evidence does not entail the known proposition, one can know that a proposition is true even though relative to a prior situation in which the proposition was not known to be true, one has gained probability lowering evidence, which allow one to now *know*. Thus, if one does not deny (NED) one would need to claim that although knowledge of the derived proposition is not possible on the basis of the evidence itself, it can be known by inferring it from a proposition that can be known under the circumstances. But then, one might ask, where does the added knowledge-promoting feature come from? It cannot come from the inference for the inference in itself does not provide added justification.

<sup>225</sup> Something along these lines is hinted at by Silins (2005).

admitted that a belief formed in contrast to one's evidence can bring about new knowledge. This I take to be a highly implausible concession. Even if (NED) fails sometimes, it seems highly questionable that it might fail systematically for all propositions that are underdetermined by the evidence that allows them to qualify as knowledge. The inference itself then, is not and cannot be considered to be an independent source of knowledge.

### 5.3. Knowledge as Evidence

A defender of epistemic closure may try a different direction. Perhaps she might suggest that the inference of  $q$  from  $p$  is itself part of the evidence. Since the truth of  $p$  clearly speaks in favor of the truth of  $q$  (in fact, it guarantees it), if  $p$  is now known, it can itself be the new evidence (or reason) on the basis of which  $q$  becomes known. (NED) then is not violated since its antecedent is not satisfied. Between  $t_0$  and  $t_1$  evidence is gained, namely,  $p$  itself.<sup>226</sup>

As far as the open knowledge proponent is concerned, there is one way of understanding this suggestion that leaves everything as it was. Another way, however, would be a proper reply to the open knowledge argument, but has undesirable implications. I will argue that this later suggestion would entail sufficiently far-reaching consequences, consequences that do not compare well with the open knowledge account. In other words, if this suggestion is right, we need to change the basic notions of probability and its relation to chance; we need to give up on justification closure; to surrender to the Lottery and Preface paradoxes; to give up on the significance for knowledge of the distinction between multi- and single premise closure<sup>227</sup> (accept conjunction introduction closure); we need to surrender fallible (or underdetermined) knowledge. This means that in saving knowledge closure we are in effect committing ourselves theoretical changes regarding knowledge and the notions that are related to it that are worse than the change we need to make when accepting knowledge openness. Moreover, we would lose an elegant

<sup>226</sup> In light of the discussion in Chapter 1, there is one way of taking the current reply to the open knowledge argument as a serious worry, another way is a non starter. Nearly everyone agrees that *modus ponens* is a valid form of inference. If  $p$  is true, and *if  $p$  then  $q$*  is true, then  $q$  is true. This claim has little to do with the current suggestion. The question is whether  $q$  is known, not whether it is true. From the *factivity* of knowledge and *modus ponens*, we can derive the truth of  $q$  but not its being known. The difference is between the following forms of inference:  $K(p \rightarrow q) \vdash Kp \rightarrow Kq$ , and  $Kp \wedge (Kp \rightarrow Kq) \vdash Kq$ . The open knowledge advocate is not contesting the second, but only the first form of inference.

<sup>227</sup> This consequence – that single-premise as well as multi-premise closure for knowledge would hold – will rightly be viewed more as an advantage than a shortcoming. It preserves the role of inference regarding knowledge generally (See Chapter 1). Nevertheless it spells trouble for the role of inference with regard to justification and has bad consequences of its own some of which will be the focus of section 5.3.3.

and simple way of coping with many problems that open knowledge can account for.

The next section spells out the unproblematic way (from the standpoint of open knowledge) of accepting the idea that  $p$  itself is what supports  $q$ . The sections that follow will focus on the problems that the second way brings with it through a criticism of the most worked out theory of this variety, namely, Timothy Williamson's safety account of knowledge.

Before proceeding I want to make it perfectly clear that I do not claim to have in my possession an a priori argument for the conclusion that any other way to respond to the open knowledge argument besides those I have proposed is moribund. Rather the claim boils down to the following: One, there is a formidable challenge to open knowledge from open evidence. Two, there is a simple way to solve the challenge which is to accept the conclusion of the argument that knowledge is open. Three, besides the other ways to respond to the open knowledge argument which I have yet to consider in detail, the prospects of such a reply are not particularly rosy. That is, a denial of (NED) across the board is not at all appealing – there is no reason to think that probability-lowering evidence is a general way to obtain new knowledge. Although there are proper replies that do not conflict with (NED) such as the claim that propositions that follow from  $p$  and are not supported by the evidence are known a priori, such replies on the face of it do not seem plausible. It does not seem like we can know a priori all the consequences that have their probability lowered by the evidence that supports our knowledge.<sup>228</sup> If we are to be fallibilists and accept (NED) the Open Knowledge Argument poses a formidable challenge to closed knowledge. Fourth, the denial of closure gives us a simple and elegant reply to many central epistemological problems some of which were the topic of Chapters 2 and 3. Fifth, we have an explanation of the oddness of denial of closure due to the oddness of the underlying structure of the *evidence for* relation. In other words, we are going to have to get used to the claim that relative to a fixed body of evidence we have evidence for  $p$ -and- $q$  but have no evidence for  $p$  and no new evidence for  $q$ . Sixth, this explanation can meet the most forceful arguments against open knowledge advanced by John Hawthorne. In sum, we have a simple and elegant way to avoid many of our problems in epistemology including the open knowledge argument and an explanation of why this track has been avoided despite its advantages. These considerations do not entail knowledge openness but they do give us excellent reason to reconsider our view of knowledge closure.

With this clarification in place, let us proceed and consider an entirely different type of reply to the argument. That is, the idea that whenever knowledge is gained it changes the evidential situation.

---

<sup>228</sup> I am not sure in any case that this reply is consistent with fallibilism.

### 5.3.1. Knowledge as Evidence: Jeffrey Conditionalization

Richard Jeffrey (1964, 1965) proposed a rule of rational belief update that yields Bayesian Conditionalization<sup>229</sup> as a special case. The basic idea is to exploit the law of total probability in order to reflect rational changes in an evidential state that do not amount to propositions becoming certain. The standard rule is that a subject who becomes certain of the truth a proposition  $e$ , should change her credence regarding the truth of proposition  $p$  in accordance with the laws of probability in the following way:

$$\text{Standard Conditionalization: } \Pr^1(p) = \Pr^1(p|e) = \Pr^0(p|e)$$

where “ $\Pr^0$ ” is a rational credence function before the evidence becomes known for certain and “ $\Pr^1$ ” after. It is understood that certain knowledge means that  $\Pr^1(e)=1$ . But the Standard Conditionalization rule does not say what a rational agent should do if the probability of  $e$  changes so that the agent is now more confident but still less than certain that  $e$  is true. Here Jeffrey advises taking advantage of the law of total probability, to get a more general form of conditionalization regularly labeled “Jeffrey Conditioning:”

$$\text{Jeffrey Conditionalization: } \Pr^1(p) = \Pr^1(e)\Pr^0(p|e) + \Pr^1(\neg e)\Pr^0(p|\neg e)$$

Trivially, Jeffrey Conditionalization is the same as Standard Bayesian Conditionalization if  $\Pr^1(e)=1$  since the posterior probability of  $\neg e$  is 0.<sup>230</sup>

Now the claim we are interested in exploring is whether when we take known propositions as our new evidence we can avoid the open knowledge argument. That is, whether (NED) can be respected by a closed knowledge infallibilist view. Let us first take the simplest case where we derive  $p$  from  $p$  itself. This is a case that the open knowledge proponent will not question. The question, however, is whether the probability of  $p$  will be promoted if we allow any known proposition to serve as evidence. The answer is negative, the new probability of  $p$  will not be promoted to 1 when we Jeffrey conditionalize on  $p$  itself, rather, it will remain the same:

$$\Pr^1(p) = \Pr^1(p)\Pr^0(p|p) + \Pr^1(\neg p)\Pr^0(p|\neg p) = \Pr^1(p)\Pr^0(p|p)$$

<sup>229</sup> For more on Bayesian Conditionalization, see Appendix.

<sup>230</sup> When we are dealing with a change in a partition (a set of exclusive and exhaustive propositions) of evidence statements (not merely the simple case of  $e$  and  $\neg e$ ), we take the sum of multiplied probabilities to determine the posterior probability of  $p$ :

$$\Pr^1(p) = \sum_{1 \leq i \leq n} \Pr^1(e_i)\Pr^0(p|e_i), \text{ where } \{e_1, e_2, \dots, e_n\} \text{ is the partition.}$$

This result is a good one, since we do not want the fact that an inference was made to change the rational credence for a proposition nor the evidential situation itself.

Now suppose you get to know that  $p$  with some evidence  $e$  at time  $t_1$  but that  $e$  underdetermines  $p$ , so that  $\Pr(p|e) < 1$ . We can assume that  $e$  is now known for certain at  $t_1$  so that  $\Pr^0(p) = \Pr^1(p|e) = \Pr^0(p|e) \gg \Pr^1(p)$ . Suppose you also know that  $q$  follows from  $p$  and you properly infer  $q$  from  $p$ , thereby basing your new posterior probability for  $q$  on the basis of this inference. What is the new probability of  $q$  given your new evidence (where your new evidence includes  $p$  now that it has come to be known)? Using Jeffrey Conditionalization the answer is:

$$\Pr^1(q) = \Pr^1(p)\Pr^0(q|p) + \Pr^1(\neg p)\Pr^0(q|\neg p)$$

Since  $\Pr^0(q|p) = 1$ , the posterior probability of  $q$  will be at least as high as  $\Pr^1(p)$ . So now the question we should consider is whether  $\Pr^0(q)$  is greater, equal, or smaller, than  $\Pr^1(\neg p)\Pr^0(q|\neg p) + \Pr^1(p)$ . If it is smaller or equal, nothing substantial will have changed relative to the cases we have previously considered. So all that is needed for the purposes of an open knowledge proponent is a case where we can show that in fact there is a reduction of posterior probability of  $q$  relative to the prior probability of  $q$  (between  $t_0$  and  $t_1$ ).

To see that this is possible let us take the following case:  $q$  will be  $\neg(e \wedge \neg p)$  which follows from what we will Jeffrey Conditionalize on which is now known  $e \wedge p$ ,

$$\Pr^1(\neg(e \wedge \neg p)) = \Pr^0(\neg(e \wedge \neg p)|e \wedge p)\Pr^1(e \wedge p) + \Pr^0(\neg(e \wedge \neg p)|\neg(e \wedge p))\Pr^1(\neg(e \wedge p))$$

Since the old and new probability of  $e$  is 1 ( $\Pr^{1,0}(e) = 1$ ), the value of  $\Pr^1(e \wedge p)$  equals that of  $\Pr^1(p)$  (which we assume is less than one). So we need to determine what the value of  $\Pr^0(\neg(e \wedge \neg p)|\neg(e \wedge p))\Pr^1(\neg(e \wedge p))$  is.

The value of  $\Pr^1(\neg(e \wedge p))$ , is that of  $\Pr^1(\neg e \vee \neg p) = \Pr^1(\neg e) + \Pr^1(\neg p) - \Pr^1(\neg e \wedge \neg p) = \Pr^1(\neg p)$  (the other probabilities are zero). Next we need to see what the value of  $\Pr^0(\neg(e \wedge \neg p)|\neg(e \wedge p))$  is. By the conditional probability definition we have

$$\frac{\Pr^0(\neg(e \wedge \neg p)) \wedge (\neg(e \wedge p))}{\Pr^0(\neg(e \wedge p))} = \frac{\Pr^0(\neg e)}{\Pr^0(\neg e \vee \neg p)} = 0$$

and so the new probability of  $\neg(e \wedge \neg p) = \Pr^1(p)$  which we know is not greater than  $\Pr^0(\neg(e \wedge \neg p))$ .<sup>231</sup>

<sup>231</sup> And relative to the time before  $\Pr(e) = 1$ ,  $\Pr^1(\neg(e \wedge \neg p)) > \Pr^0(\neg(e \wedge \neg p)) = \Pr^1(\neg(e \wedge \neg p))$ .

Now if the prior probability of  $q$  is greater than its posterior probability, (the probability of  $q$  has now gone down) the same kind of reasoning regarding the other cases will be applicable. Assuming that is, that we accept the idea that evidence cannot lower the probability of that for which it counts as proper evidence and that probabilities are raised for knowledge (at least in some cases) to less than 1. Even though, as suggested, all knowledge counts as evidence, there is good reason to conclude that since evidence is not closed, neither is knowledge. The non-probabilistic arguments of 4.3 confirm this result.

Besides having a way to answer the objection, the open knowledge advocate can agree that all knowledge is evidence. Whether this is desirable or not is another question that I will not try to evaluate here. An open knowledge advocate can accept the idea that knowledge is evidence in a way that seems to be less objectionable than on the closed knowledge view. The main benefit is that on her view a known proposition  $p$  will not be promoted to probability 1 if it was not 1 at an earlier time conditional on the evidence that allows it to be known ( $\text{Pr}(p|e)$ ).<sup>232</sup> In sum, it seems that the objection to the Open Knowledge Argument turns out to be compatible with open knowledge and poses no threat without further assumptions.

### 5.3.2. Knowledge as Evidence: Standard Conditionalization

Open knowledge has been shown to be compatible with taking all of one's knowledge as one's evidence. (NED) is not violated since gaining knowledge would mean gaining new evidence, but if this evidence does not entail the new knowledge, there are propositions that follow from what one knows that are not known. This also accords with the intuition that, e.g. looking at the airport monitor allows one to know that one's flight will stop in Chicago but not that one's plane will not crash on its way (to take one of Cohen's example).

But there is another way to understand the idea that all knowledge is evidence that would counter the Open Knowledge Argument. This way of understanding the thesis that all knowledge is evidence would allow one to take

---

There is a debate which I am not going to try to enter regarding the question of order matters for Jeffrey Conditionalization. This would appear to be the case if we take both  $e$  and  $p$  in the above example to be simultaneous. Here I merely need one case which also seems to yield the right result. The same kind of result can be obtained with regard to Carnap's example taking  $f$  to be probability 1. See 4.3.4.

<sup>232</sup> I note in passing that besides the suggestion of preserving  $E=K$  with Jeffrey conditionalization and taking  $E$  to be a subset of what is known such as what is known for certain (which is a departure from  $E=K$ ), we might also take  $E$  to be what is known directly. I.e. whatever is known not on the basis of inferring it from knowledge is evidence. Besides the conception of evidence as certainty, the other suggestions are in need of further development. But the basic idea in all of these suggestions is simply to respect the claim that there is no added evidence that somehow creeps in when a proposition is known on the basis of evidence that does not entail the (newly) known proposition.



the evidence as supporting itself. Such a conception of evidence has been developed by Timothy Williamson. Before directly considering the relevant parts of his account, let us first see what the suggestion is more generally.

The general idea is that when one properly infers from one's knowledge, new knowledge is gained not by having the evidence now apply to the inferred proposition, but rather, by treating the known proposition itself as one's new evidence. Closure of knowledge has to do with the extension of evidence to bear on new proposition only in the sense that knowledge is evidence (or rather that all that is known can count as evidence with probability 1).<sup>233</sup>

This idea conflicts with the way we usually think about evidential support. Consider the following illustration of evidential dialogue.

A: Don't worry, the car was not stolen.

B: How do you know?

A: Well, it's in the driveway.

B: And how do you know that?

A: 'cause I parked it there.

This dialogue brings out two aspects of what I will later consider more closely, namely, the backtracking structure of evidential support. The first is that we usually do not regard a belief as knowledge-promoting evidence unless it is itself evidentially supported (or supported by other knowledge-promoting justification). The second is that the degree to which a belief can support other beliefs is limited by the evidence or justification it enjoys. Thus a common continuation of the dialogue above would probably be:

B: But that doesn't prove anything. If it were stolen, you'd still remember having parked it there. In fact, cars are usually stolen from where they had been parked.

One's belief that the car is in the driveway can only be evidence that it was not stolen if it was itself supported by evidence suggesting that the car was not stolen. Similarly, a belief (or fact) that an animal is a zebra cannot support the belief that it is not a disguised mule unless it is itself supported by evidence that tells against the animal being a disguised mule. Thus, suppose A's final response was not memory of parking the car, but rather

---

<sup>233</sup> Here we need to separate two general cases. On the one hand, we have the idea that  $p$  is what allows  $q$  to be known when  $p$  becomes known, and yet  $p$  is not to be considered as evidence. This suggestion conflicts with (NED). On the other hand, the suggestion we are considering now is that  $p$  once it becomes known is part of the subject's evidence. This suggestion can either take the form of Jeffery Conditionalization or Standard Conditionalization. On the former we have just seen that it poses no new threat to the Open Knowledge Argument. On Standard Conditionalization knowledge must have probability 1. Hence, we are dealing with a view that seems to conflict with fallibilism.

A: I'm looking at it.

If A's evidence for believing that the car is in the driveway is of this sort, the doubt about his knowledge that it hasn't been stolen does not arise. The reason is obvious – perception of the car in the driveway supports the belief that it has not been stolen (so that this belief, other things being equal, will count as knowledge).

But the idea that  $p$  or the inference from  $p$  is itself evidence is not only unintuitive. If any item of knowledge is allowed to be knowledge-promoting evidence, non-conclusively based knowledge will provide conclusive knowledge. This is because if  $e$  is non-conclusive evidence enabling knowledge of  $p$ , which implies  $q$ , and  $p$  can be taken as evidence for  $q$ , then  $q$  is conclusively supported by the evidence (in probabilistic terms, given  $p$ , the probability of  $q$  is 1 assuming that the suggestion is that one is allowed to use Standard Conditionalization). Still worse, all knowledge, it may seem, will in effect be based on conclusive evidence. The reason is that once  $p$  is known – no matter what evidence it is based on – if it can serve as one's evidence in accordance with Standard Conditionalization, it will support itself conclusively. In other words, knowing that  $p$ , I can, according to the present proposal, use  $p$  as my evidence, and since I know that  $p$  entails  $p$ , I can generate for myself conclusive evidence for  $p$ . By trivial logical operations, my evidence has been upgraded from fallible/underdetermining evidence that has probability less than 1, to conclusive evidence with probability 1. Inductively based knowledge turns instantaneously into knowledge having the full support of deduction. Moreover, since I have such conclusive evidence in support of  $p$ , I can infer (and therefore know) that any evidence counter to  $p$  is misleading. In simple terms, then, allowing all known propositions to serve as evidence makes knowledge (potentially) infallible. Hence the open knowledge argument is in reality being countered by rejecting the first premise, i.e. by demanding that all knowledge have probability 1.

Adherents of this proposal must admit that a proposition can provide evidential support only to the degree to which it is itself supported. Thus if one's evidence for  $p$  raises the probability that  $p$  is true to 0.8, for instance,  $p$  can provide evidential support no stronger than that. The transmitted evidential support will not be conclusive, but at most 0.8 probability.<sup>234</sup> This is the suggestion of using Jeffrey Conditionalization, which is similar to using the evidence itself and will not serve the purposes of the proponent of closed knowledge. So in order for the suggestion to work it must be explained why in the context of knowledge the proposed theory of evidence will have a different result. If  $p$  cannot support  $q$  to a degree greater than that to which it

---

<sup>234</sup> There are other dimensions to evidential support besides the probability it confers. I will return to this issue below.

is supported, it seems, this must be because its epistemic credibility, so to speak, relies on the support  $p$  itself enjoys, i.e. the support supplied by  $e$ . A piece of evidence can support an item of belief only to the degree to which it is itself supported. Thus, while it can warrant (or even require) belief in  $q$ ,  $p$  offers no epistemic support of its own, but merely the support it received from the evidence on which it is based. In other words, forming a belief on the basis of one's evidence does not improve one's evidential situation. The reason must be because whatever evidential support one has is determined by the evidence one relies on. For otherwise it should be possible for  $q$  to enjoy a degree of support exceeding that provided by  $e$  to  $p$  (contra (NED)).

This is made evident by the following observation. Suppose, given the rate of breakdowns of your watch, the fact that it shows "3:00" raises the probability that it is three o'clock to 0.95. Suppose further that this is enough to know that it is three o'clock and that this knowledge is now your evidence that your watch is accurate. Presumably, since this latter proposition is implied by what you know, its probability is no less than 0.95. Now if you receive some weak evidence suggesting that the watch is malfunctioning, we do not say that since you have stronger evidence that the watch is accurate you know this. We do not weigh the new evidence against  $p$ . The belief that the watch is accurate, it seems, requires some independent support in order to count as knowledge. The support of  $p$  does not aid  $q$  if the latter is not itself supported by the evidence. But now if, as we have seen is possible,  $e$  provides no support for  $q$  (and assuming there's no other source of evidence), how can the mere presence of  $p$  improve one's overall evidential situation?

A further implausible consequence of the proposal that one's knowledge can serve as evidence is unreasonable inflation of knowledge. Having received the final confirmation for my invitation to speak at the departmental colloquium next fall, I know I will be presenting a paper in the fall colloquium. This knowledge, as the present proposal would have it, can support my belief that I will not suffer a fatal disease and die between now and next fall. This belief will in turn justify the belief that I will not collect on my life insurance this year. Do I now have evidence warranting cancellation of my insurance policy?<sup>235</sup> Given that I know I will live to present this paper next fall (and given that I have to put more work into it if I am to make a successful presentation), would I be warranted in canceling my physical checkup scheduled for next week? If you think something funny is happening in such cases due to the high stakes involved in them, think of the watch case. Seeing that the watch shows "3:00" I presumably know that it is three o'clock. It follows from this that if my watch is showing "3:00" it is showing the correct time. It follows further that if your watch shows something other than

---

<sup>235</sup> This example is inspired by Hawthorne (2004a).

“3:00” your watch is mistaken. Would it be reasonable of me to instruct you to reset your watch if it does not read “3:00”?

A further problem is that probability 1 seems to rob one of the resources to distinguish between logical operations on evidence that are guaranteed to preserve a given high *level* of probability and conjunction introduction that does not. If knowledge always has probability 1, then there is no probabilistic difference between knowing each conjunct and knowing a conjunction (no matter how many conjuncts we are speaking of). Even those who accept closure, may think that surely this cannot be the case. But if closure has to do solely with the propositions and not also with the evidential status of the propositions, there is no difference between knowledge of a conjunct and a conjunction even in domains that are uncontroversially underdetermined.<sup>236</sup>

These remarks may seem not to do full justice to Williamson's account, an account that proceeds in the direction outlined above. But the remarks are not meant to target his account. They are directed as criticism towards those who might be tempted to take the known proposition  $p$  itself as one's evidence without making the necessary adjustments. E.g, redefining epistemic probabilities, taking prior probabilities to be objective, rejecting fallibilism, etcetera. Following an exchange between Williamson and John Hawthorne and Maria Lasonen-Aarnio, I will argue that although Williamson makes these and other adjustments, other problems should concern the proponent of closed knowledge.<sup>237</sup>

### 5.3.3. Knowledge as Evidence: Williamson's Account

Williamson's account of knowledge in terms of safely true belief is arguably one of the most notable accounts of contemporary epistemology. Due to its scope and ingenuity it rightfully occupies the center of attention in much of the current writing in the field. I present here three central problems for Williamson's safe-belief theory of knowledge as well as several more minor problems. The central issues are: One, its relation to lottery propositions; two, its relation to Lewis's Principal Principle and to practical deliberation; and three, the difficulty of rescuing a substantive sense of fallibilism within the framework this theory provides.<sup>238</sup> Although I see no plausible way for Williamson to solve these problems while remaining loyal to the fundamentals of his account, I do not argue that such solutions are impossible. My

---

<sup>236</sup> Perhaps one would rather claim that  $p$  is not evidence but is a reason for believing that  $q$ . Thus when  $p$  becomes known it can be the basis for knowledge but is not evidence. This suggestion is in violation with (NED) and so does not address the challenge. Moreover it is very close to a blind insistence on closed knowledge.

<sup>237</sup> John Hawthorne and Maria Lasonen Aarnio, (forthcoming) and Williamson (forthcoming)

<sup>238</sup> This last issue is more of a problem for those who want to adopt Williamson's theory and be fallibilists.

argument is therefore best understood as specifying a set of challenges for Williamsonian epistemology. These challenges might make the open knowledge account seem more attractive. Below after consideration of Williamson's account I will specify some of the benefits that open knowledge enjoys.

### 5.3.3.1. Safety and Chance

According to the safety theory of knowledge, one knows what one believes on the basis of a safely true belief. Since the notion of "safely true belief" is closed under known implication, so is knowledge. Focusing on this feature of the view, Hawthorne and Lasonen-Aarnio have recently made an important contribution to a more comprehensive understanding of the weakness of Williamson's theory of knowledge. After presenting their insightful argument, I will assess Williamson's response and show that it leads to further complications. The following is a formulation of the argument adapted from Williamson's response to their argument:<sup>239</sup>

- (1)  $p_1, \dots, p_n$  are true propositions about the future.
- (2) Each of  $p_1, \dots, p_n$  has the same high chance (objective probability) less than 1.
- (3)  $p_1, \dots, p_n$  are probabilistically independent of each other (in the sense of chance).
- (4) The chance of  $p_1 \wedge \dots \wedge p_n$  is low [for large  $n$ ]. [from 2, 3, 4]
- (5) One believes  $p_1 \wedge \dots \wedge p_n$  on the basis of competent deduction from the premises  $p_1, \dots, p_n$ .
- (6) One knows each of  $p_1, \dots, p_n$ .
- (7) If one believes a conclusion on the basis of competent deduction from a set of premises one knows, one knows the conclusion ('multi-premise closure').
- (8) One knows something only if it has a high chance.

Treat (1)-(5) as an uncontentious description of the example. Relative to them, (6)-(8) form an inconsistent triad:

- (9) One knows  $p_1 \wedge \dots \wedge p_n$ . [from 5, 6, 7]
- (10) One does not know  $p_1 \wedge \dots \wedge p_n$ . [from 4, 8]

Which of (6)-(8) should we give up?

One might be tempted to reject (7) in light of this argument. The reason for this rejection – similar to the one considered in Chapter 1 – might be the

---

<sup>239</sup> The problem has obvious affinities with the preface paradox and relates to issues discussed in Chapter 1. For the original statement of the paradox see Makinson (1965).

connection between knowledge and justification. Take any threshold of justification and assume that in order to be justified in believing a proposition  $p$ , one's total evidence  $e$  must sufficiently support  $p$ . Trivially, one might be (evidentially) justified in believing  $q$  since  $\Pr(q|e) > r$  (where  $r$  is the threshold of justification in the unit interval which falls short of 1) and might also be justified in believing  $q'$ , since it too surpasses the threshold:  $\Pr(q'|e) > r$ . And yet, as the case may be,  $\Pr(q' \wedge q|e) < r$ . Hence one will not be justified in believing what one competently deduces from one's (evidentially) justified beliefs (by conjunction introduction) on one's total evidence. If one thinks, that what holds for justification holds for knowledge, the natural reaction to the (1)-(10) contradiction, is to reject (7).

A similar situation holds for single-premise closure<sup>240</sup> even within the Williamsonian framework. Suppose one is evidentially justified in believing  $p$  (where one's total evidence is  $e$ ). So  $\Pr(p|e)$  is very high but less than 1. Suppose further that evidence is so construed (as it is on Williamson's account) that necessarily if  $e$  is evidence for  $p$  then  $\Pr(p|e) > \Pr(p)$  (where "Pr( $\bullet$ )" is one's initial rational epistemic probability function). Now, for every proposition  $p$ , if  $\Pr(p|e) < 1$ , there is a proposition  $q$  that follows from  $p$  such that  $\Pr(q|e) < \Pr(q)$ .<sup>241</sup> Hence, although (by assumption) one has evidence for  $p$ , one will not have evidence for a proposition that logically follows from it (assuming Williamson's own principle of evidence). What this means is that justification is not closed under competent *single premise* deductions, or alternatively, one might say that one is justified without having evidence.<sup>242,243</sup>

---

<sup>240</sup> The distinction between multi- and single-premise closure is considered thoroughly with respect to knowledge in (Hawthorne, 2004a, 33). He there explores (without endorsing) the possibility of holding on to the single-premise version while discarding the multi-premise closure principle (e.g. Hawthorne, 2004, 141, 146, 154, 185-6). Maria Lasonen-Aarnio questions this distinction with regard to knowledge in her (2008) paper. See Chapter 1.3 for a related discussion.

<sup>241</sup> See proof in p. 98.

<sup>242</sup> Since on Williamson's view knowledge=evidence, in cases where one does not know but is merely justified, one can lack evidence. For instance, suppose my plane will stop in Chicago and although I believe it and have good evidence (it says so on the airport monitor) I fail to know it since my belief is not safe (suppose some of the monitors were not accurate though the one I was looking at is). I now have no evidence that my plane will not crash on the way to Chicago, but if I do know this by the very same evidence (supposing the monitors were all accurate), I do count on Williamson's view as knowing that my plane will not crash on the way to Chicago. So here, it seems to me, we get the worst result. Justification is not closed, and I know things that on the face of it seem that I couldn't. Airport monitors are not as informative in telling us which plane has better prospects for reaching its destination.

<sup>243</sup> Perhaps Williamson would claim that justification is a matter of passing a threshold on one's total evidence. This would allow him to say that in a sense single premise justification closure is preserved on his account. Some of his comments suggest this, e.g. (2000: 9). Yet given his view this would mean that prior probabilities confer justification, which he denies since he claims that the regress of justification ends with knowledge (Ibid). Or even worse, one can be justified in believing  $p$  even though one's total evidence relative to the prior (epistemic) probability of  $p$ , counts as a whole against  $p$ . Moreover, one is still not justified by

One might have (evidential) justification for believing that  $p$ , deduce  $q$  competently from  $p$  while retaining one's justified belief in  $p$ , and have no evidential justification for believing that  $q$ . This will follow for every view that does not commit itself to the claim that justification is a matter purely of passing a threshold, that is, that it demands that justification be tied to having evidence for a given proposition (as Williamson's account is). Moreover, as I have argued, this will follow even if one does not view evidence as probabilistic. Williamson's view, then, is already committed to a substantive separation of evidential justification from knowledge.

Notice that the proposition  $p$  that is justified but not known cannot be taken as evidence on Williamson's account without a breach of the  $E=K$  principle. Interestingly, many theorists take justification closure to be just as intuitive as knowledge closure. For them, Williamson's account will seem problematical. The principle of single-premise justification closure is regarded by epistemologists to be of great prominence.<sup>244</sup> Gettier, in his famous article, proceeds on the assumption that justification closure is a settled issue. In the Williamsonian framework, Gettier's claim is simply false:

[F]or 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  $Q$ . (Gettier, 1963)

Yet, for Williamson there is a crucial difference between what holds for one's justified beliefs and what holds for knowledge. On his account, one's evidence is one's total knowledge ( $E=K$ ), so the probability of anything that is known is 1 (since  $p$  is included in  $K$  – which is all that is known, for any known  $p$ ,  $\Pr(p|E)=\Pr(p|K)=\Pr(p|p\wedge K)=1$  by Standard Conditionalization). The principles governing justification, therefore, diverge significantly from the principles governing knowledge. No matter how many conjuncts one adds in the process of competent deduction, as long as the premises are

---

making a deduction from a justified premise (something that is made explicit in some formulations of closure), but rather by one's other knowledge. It is clear, then, that Williamson does not preserve much of the features that appeal to many epistemologists. (See main text below for some examples.)

Williamson claims that only evidence can justify belief: "One's evidence justifies belief in the hypothesis if and only if one's knowledge justifies that belief." (2000: 9). Similar remarks are made can be found in many places in the book, e.g. (2000: 208). At the very least, then, Williamson needs to modify his view in order preserve single premise justification closure. Multi premise justification closure seems to be a bad idea regardless of whether it is intuitive or not. I suppose that Williamson would not want to defend it.

<sup>244</sup> Besides Gettier (see quote in the main text) and many other theorist, Feldman (1995: 487) thinks that "the idea that no version of this principle is true strikes me, and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years". Many commentators who express sympathy with Feldman's attitude miss the fact that Feldman's claim is about justification closure not knowledge closure (though I suspect that he would make similar remarks about knowledge closure).

known, the conclusion will have the same probability as the premises, namely, 1.

The epistemic probability that Williamson appeals to is by no means standard. Particularly, prior probabilities receive a characterization in *Knowledge and its Limits*, which can, and I think, should, be questioned.<sup>245</sup> But assuming for present purposes that there is no problem with the prior probabilities as such, a consequence of Williamson's knowledge-evidence equation is that since the posterior probability of what is known is 1, the natural reaction to the puzzle – rejection of (7) – is not available. Even theorists who do not question *single premise* knowledge closure would be tempted to reject *multi premise* closure. But, since adding known conjuncts by a deductive process of conjunction introduction will on Williamson's account always leave the probability of the conjunction unscathed ( $\Pr(q)=1$ ), this natural line of reasoning is blocked for Williamson. For him, multi and single premise *knowledge* closure stand or fall together.<sup>246</sup>

So, rather than a rejection of (7), it is not surprising to find that Williamson rejects (8). This he achieves by drawing a distinction between objective chance (henceforth simply *chance*) and epistemic probability (henceforth *probability*), a distinction with quite far reaching consequences. The claim is that although the chance that the conjunction is true is low and is known to be low, its probability can be high, in fact in many cases it will be 1.

### 5.3.3.2. Lottery Propositions

How does Williamson justify the sharp distinction between chance and probability? After all, it was more or less obvious to Lewis that there is a tight connection between the two. Lewis famously claimed that objective chances should equal one's (admissibly) informed credence (Lewis, 1986, 1999).

Williamson's idea (which Hawthorne and Lasonen-Aarnio anticipate) is that what is objectively probable, i.e. chance, need not be represented in terms of close possible worlds. Suppose we represent a future event's chance of taking place as a branching out of a common history. When there is a chance that a different event, such as a quantum blip, will occur, no matter how slim the chance, there is a branch extending into the future from the

---

<sup>245</sup> There seem to me to be problems with Williamson's characterization of prior (epistemic) probabilities. He states that "P" (which is similar to a prior probability function in Bayesian terms) "measures something like the intrinsic plausibility of hypotheses prior to investigation" (2000: 113) which can vary, he says, with context. My suspicion is that there are problems with assuming that a context will determine one unique value for each hypotheses, though Williamson says that "P(p) is taken to be defined for all propositions." (Ibid). This, as well as related issues (such as the Bertrand Paradox), are beyond the scope of this manuscript.

<sup>246</sup> For reasons discussed in Chapter 1, Williamson's theory has the advantage of being supported by the intuition that one can expand one's knowledge to the conclusion of a *modus ponens* inference if one knows the premises.



actual world to one in which the other event is the one that takes place. Suppose we represent the conjunction  $p_1 \wedge \dots \wedge p_n$  in the above argument as a finite set of worlds with a common history up to a time  $t$ . We then have many branches extending from the set to worlds in which one of the events does not take place. Williamson's idea is that branching worlds are not necessarily close, where closeness is the central notion he employs to cash out his safety requirement (Williamson, 2000: 123-130). Since on his account S knows that  $q$  only if S's true belief is safe from error, there can be no close worlds in which  $q$  does not hold no matter how slim the chance is of  $q$  being true. This is how radical a divorce Williamson is advocating between objective chance and epistemic probabilities.<sup>247</sup>

Yet this account of knowledge runs into trouble with knowledge of what we can call *quantum conjunctions* (that is, conjunctions of very many propositions each having very high chance of being true adding up to low chance for the conjunction). To see how, let us slightly modify a case presented by Hawthorne and Lasonen-Aarnio:

Consider extremely unlikely and bizarre 'quantum' events such as the event that a marble I drop tunnels through the whole house and lands on the ground underneath, leaving the matter it penetrates intact. (Hawthorne and Lasonen-Aarnio, forthcoming: 3)

Let us imagine that we have a large yet finite amount of such marbles, say all currently existing marbles, such that on our best theories, it is quite likely (though not definite) that if all are dropped, at least one of them will in fact tunnel through the floor. As a matter of contingent fact, when we do drop them in some future time none of the marbles tunnel. Now the question I want to ask is this: Given that one has all and only information pertaining to the chances, does one know that none of the marbles will tunnel? On our assumption that as a matter of fact none of the marbles will tunnel, does one know that:

- (11) For all existing marbles  $x$ , if  $x$  is dropped to the floor,  $x$  will not tunnel?

Whether he denies knowledge of (11) or allows such knowledge, Williamson's theory, it seems, is in trouble. I will begin with the option of denying knowledge of (11). In Williamsonian terms, if it is not known this must be due to the fact that the basis for knowing (11) is not adequately safe. This

---

<sup>247</sup> To what extent can known chance and probability diverge? Can the known chance be 0 while the known probability is 1? As this question quickly gets entangled with issues that would take me far afield, I will leave it as an open question for further deliberation. My guess is that on Williamson's view knowledge that  $p$  is compatible with  $p$  having 0 chance of being true but whether there is a way to extend this to knowledge of the chances is a complicated questions.

would mean that if we represent (11) as a long conjunction of propositions about dropped marbles tunneling, there is at least one conjunct that is not safely believed. But which? There is no forthcoming answer. It seems implausible that one would not know (11) on the basis of reasoning that every marble has extremely high chances of not tunneling (and assuming that all will in fact not tunnel). Moreover, it is apparently false that there is some marble about which the belief that it will not tunnel is not safe.<sup>248</sup> Nevertheless, Williamson might use this reasoning in the opposite direction concluding that some of the beliefs are not safe since the belief in the conjunction is not safe. The unsafe beliefs are determined by modal reality.<sup>249</sup>

But there are further difficulties with denying knowledge of (11). If (11) is, under the circumstances, unknown, then, it would seem, so must be (12):

(12) If *this* marble is dropped, it will not tunnel through the floor.

Assuming that we know a particular marble M will be dropped over a floor, does one know that the following is true?

(13) If M is dropped over *this* floor, M will not tunnel.

To avoid skepticism about future contingents Williamson must allow knowledge of (13). But if (11)-(12) are not known, it seems hard to explain how (13) could be.<sup>250,251</sup>

---

<sup>248</sup> If conjunctions of the form of (5) and (9) are known, it seems that (11) should be too, by the same sort of reasoning. Suppose we lay out the marbles and form a belief regarding each one (that it will not tunnel) and then add them up into a long conjunction from which (11) trivially follows.

<sup>249</sup> Thanks here to Professor Pagin for noting this possible reply on behalf of Williamson.

<sup>250</sup> I wish to note in this context (but the point has already been made generally) that there is a difference between cases of perception or memory on the one-hand and knowledge of future contingents and knowledge by induction on the other. The former cases can be accounted for by an appeal to a form of epistemic externalism: if the evidence for there being a table in the room is seeing the table in the room, there is no mystery. When the fact itself is part of the evidence it is clear why the probability is 1. (Whether this form of externalism is plausible or defensible is another matter.) In the case of future contingents and induction it is clear that the evidence does not entail the known proposition. Supposing, as Williamson does, that one theory can account for these two sorts of knowledge is, perhaps, the heart of the matter. The epistemic situation seems to be very different when a proposition is entailed by the evidence (in which case assigning it epistemic probability 1 makes sense) and when it is not. There would be no need for induction if the inductive base entailed the known proposition. The idea that  $\Pr(p|e)=1$  when  $e$  does not entail  $p$ , to which Williamson is committed, seems to be the root of many of the problems. There is room then for bifurcating one's account of knowledge so that mathematical, perceptual, and knowledge based on memory is explained Williamson's way, while inductive knowledge, for example, is explained by a more Lewisian conception of evidence. Some remarks by Williamson suggest that he may be more sympathetic to this kind of bifurcation than one might think:

But even disregarding the behavior of epistemic probabilities, it seems very strange to set the knowledge anywhere between (11) and (13), either all are knowable, or none are. Knowing none is skepticism, knowing all means knowing lottery propositions. Or so I will subsequently argue.

I have argued that preserving knowledge of claims about particulars while denying knowledge of related general claims is problematic on Williamson's view. A similar problem arises for knowledge by induction. I observe several ravens and note that they all appear black. Suppose that all ravens are black, and that at some point in the sequence of observations I come to know this (assuming induction is a method for obtaining knowledge). So at this point, I go from having a true belief that the probability of all ravens being black, such that  $\Pr(\text{For all } x \text{ (if } x \text{ is a raven, } x \text{ is black)}) < 1$ , to a state in which the probability equals 1. Although the transition from non-knowledge to knowledge is problematic on any account (Hume famously questions the very justification of induction), there is an added mystery in Williamson's account. My prior conditional probability of all ravens being black on my observing the next raven to be black was less than 1 and it increased steadily as evidence came in. But what is it about actually observing the next raven that changes the probabilities of all ravens being black to 1? Presumably, all theories of inductive knowledge will have to explain how before observing the raven I didn't know that all ravens are black, and now, after observing the relevant raven, I do. But for Williamson there is another difficulty stemming from the shift in probabilities. We are faced with the situation where we know that the proposition arrived at inductively does not follow from the evidence. The prior conditional probability of the hypothesis on the evidence is less than 1 (and we may even know this). Yet by receiving the evidence

---

On the view above of evidence, when they constitute knowledge, they are part of our evidence. Moreover, they may constitute knowledge simply because perceiving counts as a way of knowing; that would fit the role of knowledge as evidence...I certainly did not *perceive* that your ticket did not win. There is no valid argument from the denial of knowledge in the lottery case to its denial in perceptual and other cases in which we ordinarily take ourselves to know. (Williamson, 2000, 252)

<sup>251</sup> Of all the options the knowability of only (13), seems to be the worst. The relevant information about M being dropped should, if anything, lower its probability, since the information verifies its antecedent (making it more probable (and on Williamson's view, raising it to probability 1)). Call the antecedent of (13)  $e$  and its consequent  $p$ . The probability of if  $e$  then  $p$ , is lower given  $e$ , than the unconditional probability of if  $e$  then  $p$ . On a standard Bayesian picture,  $\Pr[(\neg(e \rightarrow p))|e] = \Pr(e \wedge \neg p) / \Pr(e) = \Pr(e \wedge \neg p) / \Pr(e) \geq \Pr(e \wedge \neg p) = \Pr[\neg(e \rightarrow p)]$  and hence,  $\Pr[(e \rightarrow p)|e] \leq \Pr(e \rightarrow p)$ . Given these standard assumptions  $\Pr[(e \rightarrow p)|e] < \Pr(e \rightarrow p)$ . Williamson rejects these assumptions, but given the case above this is to count against this rejection not against the standard assumptions. His rejection depends on counting  $\Pr[(e \rightarrow p)|e]$ , as having probability 1 once  $e$  becomes known.

(which we know does not entail the proposition) we somehow arrive at probability 1 for that proposition.<sup>252</sup>

Leaving aside the issue of induction, the idea that propositions like (13) are not known seems to be entangled with too many problems. It seems, then, that Williamson's theory would incline him to treat universal statements such as (11) in the same way he treats conjunctions of future contingents, namely as cases where, although chances are low, the epistemic probability is 1.

But suppose now that we have a lottery drawing in which 1 of a billion tickets will be drawn. Suppose further that all but one ticket have been sold and, coincidentally, it will be the one unsold ticket that will be the winner of the draw. So, for each of the sold tickets it is true both that its chances of losing are very high and that it will in fact lose. Is the belief that one of these tickets will lose safe? Williamson, like most epistemologists, thinks that lottery propositions are not known.<sup>253</sup> This is required for his explanation of the unassertability of lottery propositions in terms of their unknowability (Williamson, 2000, 224-229).<sup>254</sup> Merely having probabilistic reasons that a losing ticket in a large lottery will lose does not allow me to know, no matter how good my information is about the chances, that the ticket will not win. On Williamson's conception of safety, if belief in a conjunction is not safe,

---

<sup>252</sup> For some, this territory is already familiar - it is a version of the consequence endorsed in *Knowledge and its Limits*, namely the rejection of van Fraassen's Reflection Principle: if  $x$  is one's current rational credence in the truth of a proposition  $A$  at some later time  $t$ , then the current rational probability for  $(A)$  should be  $x$ ;  $\Pr(A|\Pr^t(A)=x)=x$ . (Van Fraassen, 1984). In later writings van Fraassen articulates the principle differently (1995, 1999), but the differences do not concern the current discussion.

In the main text Reflection was not assumed. With respect to a specific case of induction suppose I know I am going to see a raven, call him Max, and as a matter of fact this raven will be the last in a long line of observations resulting in knowledge of  $B$  (=all ravens are black) but I am unaware of this fact (I don't know that I will know). What should my current credence be conditional on Max being black? Answer: on Williamson's (and most everyone else's view except Karl Popper's perhaps) it should be high but less than 1. What about after I view black Max? Answer: on Williamson's view, the probability is now 1. But what happened? Why?

Having probability 1 that all ravens are black, I also can know that there is a chance that they are not all black assuming I know that there is an objective chance of there being (past, present or future) e.g. non-black raven mutations. So again, I know that there are no non-black ravens but I also know that there is a chance that there are (perhaps even a high chance). Read in any way one wants (epistemic or metaphysic), it sounds very odd to say that *I know that there are no non-black ravens and I know that there very well might be*.

<sup>253</sup> "[H]owever unlikely one's ticket was to win the lottery, one did not know that it would not win, even if it did not ... No probability short of 1 turns true belief into knowledge." (Williamson, 2000: 117)

<sup>254</sup> On page 255, Williamson (2000) connects the case of lotteries to the unassertability of the belief that one will not be knocked down by a bus tomorrow. It is hard to see how Williamson would separate this belief from beliefs about the non-occurrence of quantum events. If you don't know that you will not be knocked down by a bus, how can you know that a quantum blip will not happen in this instance?

there must be at least one conjunct belief in which is not safe.<sup>255</sup> But as all conjuncts in this case are on a par, if belief in one is not safe, belief in any isn't. Thus, this would commit Williamson to a substantive distinction between quantum propositions – which are known, and lottery propositions – which are not. But what could be the difference? If I cannot know that a lottery ticket is a loser, how can I know that a quantum blip will not occur, let alone know that the negation of a long disjunction about quantum events is true? If beliefs regarding falling marbles are safe, why not lottery beliefs?

To make the connection even tighter, assume we match each of the lottery tickets to a marble-dropping event (suppose we write the numbers of the tickets on marbles which are then dropped, and the winner is the holder of the ticket whose number is on the marble that tunnels through).<sup>256</sup> It does not seem plausible in this case to say that although I know the marbles will not tunnel, I don't know my ticket is a loser.

There are two possibilities here: either knowledge is lost by when a connection is made between the marbles and the lottery, or knowledge of the lottery propositions is gained by this connection. If it is gained, then one can know by closure that the unsold ticket will win. Loss of knowledge is equally dubious. Why would the fact that the quantum events are used as lottery mechanisms make them unknowable, if other quantum conjunctions are known?

In general, it is hard to see why the world in which I win the lottery should be regarded more similar to those in which I lose, than the worlds in which a marble tunnels is to those in which none tunnel. The lack of clarity regarding the similarity (or closeness) relation at play in his account is related to a further lacuna in Williamson's presentation. Not often does Wil-

---

<sup>255</sup> Advocating what he calls the ordinary conception of safety, he claims:

Suppose that I am not safe from being shot. On the ordinary conception, it follows that there is someone  $x$  such that I am not safe from being shot by  $x$  (assume that if I am shot, I am shot by someone). On the high chance conception of safety, that is a *non sequitur*. For each individual  $x$ , the chance of my being shot by  $x$  may be low enough for me to count as safe from being shot by  $x$ , even though the chance of my being shot by someone or other may be too high for me to count as safe from being shot. (Williamson, forthcoming: 23-4)

Besides the problems as outlined in the main text for Williamson's theory that this safety conception generates, I agree here with Maria Lasonen-Aarnio that this is not the ordinary conception of safety. Let me take this opportunity to thank Maria Lasonen-Aarnio for helpful conversations concerning this as well as many other matters related to closure.

<sup>256</sup> How is this possible? After all, we might have more than one tunneling. One answer is that we might suppose that we have one possible winner that is selected by the *first* marble to tunnel. To equal the chances the first will be the one that appears first on a list the order of which is hidden from the lottery players. This way each ticket has the same chance of winning and there is only one possible winner. One who knows that no marble will tunnel will be in a position to know that no one will win the lottery. This is Professor Hawthorne's suggestion. I am indebted to him for many helpful conversations on this and related topics.

Williamson specify concrete instances of what by his lights would amount to knowledge. Can we know things by induction, or is the scenario in which they fail to be true too similar for such propositions to ever have probability 1? What can we know about the future? If he is not to slide too far on the way to skepticism, Williamson must allow that at least some knowledge of these sorts is possible. But then what could be the constraints on the similarity relations such that we get only the “good” cases and none of the “bad”?

A simple statement of the challenge is this: Are lottery propositions known or not? If they are, this would create problems for Williamson’s thesis that knowledge is the norm of assertion (Williamson, 2000: 224-229) and commit Williamson to what is widely considered an unfavorable position (together with multi-premise closure this would mean that one would know beforehand which ticket is a winner if one made the right choice of deduction). If lottery propositions are not known, what is the relevant difference between them and quantum propositions? Specifically, if the lottery mechanism is just the quantum events, how can the latter be known while the former are not?

### 5.3.3.3. The Principal Principle and Practical Deliberation

The divorce of epistemic probability from chance is intuitively problematic. The following is one way to give this intuition some substance. Since the chances are known, in place of Williamson’s (9), we might just as well have:

(9’) One knows that the objective chance of  $p_1 \wedge \dots \wedge p_n$  is low.

Is knowledge that the chance that some proposition is true is extremely low compatible with knowledge of that proposition? The answer to this question depends on the validity of a weakened version of Williamson’s (8):

(8’) If one *knows* that the objective chance of a proposition  $q$  is low, one does not know  $q$ .

Williamson must reject (8’). Yet its rejection entails the truth of Moore-type future sentences of the form (for ease of exposition I use a higher order form):

(14) I know that (the chance that  $q$  is true is low, but it will happen).<sup>257</sup>

Given Williamson’s knowledge account of assertion, the following instance of (14) is assertable: *the chance that my book contains no mistakes is very low, but it doesn’t!*. Strictly speaking, there is no contradiction in (14)

---

<sup>257</sup> Parentheses here are for avoiding scope ambiguity.

or in any of its instances, just as there is no contradiction (at least no trivial contradiction) in any of the Moore-type sentences (many Moore sentences are true). With a further seemingly plausible principle, we can derive other variants of Moore sentences from (14) that sound equally odd.

- (15) If S knows that the objective chance of a future event is very low, then S knows that the future event might not take place.<sup>258</sup>

From (15) we can derive:

- (16) S knows that ( $q$  might not be true but it is).

Note how odd that sounds in the case of future contingents: “I know it might not rain tomorrow, but I know it will.” Obviously Williamson would prefer to regard (16) as unassertable. But since he is committed to the assertability of (14) this would mean he must take (15) to be false. The point of raising this issue (aside from the difficulties associated with rejecting (15) and (8’)), is to shed more light on the radical rift Williamson is imposing between chances and epistemic probabilities.

Beyond intuition, however, there is a theoretical strain here. The sharp split between chance and epistemic probability conflicts – if not in letter, certainly in spirit – with the central idea motivating Lewis’s Principal Principle which says:

Take some particular time – I’ll call it “the present”, but in fact it could be any time. Let  $C$  be a rational credence function for someone whose evidence is limited to the past and present – that is, for anyone who doesn’t have any access to some very remarkable channels of information. Let  $P$  be the function that gives the present chances of all propositions. Let  $A$  be any proposition. Let  $E$  be any proposition that satisfies two conditions. First, it specifies the present chance of  $A$ , in accordance with the function  $P$ . Second, it contains no “inadmissible” evidence about future history; that is, it does not give any information about how chance events in the present and future will turn out. ...Then the Principal Principle is the equation:  $C(A|E)=P(A)$  (Lewis, 1999: 238)

---

<sup>258</sup> It’s hard to see how (15) could be false. Suppose I know that the chance for rain tomorrow is very low. Does it not seem adequate to say, at least, that I know that it might not rain tomorrow?

Compare this claim to Hawthorne’s *Objective Chance Principle*: “If at  $t$ , S knows that there is a nonzero objective chance that  $p$  at  $t$  (where  $p$  supervenes on the intrinsic facts about the future relative to  $t$ ), then, at  $t$ , there is nonzero epistemic probability for S that  $p$ .” (Hawthorne 2004a: 92). He then notes that if the principle is incorrect there should be an epistemic reading of the conclusion “according to which the inference is faulty. But no such reading occurs to us.” (Hawthorne, *Ibid.*) The inference I am appealing to here is weaker since it only pertains to knowledge of high chance, but I agree with Hawthorne that it should hold for low chances as well.

According to this principle one's credence in conjunctions with low chances should be just as low. Yet Williamson – along with others that appeal to the same defense of epistemic closure – is committed to the claim that one knows them. To know a proposition one must, presumably, believe it, which means that one must assign the proposition sufficiently high credence. It seems then that Williamson's desired conclusion requires abandoning the Principal Principle.

Williamson's response is to preserve the principle by allowing updating one's credences on evidence that Lewis regards "inadmissible."<sup>259</sup> Specifically, by conditionalizing on the future contingents comprising the conjunction, which one is assumed to know, one's credence in the conjunction will be 1. Since for Lewis future contingents do not count as evidence such knowledge is inadmissible and is therefore not part of the formulation of the Principal Principle, which, as Williamson says, "is logically neutral as to the results of conditionalizing on inadmissible evidence, despite the forbidding connotations of the word "inadmissible"."<sup>260</sup> For Williamson, all knowledge counts as evidence, including knowledge of the future. So one can conditionalize on this knowledge and update the credence assignments in accordance with the epistemic consequences of closure and the Principal Principle remains unviolated.<sup>261</sup>

To the extent that Williamson's idea of Standard Conditionalizing on all knowledge succeeds, it is a technical victory at best. Surely, even if Williamson manages to avoid violating the letter of Lewis's principle, he still undermines its spirit. The rationale behind the Principal Principle is that one should apportion one's credence in a proposition to what one (rationally) believes are the chances that that proposition is true.

...we have some very firm and definite opinions concerning reasonable credence about chance. These opinions seem to me to afford the best grip we have on the concept of chance. Indeed, I am led to wonder whether anyone *but* a subjectivist is in a position to understand objective chance! (Lewis, 1986: 84)

Thus, although Lewis's formulation of the Principal Principle is silent regarding inadmissible evidence, it is clear that its point is to articulate a tight connection between credence and chance. Williamson's position runs counter to this idea.

---

<sup>259</sup> Williamson, forthcoming: 16.

<sup>260</sup> Williamson, forthcoming: 16.

<sup>261</sup> I must admit that I am not sure how exactly to understand the relation between rational credence and epistemic probability on Williamson's view. It seems to me that he is committed to the idea that there are three kinds of probability: epistemic probability, chance, and credence. His presentation, however, makes it look as if there are only two: chance and epistemic probability.



To make the point more explicit: in his reply to Hawthorne and Lasonen-Aarnio, Williamson goes to great lengths to show that his view does not commit him to any implausible principle. But there is at least one highly plausible principle he seems to be forced to reject, call it the Weak Low Chance Principle:

(WLC) If in  $w$ , at time  $t$ ,  $S$  knows that  $p$  has a low chance of being true,  $S$  does not know  $p$  in  $t$  at  $w$ .<sup>262</sup>

Given Williamson's divorce of epistemic probability from chance, he cannot endorse (WLC). According to the safety theory of knowledge, one knows conjunctions of future contingents, for example, even when one *knows* their chances of being true are very slim.<sup>263</sup>

To see just how problematic this is, consider the practical consequences of this commitment. Suppose the truth of some long conjunction of propositions is of crucial importance for you – if it is true you must  $\phi$  and if false you must not. Now suppose you know each of the conjuncts is highly probable and believe it to be true. Assuming that this belief is safely true, according to Williamson's theory you know the conjunction. But you also know that there is very high chance that the conjunction is false. Should you, or should you not  $\phi$ ? Assuming knowledge is a rational basis for action, that it warrants taking something as a premise in practical deliberation (Hawthorne, 2004a, 29), it seems that Williamson's theory would entail that you ought to  $\phi$ , despite the fact that you know there is very high chance that the conjunction is false. The reason for this, according to the theory is that although the chance that the conjunction is true is (and is known to be) low, since it is true and since you know this, you should act on this knowledge. But one can always turn the table on this argument. If you know that a proposition has high chance of being false you ought not act on it. After all, this too is knowledge. And acting on it is what the Principal Principle seems to suggest, again, if not by strict letter, than by spirit. Viewed from this perspective, a theory of knowledge that has consequences inconsistent with WLC is so much the worse for it, not the better. But even if it is discarded, we still face

---

<sup>262</sup> Compare this principle to Hawthorne's and Lasonen-Aarnio's Low Chance principle (forthcoming, 5) modified by Williamson (forthcoming: 17). I note that I am not advancing this as the most informative connection between chance and epistemic possibility but rather as one we could feel sure of. Thanks here to Professor Hawthorne.

<sup>263</sup> I think that the situation would be bad enough for Williamson's theory even if one's inductive general evidence counted in favor of the chance being very slim and yet one did not take this evidence into account. We might, imagine that we have a quantum physicist who looks at the chances as apposed to an epistemologist who looks at the events. The point is that these two might be the same rational person, while in the case of past events we would have to have two people since the evidence set could not be the same. The difference is that regarding the past the evidence disqualifies the inductive evidence being applied for the same rational individual. I am indebted here to Professor Hawthorne.

the problem of how one ought rationally to act when one has two incompatible directives stemming from one's knowledge. The point is not merely that we might not know what to take as a premise in practical deliberation because we don't know that we know. The point is that even if we do know that we know, we don't know what to take as our premise.<sup>264</sup>

At this stage I think it would be fair to say that the problems associated with taking knowledge as having probability 1 – the idea that one can standard conditionalize on all known propositions – are formidable. The probability 1 account does address the challenge set by the open knowledge account by treating knowledge as if it were always based on conclusive evidence. This in turn leads to the inability to separate knowledge closure from multi-premise knowledge closure and leads in cases where the evidence clearly does not entail the known propositions to substantial theoretical maneuvering that leads to further difficulties and seem to conflict with common sense.

#### 5.3.3.4. Fallibilism

The problems we saw Williamson's theory of knowledge faces – the dilemma regarding lottery propositions and the relation between chance and credence – both seem to arise from the same feature of this theory, namely the attempt to treat human knowledge as infallible in all domains. While the precise formulation of epistemic fallibilism is a matter of contention, most epistemologists agree that human knowledge is, in some sense, fallible. Since the conception of knowledge as safely true belief has the consequence that what is known has an epistemic probability of 1, it is hard to see how any notion of fallibility can apply to knowledge under this conception (specifically if evidence is identified with knowledge). Consider some of the leading formulations.

“Fallibilism is the doctrine that someone can know that  $p$ , even though their evidence for  $p$  is logically consistent with the truth of  $\text{not-}p$ ” (Stanley, 2005, 127). For Williamson, remember, one's knowledge is one's evidence, so surely fallibilism in this formulation by Jason Stanley is not consistent with it. If you know  $p$ , then  $p$  is now part of your evidence and therefore, the evidence you have is inconsistent with the falsity of  $p$ .<sup>265</sup> If one knows that  $p$ , one's evidence is inconsistent with the falsity of  $p$ . Stewart Cohen's formulation is even more directly at odds with Williamson's theory: “a fallibilist

---

<sup>264</sup> Here as in many other places I am indebted to Assaf Sharon.

<sup>265</sup> Perhaps Stanley (who is sympathetic to the idea that knowledge has probability 1 – See Hawthorne and Stanley, forthcoming, 10) could change the formulation thus: Fallibilism is the doctrine that someone can know that  $p$ , even though the evidence one *had* for  $p$  and by which one came to know that  $p$ , is logically consistent with the truth of *not-}p*. If this is true, it will cause other trouble for the view. It would be better if he were to decide on whether he prefers fallibilism or the Williamsonian view.

theory allows that S can know  $q$  on the basis of  $r$  where  $r$  only makes  $q$  probable” (Cohen 1988: 91). Clearly, according to Williamson, this is false. Jim Pryor says: “a fallibilist is someone who believes that we can have knowledge on the basis of defeasible justification, justification that does not guarantee that our beliefs are correct” (Pryor 2000: 518) The notion of defeasibility employed here requires clarification. Nevertheless, it is clear that under any plausible theory, justification will be defeasible in the sense that future evidence can undermine it. Anything else would be an objectionable form of dogmatism.<sup>266</sup> The defeasibility associated with fallibilism must be something to do with the inconclusive nature of that on which ordinary knowledge is based. This, presumably, is what Pryor means when he speaks of “justification that does not guarantee that our beliefs are correct.” Feldman articulated this idea more explicitly. Fallibilism, he says, is the view that “it is possible for S to know that  $p$  even if S does not have logically conclusive evidence to justify believing that  $p$ ” (Feldman 1981: 266). As he explains, this amounts to the claim that knowledge can be had based on less than deductive inference, that one’s evidence need not entail what is known. But if knowledge is safely true belief, belief that has epistemic probability 1 on one’s evidence, then it *is* guaranteed – evidence is conclusive and entails what is known (at least epistemically).

There is, perhaps one kind of fallibilism that is compatible with Williamson’s safety theory of knowledge. Since epistemic probabilities are for Williamson divorced from objective chances, it is consistent with his theory that one can know things which have less than perfect chance of being true. Indeed, as we have seen, this is a desiderata of his theory. It may be claimed therefore that epistemic fallibilism consists in the fact that we can know propositions which – objectively speaking – have some chance of being false. For one thing, it is not clear how this conception can be applied to propositions not about the future (which, if true, presumably have objective probability of 1). The important thing to notice about this proposal, however, is that it allays the epistemic bite of fallibilism. As the attempts to define fallibilism all indicate, the idea that knowledge is fallible is supposed to capture something about the relation between knowledge and the evidence or justification on which it is (or can be) based. This is lost in the kind of fallibilism compatible with Williamson’s theory. With this definition all that the fallibility of knowledge comes to is the uninteresting claim that we can know truths that have some chance of being false, although they aren’t, and our evidence guarantees that they aren’t. If this is as robust a notion of fallibilism as Williamson can endorse then those convinced by the idea that

---

<sup>266</sup> The kind of dogmatism that is dealt with quite easily in chapter 3 by an open knowledge account, not dogmatism in Pryor’s (2000) sense.

knowledge is fallible should be left unsatisfied and the problems discussed in previous sections remain intact.<sup>267</sup>

Moreover, consider the consequences of this infallibilism for belief revision. Surely there are psychological facts about a subject that might undermine knowledge of  $p$  such as ceasing to believe it. But loss of knowledge and change of belief is also sometimes rationally required, specifically, when proper counter-evidence presents itself. On Williamson's account, once something is known such change is not rationally mandated. This is because one can always take one's knowledge  $p$  as evidence ruling out any evidence to the contrary. To illustrate this point consider Williamson's own example:

I put exactly one red ball and one black ball into an empty bag, and will make draws with replacement. Let  $h$  be the proposition that I put a black ball into the bag, and  $e$  the proposition that the first ten thousand draws are all red. I know  $h$  by standard combination of perception and memory, because I saw that the ball was black and I put it into the bag a moment ago. Nevertheless, if after ten thousand draws I learn  $e$ , I shall have ceased to know  $h$ , because the evidence which I shall then have will make it too likely that I was somehow confused about the colours of the balls. (Williamson 2000, 205)

This surely seems to be the rational way to go. But  $e$  is never simply given. After ten thousand draws one is faced with two inconsistent pieces of information,  $h$  and  $e$ . If  $h$  is known and therefore has probability 1, it would be just as reasonable to question one's memory, which, presumably, is the basis of one's belief in  $e$ . One can always conditionalize on one's knowledge (evidence) that there is a black ball in the bag and conclude that one is confused about  $e$  ( $\Pr(e|h)=1/2^{10,000}$ ), not about  $h$ . This is a further sense in which knowledge must be fallible in a more substantial way than Williamson's view allows.

In sum, then, and relating back to Chapter 1, a position that considers knowledge to be incompatible with any rational measure of epistemic doubt can support the idea that basic inference patterns such as *modus ponens* are closed. This is a significant advantage that makes room for a true extension of knowledge through valid inference. It blocks the accumulation of epistemic doubt and thus by knowing the premises of a *modus ponens* inference one will know, by proper inference, its conclusion. The idea that inference is a general way to extend one's knowledge, perhaps mistakenly, is considered by many to be a central motivation for the view that knowledge as closed. However, it is this feature that leads to considerable difficulties precisely when we have independent means, e.g. by our knowledge of chances, to establish a gap between evidence and knowledge. This gap, which the infallibilist denies, runs into trouble with induction, future contingencies, chances, practical reasoning, justification closure, lottery propositions, quan-

---

<sup>267</sup> Here too thanks to Assaf Sharon.

tum events, and the like. The trouble for the infallibilist is more vivid when items of knowledge are conjoined to form would be knowledge that runs into conflict with our knowledge of chances. Thus, for instance, when we connect knowledge of a long conjunction to a lottery setting, it becomes evident that an unreasonable expansion of knowledge is being postulated that does not take into account the accumulation of epistemic doubt – the gap between our evidence and our knowledge.

#### 5.3.4. Knowledge as Evidence: Conclusion

I have claimed that avoiding the Open Knowledge Argument can take the form of denying that (NED) is applicable with regard to the pertinent cases. New knowledge of a proposition  $p$  changes the evidential situation so that now there is good evidence for a proposition that is known to follow from  $p$ , namely, the new evidence that consists of  $p$  itself. One way of spelling out this suggestion does not work for the purposes of closed knowledge (Jeffrey Conditionalization) another way does avoid the argument. The latter is Williamson's view.

To be sure, I have said nothing to conclusively refute Williamson's theory of knowledge. Choosing between competing theories in such matters is more often a matter of balancing costs against benefits than of conclusive refutation and proof. The suggestion that knowledge is open has intuitive costs of its own. Yet in trying to avoid it by taking the proposition known as one's evidence and treating it like certain knowledge, I have argued that one needs to redefine epistemic probability; face preface and lottery paradox type challenges with regard to knowledge; rejecting most, if not all, of the more plausible characterizations of fallibilism; reject van Frassen's reflection principle and face a formidable problem of knowledge by induction; reject Lewis' Principal Principle (at least in spirit); face problems of practical reasoning (that several adherents of this view would find troubling); reject justification closure; face one problem if it is accepted that lottery propositions are known; and another if they are unknown. Surely, even Williamson would admit that these are formidable costs. The question is whether these costs are worth paying in order to avoid the consequences of open knowledge or for any other reason short of necessity. To my mind, the most problematic feature of Williamson's account (or any account which takes knowledge to have probability 1) is the idea which is exemplified in the preface paradox for knowledge, e.g., that one can know that a book contains no mistakes and yet know that it is highly probable that this book does contain mistakes. It is the idea that knowledge always has probability 1 that is at the root of this consequence, yet I think it is the most plausible way to avoid the open knowledge conclusion. In this respect, one might be tempted to think of the whole situa-

tion as a paradox rather than a challenge.<sup>268</sup> Yet in light of the simple and elegant way in which knowledge openness can deal with many of the current challenges of contemporary epistemology, open knowledge seems to be in a reasonably good position.

## 5.4. Evidential Knowledge

The attempt to provide the Open Knowledge Argument with grounds for viewing it as sound, are in effect complete. My defense of open knowledge against Hawthorne's arguments has given rise to reasons for thinking that knowledge is open, namely, since evidence is not closed under known entailment. Given even a moderate dependence of knowledge on evidence, I have claimed, this conclusion leads to the idea that knowledge too is not closed under known entailment, at least knowledge of a disputed domain. But where exactly does this argument leave the debate on epistemic closure? Advocates of closure might reject (NED) and argue that although the evidential principles we have considered are not valid, their epistemic counterparts are (e.g. (EAD) is invalid but (AD) is). Of course, one does not merely need to claim this, one has to show how this claim might be made to work. The previous sections of this chapter criticize suggestions of how to defend closure in light of the Open Knowledge Argument (which is not to say that there is no other way I am unaware of that may work). In this section – 5.3 - I provide further reasons for thinking such suggestions are not as appealing as the open knowledge view.

I argue for this in two steps. First, I underscore the theoretical advantages of knowledge openness. The benefits of epistemic openness reach far beyond the foregoing evidential considerations, which I take to be the primary basis for epistemic openness, and bear on many of the central issues of contemporary epistemology. Second, I show how this position can accommodate one of the main ideas behind closure, i.e. the idea that a belief formed on the basis of competent inference from a justified belief is itself justified (i.e. single-premise justification closure). I do this by appealing to a distinction between two types of justification, one of which is closed but does not facilitate knowledge, while the other is knowledge-conducive but not closed.

### 5.4.1. The Benefits of Epistemic Openness

The openness of evidence, as I have said, provides the advocate of epistemic openness with a reasonable positive account for her position and a defense

---

<sup>268</sup> Professor Ruth Weintraub suggested this to me in conversation.

against attacks of the sort mounted by Hawthorne (together with an explanation of why his argument seems so compelling). There is also the argument for open knowledge. But, as in urban planning, there are other, environmental reasons for preferring openness to closure. In the case of knowledge the relevant environment consists of a host of epistemological problems that have seemed quite resilient to proposed solutions, that are easily solved, or rather dissolved, once epistemic closure is denied.

Skepticisms of various sorts rely on the validity of closure. These are not merely Cartesian skeptics, i.e. skeptics undermining entire realms of knowledge, but also Mundane and Live skeptics. Skeptics of these brands, as I claimed in Chapter 2, argue from the admitted lack of knowledge of an implied proposition to the dismissal of ordinary knowledge claims. It is easy to see that this maneuver cannot get off the ground without closure. Epistemic Dogmatism is the idea that, since  $p$  implies that evidence counter to  $p$  is misleading, knowing that  $p$  one can also know by mere reflection that any counter-evidence is misleading and thereby be – absurdly – warranted in disregarding evidence counting against what one believes. Again, if closure is denied, the odd knowledge claim is avoided. Similar considerations apply to lottery propositions and the Lottery paradox for knowledge.<sup>269</sup> Knowing mundane propositions about the future does not commit one to knowledge that one's lottery ticket is a loser or that one will not be one of the unfortunate victims of sudden heart attacks etc. Easy Knowledge (Cohen 2002,

---

<sup>269</sup> As presented in Hawthorne (2004a) .I have not said much in this manuscript about the Lottery Paradox for knowledge. One reason is that Hawthorne (2004a) has a presentation that is better than any I can offer. I have only one additional thing to say. Lotteries seem to pose a difficulty with regard to rationality as they do with knowledge and closure. Suppose that I receive two lottery tickets as a present one for the Stockholm lottery that has odds of 1:1,000,000 of winning, and the New York State Lottery with odds 1:100,000,000. I place both tickets on my desk. One week later I read in a local Stockholm newspaper that a ticket with a different number than my ticket has won. It seems rational for me to throw away the Stockholm ticket even though the probability that there was a mistake in the paper and that I have won, may be greater than the probability that I won the New York State Lottery. How are we to account for this? It seems irrational to throw away the New York State Lottery ticket before receiving news that I have lost yet with better odds of winning it seems rational for me to throw away the Stockholm Lottery ticket.

This is not a mere peripheral problem since, as Avishai Margalit noted in conversation, a similar problem arises with regard to legal evidence. Defendants are sometimes convicted on the basis of eyewitness reports even though the probability of mistake is greater than when using statistical evidence. Suppose we have a defendant accused (among other things, perhaps,) of participating in a prison-riot that all but a few prisoners participated in. The evidence that few of the prisoners did not participate might be as conclusive as we wish yet few courts would convict the defendant of this offense on this evidence. Yet, if a prison guard claims he saw the defendant participating in the prison-riot, we would expect a conviction even if the probability of convicting the innocent on this evidence is greater than on the previous statistical evidence.

I am not convinced that the cases are exactly similar. In the prison case, if we had convicted all the prisoners we would be sure to have convicted some innocent convicts (innocent of rioting, that is). In the case of an eyewitness we have no such guarantee. And if the case is such that we do, I think our expectation that there will be convictions will drop considerably.

2005) of the reliability of one's faculties is also blocked once closure is discarded. The correlation between what I believe is true and the deliverances of my faculties do not permit the inference that my faculties are reliable. All these problems, as well as others I will not rehearse here,<sup>270</sup> do not so much as arise once closure is given up and they are explained by the failure of evidence to support propositions that one knows to follow from what one knows (as well as accounting for the hold these problems have on us).

It should be noted that the evidence account of why closure fails is readily applicable to each of these cases. Perception of<sup>271</sup> my hand provides me with evidence that I have a hand but not that I am not a brain in a vat misled to believe that I have a hand. Evidence for  $p$  can support my belief that  $p$  is true, but does not indicate that evidence against  $p$  is misleading. My promise to meet you at the movies does not make it more probable that I will not fall on the way and break my leg, or that my folks will not show up for a surprise visit. Equally, experiencing perception of red patches makes it more likely that there are red patches before me, but not that my perceptual faculties are functioning well. The point is that a single account that both explains and dissolves a wide range of what were previously considered detached problems, is surely very attractive and deserving of serious attention.

#### 5.4.2. Denying Closure: Not as Bad as You Think

Giving up epistemic closure surely has its costs. Strong intuitions support the principle of closure, not least among them is the idea that inference provides justified beliefs. Insofar as knowledge requires justification (whether it is understood as internalist justification or as externalist reliability) a belief formed via proper inference should be a candidate for knowledge. Regardless of whether there's evidence, it would seem, anything properly inferred

---

<sup>270</sup> It also avoids the problems associated with the semantic externalism/first-person access (see Brown 2004: 239-42). This is perhaps the place to note that the watch example represents a type of case not covered by the standard defense of closure that relies on denying the assumption of a known premise as soon as a counter instance is proposed: even those who think warrant for believing an animal is not a disguised mule is a necessary precondition for knowing that it is a zebra, will, I presume, agree that to know that it is three o'clock one does not need to be already know that even if the watch is broken it is showing the right time.

Another paradox that is avoided is the Knower (Kaplan and Montague, 1965). I am not sure whether closure is the way out, but it should be considered a serious contender given the other options. See the exchange between Cross (2001) and Uzquiano (2004). Of course open knowledge also avoids difficulties having to do with conjunction introduction (multi-premise closure), such as the preface paradox for knowledge. For more on the preface, see 5.3.3.1. Other cases which have closure as one of the premises I am not so optimistic about. I do not think, for instance, that the key for resolving the surprise examination paradox is rooted in epistemic closure. I am also not optimistic about resolving Kripke's belief puzzle (1979) that appeals to rational belief closure (though Frances (1999) has claimed that it does).

<sup>271</sup> Assuming that "perception of" is non-factive.



from a known belief is justified and hence can be known. In this section I claim that epistemic openness need not conflict with this idea. By appealing to a distinction between warrant for belief and knowledge-promoting justification, epistemic closure can be denied without thereby undermining the justificatory capacity of inference. The issues pertaining to epistemic justification are copious and convoluted and surely cannot be exhausted here. My aim is merely to tease out some intuitions and common conceptions about justification that can go some way towards clarifying and supporting the distinction between justification for belief – doxastic justification, and knowledge-promoting justification – epistemic justification. Given this distinction, epistemic openness will seem not as alarming as it may appear.

It is widely accepted that if one has justification for  $p$ , but forms the belief that  $p$  not on the basis of justification, one does not know  $p$ . Gettier cases demonstrate an additional feature with regard to knowledge. Russell's example, for instance, of forming a correct belief regarding the time of day on the basis of a faulty clock illustrates that even if the belief is based on one's justification – and is thus justified – still, it might not amount to knowledge.<sup>272</sup> Some philosophers believe that different types of belief require different types of justification. Knowledge of a mathematical theorem's truth, according to these philosophers, requires knowing its proof. While believing it, say, on the basis of testimony may be sufficient to justify this belief, it cannot provide sufficient grounds for knowledge. But even those who dispute such a distinction between types of beliefs tend to agree that reasons to ascribe high probability do not always promote knowledge. Presumably, one knows it is highly probable that a lottery ticket will lose, and is thus warranted in believing it will lose, yet we are not inclined to say that one *knows* the ticket is a loser. A belief that is (known to be) highly probable is surely justified. But if justification in the sense of reason-to-ascribe-high-probability could promote knowledge, then lottery propositions would be known. Or take the example of believing there is a sheep in the field based on seeing a sheep-shaped rock behind which a sheep happens to be grazing. Perception of a sheep-shaped object in the field surely raises the probability that a sheep is in the field, thus making it reasonable to believe it, and yet under the circumstances one would not be said to know as much. Knowing my financial state, it would be reasonable of me to believe that, despite my life-long dream, I will not buy a classic estate in Provence in the next year. But if my long-lost uncle has just tracked me down and is planning to bequeath me a large sum of money, my belief does not amount to knowledge, even if eventually I do not receive the money.<sup>273</sup>

---

<sup>272</sup> Russell (1948: 154). Russell mentions similar Gettieresque worries about knowledge much earlier, see his (1912: 132).

<sup>273</sup> Hawthorne ascribes a similar example to Joseph Raz (Hawthorne 2004a: 65) and Harman (1973) presents similar examples.

To gain some clarity, we may distinguish between different conceptions of justification here. One can be warranted in believing  $q$  on the grounds, for example, that this is what one must believe in order to retain coherence amongst one's beliefs. Thus, we may have reason to believe that there are external objects if we are to maintain coherence without revising a wealth of our beliefs. In this sense one can be said to be warranted *in believing  $q$* . But does this entail that  $q$  is justified? That there is a consideration telling in favor of  $q$ 's truth? Not necessarily. The fact that coherence amongst our beliefs requires us to believe that the external world is real does not constitute a reason telling in favor of it being real. Yet, it does warrant us in believing the external world is real.<sup>274</sup> Let us call this kind of justification *warrant for believing*.<sup>275</sup> The notion of warrant for believing is an evaluative notion pertaining to epistemic agents. A second notion of justification – *warrant for a belief* – pertains to beliefs.<sup>276</sup> A belief is justified when, for instance, it is supported by the evidence or has been formed in the right way (by reliable method or whatever). Thus, if one believes something on the basis of a false belief, one can be warranted in believing it while the belief itself is unwarranted. Warrant for believing surely does not suffice for knowledge. Even those who think justification is a necessary condition for knowledge will agree that being warranted in believing something does not always guarantee knowledge, even if the belief is true. As lottery and other cases show, even warranted true beliefs may not amount to knowledge.<sup>277</sup>

What about beliefs justified by inference? Surely, the mere fact that a belief is the product of a valid inference does not suffice for it to count as knowledge. The inference has to be from a true and *justified* belief. But then if the justification of the original belief is evidential, and evidence is not closed under implication, what reason is there to think that the inferred belief is evidentially justified?<sup>278</sup> Inference, it seems, is not an independent source of justification, if anything, it *transmits* justification from beliefs to inferred

---

<sup>274</sup> Even coherentists may subscribe to the idea that coherence alone does not suffice for knowledge conducive justification. In other words, although coherentists commit to coherence being a necessary condition for knowledge, they need not advocate coherence as a sufficient condition. Thanks to Mikael Janvid here.

<sup>275</sup> This is not to say that there is no consideration that tells in favor of external objects. What it does mean is that in order to be justified in believing that these exist, one does not have to base one's belief on such a consideration.

<sup>276</sup> For a similar distinction see Engel (1992).

<sup>277</sup> The distinction does not relate to the degree of justification. Few of our beliefs are as justified, probabilistically speaking, as our beliefs in lottery propositions.

<sup>278</sup> I am assuming that the KK principle (that necessarily if one knows that  $p$ , then one is in a position to know that one knows that  $p$ ) is false and that the JJ principle cannot be assumed (in a sense it is this principle that is at issue). For a related discussion about necessary conditions of knowledge and their relation to closure see Brueckner (1985, 2004), Murphy (2006) and Warfield (2004).

beliefs.<sup>279</sup> But, as the open evidence argument has shown, at least one type of justification, namely evidential justification, does not transmit across inference. Therefore, to insist that inferring a proposition supplies one with knowledge-promoting justification for its truth is, in the present context, to beg the question.

Still, it might be urged, if we know that  $p$  is highly probable and that  $q$  is implied by  $p$ , we know that  $q$  is highly probable and can therefore know  $q$ . But the fact that one may have reason to assign a high degree of probability to some proposition's truth, does not suffice for knowing it. We have just seen that other reasons to ascribe high probability do not always amount to knowledge: I know the probability that my lottery ticket is not the winning ticket is extremely high, and yet, it is widely acknowledged that I do not know that my ticket will lose.<sup>280</sup> Believing with high credence that I have a hand, I should assign high probability to there being external objects, yet, pace Moore, I do not know this.<sup>281</sup> My uncle's generosity eliminates my justified belief about my prospects of acquiring French real estate from counting as knowledge. Indeed the claim is hardly novel. The point is that these observations support the distinction between the types of justification. Doxastic justification (of both brands – warrant for believing and warrant for a belief) may not be enough to make true belief count as knowledge.

But if doxastic justification is not enough for knowledge what else is needed? The following is one proposal that relates to the reflections on evidence and a probabilistic conception of justification. The relation of evidential support, we might say, has at least three dimensions. The *degree* of support, i.e. the conditional probability that a proposition is true given the evidence, is just one. A second dimension can be called the *direction* of support, i.e. whether the evidence raises or lowers the probability; and a third dimension is the *magnitude* by which the evidence changes (raises or lowers) the proposition's rational probability assignment.<sup>282</sup> In each of the cases of

---

<sup>279</sup> This is shown by the following consideration. Suppose S has justification for  $p$ . Forming the justified belief that  $p$ , S then infers from it that  $p$  is true. Surely her inferred belief does not enjoy a greater degree of justification than her original belief. Inference does not itself provide justification; rather it is supposed to be a mechanism of transmitting justification from premises to conclusion.

<sup>280</sup> If knowing the odds alone would suffice for knowing that a ticket will lose buying a lottery ticket would be not only irrational but downright crazy. Moreover, absurdities of the form "I know that  $p$ , but I also know that  $p$  might be false" would be assertable.

<sup>281</sup> Even those inclined to grant knowledge of this sort may not wish to base their position solely on its high probability assignment as entailed by the high probabilities assigned to ordinary truths such as that I have hands. Also, they must at least recognize the intuitive cost of their position. It is far from clear, in any case, how to account for the probability of such propositions. How probable is it that there are external objects? If asked how much to bet, one would be warranted in betting the farm. If one loses, one would not really be losing the farm.

<sup>282</sup> There are other dimensions as well. James Joyce (2005) makes a very compelling case for what he calls the "weight" of evidence. Roughly, the idea is that if one believes that a coin is fair on the basis of 100 coin flips, one would rationally change one's mind if the next 100 flips were highly biased towards Heads. Not so if one based one's belief on 100,000 coin flips. Yet

closure failure I have canvassed, the evidence functions properly only along the first dimension.<sup>283</sup> It is only the first dimension – the degree of support – that is preserved through inference. If  $p$  implies  $q$ , then, necessarily, the probability of  $q$  is equal to or higher than that of  $p$ . The open knowledge advocate argues that empirical knowledge (or any type of knowledge under dispute) requires that the second dimension of support also be satisfied, i.e. that the probability that the proposition is true be raised by one's evidence.<sup>284</sup>

The following example I find instructive. Suppose a scientist is wondering whether to invest money in an experiment that, if successful, will support the hypothesis that  $p$ . Suppose further that the scientist is not interested in  $p$  but rather in  $q$  that is entailed by  $p$ , the probability of which will be lowered if the experiment is successful (supporting  $p$ ). Now imagine the scientist reasons as follows: "I am well aware that if the experiment turns out the results I expect it will lower the probability that  $q$  is true. So I know I will not gain evidence for  $q$ . Nevertheless I will have evidence for  $p$ , and will then infer  $q$  from  $p$  and thus acquire justification for believing  $q$ . So, granted, I will have no new evidence for my desired conclusion, but still, who needs evidence when there's justification?" I take it that such reasoning is untenable.<sup>285</sup>

The example suggests that the point may be more general than the question of whether the evidence raises or lowers the probability of some proposition; that knowledge requires something qualitatively different from what justifies belief. This is reflected in some of our most entrenched linguistic

---

the initial credence might be the same and so would be the posterior probability after the 100 flips in the first case, and the 100,000 in the second.

<sup>283</sup> I think that an account which takes the raising of probability as part of a sufficient condition for knowledge, is ultimately unworkable. Suppose I bought a ticket to a 100,000 ticket lottery, only to discover that there has been a change and the winning ticket will be randomly selected from one million tickets. Although this evidence raises the probability that my ticket is a loser, I do not know that my ticket is a loser.

In conversation Professor Hawthorne proposed the following case related to the idea that probability raising evidence and truth can be part of a sufficient for knowledge. Suppose there are two tribes A and B. A has a tribe member that is shorter than 5 feet but all the rest (perhaps many tribe members) are over 8 feet tall, while the other tribe - B - has tribe members that are all between 6 and 7 feet. I learn that an A tribe's member, call him/her  $a$ , has been selected for a special ceremony. If this evidence is good enough for knowledge that  $a$  is more than 8 feet tall, I would know that the tribe member selected for this ceremony taller than all the members of tribe B. However, since the probability that  $a$  is shorter than 5 feet has grown, I would know that  $a$  is over 8 feet tall but not that s/he is not shorter than 5 feet tall.

The open knowledge advocate can either say that this is in fact the case. I.e. that one can, in fact, know that  $a$  is over 8 feet tall but not know that  $a$  is not shorter than 5 feet. The explanation would be that the evidential situation is what makes for this odd result. Perhaps, however, this case is too similar to the lottery case to count as knowledge and is just one more instance where probability-raising evidence is insufficient for knowledge. I am not sure. Thanks to Professor Hawthorne for suggesting this example.

<sup>284</sup> Or at least that the evidence on which it is based (or which supports the proposition from which it is derived) not lower the probability that one's belief is true.

<sup>285</sup> The example is inspired by Kaplan (1996: 45).

practices regarding knowledge and belief. While questioning, “how do you know?” is perfectly natural and intelligible, the question “how do you believe?” is hardly either of these. Conversely, the question “why do you believe that  $p$ ?” is commonplace, whereas the question “why do you know that  $p$ ?” is unheard of (besides some contexts where it is assumed that *you* were not supposed to know that  $p$ ). Notice that both questions pertain to justification. When asked why one believes something one is prompted to provide a justification for one’s belief. When asked how you know something, likewise, you are required to come up with the grounds or justification for your knowledge claim. In both cases the question is what it is that supports one’s belief/knowledge. And yet the question takes on significantly different forms in the context of belief and in the context of knowledge. We use different notions of justification in these respective contexts. When referring to beliefs we ask for one’s reasons for believing it. Referring to knowledge we ask *how* it is supported. We ask for evidence.

This suggests that knowledge is governed, among other things, by objective external constraints (such as evidence), while belief is also sensitive to rational constraints such as reasons and coherence with other attitudes. As the previous reflections suggest, doxastic justification can be based on agent-relative reasons such as coherence.<sup>286</sup> Being justified in believing something depends on how it relates to the rest of one’s attitudes. But this does not always suffice for knowledge. That is why if someone were to ask “how do you know the external world is real?” answering “well, it follows from the fact that I have hands” or “it coheres with many of my beliefs” would hardly seem appropriate. When it comes to knowledge, it matters *how* the belief is justified.<sup>287</sup> Epistemic justification, we might say, is backtracking – it tracks *how* the justification was acquired or based.

---

<sup>286</sup> My use of this notion is akin to Parfit’s, despite the obvious difference in context. As Parfit says, agent-relative reasons “are reasons only for the agent...When I call some reason agent-relative, I am not claiming that this reason cannot be a reason for other agents. All that I am claiming is that it may not be” Parfit (1986: 143). The fact that  $p$  coheres with my beliefs may be a reason for me to believe it, but might not be a reason for you if your doxastic repertoire is different from mine. It is interesting to note that in the cases above whether one’s evidence supports  $p$ , and thus provides reason for believing  $q$ , depends on one’s belief states. Since the evidence in each case supports both  $p$  and *not- $q$*  (e.g. that I have a hand or that I am experiencing vat hands), whether it counts as a reason for believing  $q$  or not, depends on whether one believes that  $p$  is true. In general epistemologists neglect the fact that there are those who hold such things as true. Gnostics, for instance, believed that our world is governed by an evil deity while the benevolent God is in exile. Berkeley believed that there are no external material objects. Taking these and other positions more seriously would perhaps facilitate greater appreciation of the kind of justification I am trying to demarcate. While you might be justified in believing that there are material external objects, Berkeley might not have been. But this does not mean you have better evidence than he did.

<sup>287</sup> The same thought, I take it, is behind some forms of reliabilism and sensitivity theories of knowledge – it is not enough that one has reason to believe something is true, or that the belief is in itself justified (perhaps it is not even necessary), one must stand in a certain epistemic relation to it.

To explicate the notion of backtracking consider the following scenario. If asked why he believes the world was created in six days, a believer in God and in the truth of all that is asserted in the bible can present the biblical creation story as his justification. In fact, had he professed his religious beliefs while denying that the world was created in less than a week, he would probably be deemed irrational – failing to appreciate what his beliefs commit him to (assuming the person is familiar with the biblical myth of creation). But claiming to know this, invoking the biblical myth would hardly seem appropriate. The credibility of his source will immediately be questioned.

In the same way, seeing your car in the driveway justifies your belief that it has not been stolen. Remembering where you parked it warrants the belief that it is where it was parked and this belief in turn warrants the belief that it hasn't been stolen. Knowing that your plane will land in Chicago, requires a belief that it will not crash on the way to Chicago. But none of them epistemically justifies this latter belief. For this it matters *how* the justification was received. If – backtracking your justification – we find that your belief is based on your looking at your car, we would not question your knowing that it has not been stolen. But if it was based on memory of parking the car, or in the second case looking at the flight schedule on an airport monitor, we do not ascribe to you such knowledge. To doxastically justify  $q$ , suffice it that  $p$  stand in the appropriate logical relation to  $q$ . To justify it epistemically, the way in which  $p$  was evidentially established must be taken into account as well. The point is a simple one. Just as there can be practical reasons for believing something, which provide practical, but not doxastic, justification for one's belief, so too there may be reasons providing doxastic warrant (that is justification for believing), but not epistemic justification (the kind needed for knowledge). Epistemic justification is backtracking – sensitive to the ways in which it was formed or acquired. Therefore, when one's belief is based on evidence lowering the probability that it is true, the belief may be warranted (if the probability is high enough), but one does not know it.

Surely, a lot more than I am able to provide here needs to be said about the details of the distinction.<sup>288</sup> What I have tried to show, however, is merely that with the aid of a reasonable distinction between doxastic warrant and epistemic justification – a distinction that is in line with intuition – the idea that knowledge is open can be sustained, providing its many epistemological benefits without sacrificing the idea that a belief properly inferred from knowledge is justified. The proposal is that an account of epistemic openness can be given while retaining (at least some version of) closure of justifica-

---

<sup>288</sup> I will have more to say in the Appendix.

tion.<sup>289</sup> Other proposals on behalf of knowledge openness are surely possible and perhaps some of them will turn out to be more appealing.

## Chapter Summary

The current state of the debate suggests that any position regarding the validity of epistemic closure must pay some intuitive cost. I have therefore tried to steer the debate about closure away from the battleground of intuitions and counter-intuitions and into the realm of theoretical considerations (as much as I could). Traditionally, such reasons for rejecting closure were advanced by externalist epistemologies. Philosophers such as Dretske and Nozick are famous (or infamous) for having argued against closure not on the basis of its unintuitive consequences, but rather substantive epistemological positions.<sup>290</sup> In contrast to this traditional setting of the debate, the argument suggested here on behalf of the open knowledge advocate is that the logic of evidence supplies the most favorable grounds for epistemic openness. Rejections of closure grounded in the subjunctive nature of knowledge, do not stand up to Hawthorne's charges of inconsistency. Furthermore, such positions, fail to appreciate the evidential nature of knowledge and the backtracking structure of epistemic justification. It is these features of knowledge, the open knowledge advocate argues, that give rise to and explain its openness. The open knowledge position advanced here thus provides a unified, simple and elegant account of the failure of various seemingly intuitive epistemic principles and offers a systematic foundation for reaping the numerous theoretical fruits of epistemic openness.

The open knowledge advocate must admit that the denial of closure has its costs. Yet at least some of its unintuitive consequences are grounded in the unintuitive logic of evidence which all must accept, and can be (at least partially) accommodated by distinguishing between doxastic and epistemic justification. Since belief is governed by rationality, most prominently by (probabilistic) coherence, believing that  $p$  and that  $p$  implies  $q$ , one ought to believe  $q$ . Knowledge on the other hand, depends on justification and, in the case of empirical knowledge, on evidential justification. If the evidence one has lowers the probability that something is true, one does not know it. This oft-conflated disparity can explain the inclination to dismiss epistemic openness. Whether one ought to believe something depends on its relation to other things one takes to be true and thus on the inferences one makes. But

---

<sup>289</sup> I have argued that Williamson trades justification for knowledge closure. Similarly, I am saying that the trade might be a good one, only that what should be kept is justification not knowledge closure.

<sup>290</sup> See Chapter 1.5.2 and 1.5.3.

this should not be confused with the question of whether what one has derived enjoys evidential support requisite for the status of knowledge.

The arguments for open knowledge do not depend on the contentious definition of evidence by purely probabilistic notions. Rather, I have only assumed that evidence does not lower the probability of that which it is evidence for. Even this modest assumption is not needed. By accepting that there are cases of (weak) underdetermination, we are already committed to the rejection of evidence closure, addition (EAD) and distribution (EDIS). What counts as evidence for what and to what degree, is an extremely complicated issue, perhaps no less complex than reasoning itself and no less elusive than the ingenuity of our multifarious attempts at reaching truth. This should not deter us from illuminating some aspects of evidential support by identifying and drawing out connections between evidence and principles of probability of which, arguably, we have clearer understanding. The idea has been that without pretending to know what evidence ultimately amounts to, we can show something about the logic of evidence and use it to draw conclusions about knowledge and the principles it is governed by. The evidence-knowledge link (NED) provides good ground for being suspicious of principles that do not coalesce with the features of evidence on which, presumably, some knowledge depends. This suspicion can be formulated as a challenge. If evidence is not closed under implication, how can empirical knowledge be so closed? What allows knowledge to break free from that which it is based on? How can inference provide what the evidence enabling it cannot?

In the course of this argument I have also provided an analysis of why evidence fails to be closed under different logical operations. The basic idea was that although the conditional probability of the implied proposition given the evidence is not low (not lower than that of the proposition supported by the evidence), given high initial probability (relative to the known proposition) the evidence can, and often does, lower the probability that the proposition is true. Thus, the evidence may change what we might call the “direction” of support. Evidence is basically directional, it points in favor of the truth of some proposition or against it. Evidence pointing in favor of one proposition may point against a proposition it entails.

Using this characterization of evidence, I have claimed on behalf of open knowledge that various epistemological issues that are often considered distinct are, at bottom, one and the same phenomenon, namely, the openness of evidence. The puzzle of dogmatism, “lottery propositions”, the problem of easy knowledge, purported knowledge of “heavyweight propositions” or intuitively implausible knowledge of “lightweight” propositions, are different manifestations of the queer structure of knowledge owing to the openness of evidential support. The implausible ramifications of epistemic closure in the different types of cases discussed in the literature are all one and



the same. They all share a common feature, namely, exceeding the scope of the evidence on which the propositions from which they are derived is based.

The basic argument of knowledge openness has been this. If knowledge in any domain is underdetermined, e.g. propositions known about future events, inductive propositions, or even propositions about objects and events in one's immediate perceptually accessible surrounding, then this knowledge will entail propositions that are not supported by the available evidence on the basis of which knowledge was acquired. This claim was based on an argument showing that for any proposition that is underdetermined by a body of evidence, there will be propositions such that those propositions follow from the propositions that are not supported by the body of evidence. Now if we accept that one cannot go from not knowing to knowing without having evidence (or even that this is at least sometimes impossible), then we have a *prima facie* reason to doubt closure of knowledge.

Without pretending to have covered the entire spectrum of possible responses, two main strategies were considered for countering the challenge to closed knowledge. One was to view the inference itself as adding justification to the inferred consequences of known propositions. Second, that whenever a proposition is known it changes the evidential situation. Now that the proposition is known it can be taken as evidence and hence will support its properly inferred consequences. The first option is gravely flawed, since if anything, inference transfers warrant, it does not create any new warrant that was not there to begin with. The second option for countering the openness argument was further divided into two. The first way is congenial to open knowledge, it appeals to Jeffrey Conditionalization and was shown not to support any defense of closure against the open knowledge argument. The second way, Williamson's way, is to use Standard Conditionalization with the understanding that anything that is known is evidence. Since propositions trivially follow from themselves, all knowledge must have probability 1. This suggestion involved a non-standard understanding of prior probabilities to allow promotion of known propositions to epistemic probability 1 while chances are clearly lower (perhaps even very close to 0). This promotion of probabilities (relative to the prior conditional probability) in turn leads to several unhappy consequences that in light of the benefits open knowledge can offer, reveals its many shortcomings. As far as I can see, Williamson's knowledge account is as plausible regarding a domain of interest as the plausibility of regarding this domain as infallible. Where a domain is infallible, Williamson's theory seems to deliver the right results, but when it is clear that evidence does not entail known propositions in a given domain (like in the case of future events, induction, car parking knowledge and the like) the problems associated with the rule of conjunction introduction come to haunt the theory.

Our conception of knowledge includes the following ideas. First, that we have knowledge of truths. Second that we gain such knowledge by way of

evidence that more often than not, is not conclusive (the evidence is compatible with the falsity of what we know by it). Third, that we do not know certain empirical truths that are implied by what we do know (either because given our epistemic limitations we cannot know them as in the case of heavyweight propositions, or because the grounds we have do not suffice for knowing them, as in the case of ordinary propositions exemplified by the watch, zebra and car cases). And fourth, that knowledge can always be extended by deduction. Combined, these ideas generate a contradiction giving rise to a host of problems and examples that amount to what is perhaps of the most pertinent problems of contemporary epistemology.

Various ways have been proposed of how to modify or deny each of the above stated ideas. Skepticism opts for denial of the claim that we have knowledge even of the most mundane sort. Infallibilists deny that knowledge can be had on the basis of inconclusive reasons. Others claim that we have *a priori* knowledge of heavyweight propositions and perhaps even of non-heavyweight implications of what we know, or that by having knowledge of ordinary truths we *ipso facto* gain knowledge of their implications. The costs and shortcomings of each of these proposals are by now familiar. I have tried to show that the variety of problems arising from our ideas about knowledge are owed to the unintuitive features of evidence and that a proper understanding of these features supports the resolution of these problems by rejecting epistemic closure. By sustaining a distinction between doxastic and epistemic justification the open knowledge advocate is able to account (at least partially) for the intuitive pull of closure – believing that *p* and that *p* implies *q* one is normally justified in believing *q*. Yet beliefs justified in this way might not amount to knowledge.

I, of course, do not pretend to have any way of getting behind the epistemic scenes to check and see if knowledge is closed or open. To be sure, many epistemologists viewing certain closure-friendly features of knowledge as the most central, will remain unconvinced. My purpose is not to make a case that would convince them; I am not sure if I am convinced myself. Yet I think that the case made in this chapter for open knowledge is an improvement on those I am familiar with and is a strong one at that.

## Chapter 6: Concluding Remarks

To cast additional light on the basic line of reasoning of this manuscript I will utilize these concluding remarks without trying to improve or alter the arguments themselves. Rather, placing them in a different order with change of emphasis might facilitate a better understanding of how they relate.

In trying to systematize thinking about knowledge, a good place to start is with our intuitions regarding the principles that govern this notion. We think intuitively that knowledge entails truth; that it entails rational belief; that it must rely on evidence or reasons; that besides special cases it does not require the evidence or reasons to be conclusive; that deductive valid inferences are a general way in which knowledge is extended; and more. Yet upon closer inspection we find that there are tensions between these intuitions. Specifically in light of the idea that we have non-conclusive knowledge, the intuitions regarding the extension of knowledge by inference and that one cannot know a belief that one is rationally required to disbelieve, are in tension. Rational doubt or uncertainty can accumulate, as we have seen (Chapter 1), and so valid inferences can have as their conclusion beliefs that we are required by rationality to doubt. Thus we are forced to either abandon the idea that knowledge can be fallible and tie knowledge to epistemic certainty, or jettison the idea that inference from knowledge is a general way of extending knowledge. Pagin's argument of Chapter 1 shows that the persistent intuition that basic *modus ponens* inferences are immune to the kind of trouble we face in Preface paradox type cases is deeply flawed. We must either give up the extension of knowledge by proper inference or give up epistemic uncertainty.

To my awareness the best attempt at the latter idea is Timothy Williamson's safety account of knowledge. By eliminating epistemic uncertainty it succeeds in preserving the role of valid inference in extending knowledge. This solution has a heavy cost, however, the cost of an infallibilist epistemology. Besides the problem of giving up the idea of fallibilism itself, infallibilism leads to further formidable difficulties (as we saw in Chapter 5.3.3). One problem is directly connected to the rift between knowledge and rational uncertain belief. Following Williamson's proposal entails accepting the possibility that one will know that a long conjunction is true and yet know for certain that its chance of being true is minute. Not only is this a formidable problem in its own right, it also shows that this view does not

connect practical reasoning with knowledge as one might have hoped.<sup>291</sup> A further (though related) problem arises with respect to knowledge of lottery propositions. The infallibilist has two alternatives: The first is to admit that one can know that a ticket of a large future lottery has a very good chance of winning (by inferring this from knowledge that other tickets will lose).<sup>292</sup> The second alternative is to accept equally implausible consequences, for instance, that knowledge of a long conjunction is destroyed simply because the events described by this conjunction will be used as a mechanism for selecting a winner of a prize among many tickets.<sup>293</sup>

These and other problems, which were the subject of Chapter 5.3.3, should persuade us to look for a fallibilist notion of knowledge, one that will avoid these theoretical difficulties. But if knowledge is fallible, rational doubt can accumulate over valid inferences. There seems to be no other choice but to forfeit closure of valid inference altogether or to limit closure to inferences from single items of knowledge, i.e. valid inferences from what falls within the scope of one knowledge operator. Both of these paths encounter difficulties. The first – open knowledge – requires giving up one of the basic intuitive ideas we started out with: how can one know that  $p$  is true and competently infer  $q$  from  $p$  and yet fail to know that  $q$ ? After all, the subject seems to have the best kind of epistemic standing with respect to  $q$ . She believes that  $q$  by properly inferring it from knowledge. The second path, which is a limited form of knowledge closure, leads to a host of problems that have been the main concern of this manuscript. But before rehearsing those, it must be made clear that this limited closed knowledge suggestion shares the difficulty faced by the first option. Anyone who accepts this option (as do many contemporary epistemologists) must admit that one can know that  $p$ , know that  $q$  follows from  $p$ , and yet fail to know that  $q$ . The reason is simple: knowledge that  $q$  in this case is based on an inference from two premises. So the popular strategy of restricting epistemic closure to inferences from a single premise fares no better with respect to the basic idea that proper modes of inference such as *modus ponens* always serve to extend knowledge.

As noted, accepting the restriction of closure to a single premise together with the thesis of fallible (underdetermined) knowledge that is compatible

---

<sup>291</sup> For several adherents of Williamson's view, e.g. Hawthorne and Stanley, this is a difficulty.

<sup>292</sup> Other examples are: inferring and coming to know that one's ticket is a loser on the grounds that one knows that one will not have money to go on Safari (Hawthorne: 2004a); inferring that I will not suffer a fatal heart attack today from my knowledge that I will see a certain movie tomorrow, inferring and coming to know from my knowledge that my plane has a layover in Chicago that it will not crash on the way (Cohen 1999: 58); etc.

<sup>293</sup> Here I am referring to the idea that one will know that many quantum events that have a very high chance of taking place (such as the tunneling through the floor of one among many marbles) will not take place, and will lose this knowledge once these events are used for selecting a lottery winner. See 5.3.3.2 for further detail.

with some measure of epistemic doubt, leads to several problems that constitute the main concern of this manuscript. These problems can be described more generally as two types of challenges. The first type is the challenge of how to respond to implications of closed knowledge. How to respond to Cartesian Skeptical implications, to the implication of Mundane Skepticism of Live Skepticism, Dogmatism, Bootstrapping, Easy Knowledge and other Epistemic Ascent implications. The second type of challenge to closed knowledge is the theoretical difficulties that come from the logic of evidence and its relation to knowledge.

Chapter 2 focused on the first type of challenge: Skeptical implications of closure. The most notorious of these is Cartesian Skepticism and it received a more systematic treatment than did the other skeptical challenges. In brief the problem was here analyzed as one where the Cartesian skeptic has a way to show that her opponent has no evidence against the possibilities she raises. Worse, she has an argument showing that one's evidence counts in favor of such possibilities relative to any initial probability one chooses. The Cartesian Skeptic has a systematic way of generating such propositions that intuitively and on very modest assumptions are not known. She can then use closure for a *reductio* argument discrediting the opponent's claim to know propositions of a disputed domain (e.g. knowledge of past events). A second skeptical argument – Mundane Skepticism – ensues on the basis of a non-systematic ability. Its main virtue is that it appeals to possibilities that in contrast to the Cartesian far-fetched scenarios, are quite mundane. The contrast is not accidental, the Cartesian Skeptic relies on cases that are not presently known to be false by all subjects, e.g. that the world was not created five minutes ago. Mundane Skepticism appeals to cases that are only beyond the evidence of a specific epistemic agent in a certain context. For instance, one does not know that one will lose tomorrow's lottery. A third type of skepticism appeals to modesty regarding scientific/philosophical theories that are not regarded as prominent and yet have some chance of becoming the received view. This is Live Skepticism advanced by Frances (2004). It relies (despite his claim to the contrary) on closure. When successful it targets a subject's everyday knowledge by appealing to the fact that we don't (at present) know such theories to be false. Claiming otherwise is not only immodest (especially if experts despite believing that they are false do not think they know them to be false), it also seems false. Our present evidence (evidence that the non-standard theorists are aware of) is not good enough for us to *know* that these non-standard theories are false. These three skepticisms pose a challenge to those who both believe that knowledge is fallible and closed.

A second type of problematic implications for closed knowledge falls under the heading of Epistemic Ascent. We saw in Chapter 3 that Kripke's Dogmatism remains a formidable problem for those who assume closed knowledge since one who fallibly knows that a proposition  $p$  is true can de-

duce (and by closure come to know) that any evidence that counts against  $p$  will be misleading. Not only does this knowledge sound dubious under the circumstances (as some closure advocates candidly admit), but there are Gettier type considerations that urge the conclusion that in fact these are not cases of knowledge. For instance Hawthorne's (2004a: 73) Manchester United example is one where a false belief is necessary in maintaining the belief that is known if closure holds.<sup>294</sup>

The similarity between the different kinds of Epistemic Ascent is not accidental. Kripke's puzzle relates to the inference that if  $p$  is true, any counterevidence to  $p$  will be misleading. Vogel's Bootstrapping argument, for instance, is an inversion of the consequent – the evidence I have for  $p$  is not misleading.<sup>295</sup> The same kind of reasoning is exemplified in Cohen's Easy Knowledge example (I know that this table is not a white table lit by a red light to look red) and in analogous cases of belief (I know that my belief that Churchill was born November 1974 is not mistaken). Under the heading of Epistemic Ascent these cases were further analyzed as propositions regarding one's own epistemic state that are inferred from a known proposition and are not supported by one's evidence. The ascent is made purely by reflection and supposedly improves one's epistemic standing regarding the relation between one's epistemic subjective state and the world. Intuitively, these are not ways in which knowledge can be gained and yet closure tells otherwise.

The second type of challenge, the theoretical challenge, was the main focus of Chapters 4 and 5. Chapter 4 started by substantiating the claim that evidence is open, i.e. if  $e$  supports  $p$  to any degree short of an absolute guar-

---

<sup>294</sup> I note a contrast between the Dogmatism Puzzle (and cases of Epistemic Ascent more generally) and Cartesian Skepticism. The skeptic invokes propositions that are alternative explanations of one's evidence. For instance, that your evidence that you have hands (whatever it is) is explained by your being a brain in a vat that is fed sensory stimulations as if you had hands. That these possibilities are alternatives allows a response to Skepticism that is not available as a response to the Dogmatism puzzle. This is the idea that if an alternative proposition becomes salient to the epistemic subject (or to a knowledge ascriber), the standards of knowledge alter and if the evidence does not rule-out these alternatives, knowledge is lost (or the knowledge predicate can no longer truly be ascribed to the subject-proposition pair). The dogmatism case is not one that lends itself to this tactic. That there is evidence that I have not encountered counting against the proposition I know, is *not* an alternative explanation of my evidence. Thus, besides the fact that in some cases of dogmatic inference error becomes salient, this need not be the case and so the above tactic is not enough to accommodate lack of evidence with closure. Moreover, some Epistemic Ascent share an important feature exemplified in Kripke's Dogmatism puzzle. There is no intuitive inclination to claim that one does not know the proposition from which the questionable knowledge was inferred. Thus, even if one thinks that intuition and linguistic practice should convince us that I don't know that there is a zebra in the pen if I don't know that there is no disguised mule in the pen, one will have to find a different reason for maintaining closure in face of these Epistemic Ascent cases. There is no intuitive inclination to question one's knowledge that  $p$ , if one does not know that any future counterevidence to  $p$  will be misleading.

<sup>295</sup> It was shown in Chapter 3 that one does not need to know what one's evidence is in order to infer the kind of objectionable knowledge that is at the heart of Vogel's Bootstrapping argument.

antee,  $e$  will not support all the propositions that follow from  $p$ . This fact (at least since Carnap and Hempel see Chapter 4.3.4) has been noted by many theorists but I suspect that its potential as a threat to knowledge closure has been underestimated.<sup>296</sup> Weaker principles than evidence closure were found to be false, such as the idea that if  $p$  is supported non-conclusively by  $e$ ,  $e$  will support  $p$ -or- $q$ . The watch case was proven (within a fallibilist framework) to be a case where one has good evidence for the time being 3:00 and yet one does not gain any evidence for the claim that *either my watch is not showing “3:00” or it is showing the correct time*. Whether or not one thinks that the evidence that allows one to know that the time is 3:00 (i.e. the watch showing “3:00”) intuitively supports this disjunction, this evidence does not in fact support it. And yet, closure entails that it is known as long as it is known that the time is 3:00. If it is in fact known, then one must have known this beforehand since the added evidence – the evidence that allowed one to know that it is 3:00 - if anything, counts against this proposition that follows from the time being 3:00.

The argument for open knowledge in Chapter 4 spells out the challenge (page 97). Roughly the theoretical challenge was this. If one comes to know that  $p$  on the basis of evidence  $e$ , no matter how high the probability of  $p$  is on this evidence, as long as the conditional probability of  $p$  on  $e$  is less than 1, there will be propositions that follow from  $p$  that will not be supported by this evidence. Looking at the situation as a whole, closure entails that one can go from ignorance to knowledge without evidence (and in the face of counterevidence) or that our *a priori* knowledge is much more extensive than we have supposed. As stated, the theoretical challenge can be used to analyze the first type of challenge, problematic implications of closure – Skepticisms and Epistemic Ascents – in terms of the principles of evidence and/or its probabilistic properties. Cartesian Skepticism (2.2.2), Mundane Skepticism (2.3), Live Skepticism (2.4), Dogmatism (3.1), Bootstrapping (3.2.1), Easy Knowledge (3.2.1) and other Epistemic Ascent cases (3.2) are all instances of inferring propositions from knowledge that the evidence for this knowledge does not support. Thus, the challenges combine into a comprehensive argument against closed knowledge and an analysis of its failure.

Two ways to avoid the argument are wrongheaded. One is to claim that evidence is not analyzable probabilistically, two is to claim that there are items of evidence that count in favor of propositions that have their probability lowered by those items. The Underdetermination Argument (page 104) was introduced as a way to counter these ways of responding to the Open Knowledge Argument. It shows – without appeal to probabilities – that valid

---

<sup>296</sup> White (2006), in a sense, is an exception. His argument makes use of evidence openness on (Bayesian) considerations similar to the probabilistic argument advanced here. His target, however, is limited to the Dogmatism account of justification closure. Justification closure in the sense endorsed in that paper, by the way, is compatible with open knowledge. See 5.4 and the Appendix for more detail on probabilistic coherence.

inference does not universally preserve evidence. This conclusion holds for very intuitive principles that are weaker than closure. Moreover, it also shows that an open knowledge account that is based on open evidence is successful where other open knowledge accounts (Nozick's and Dretske's – Chapter 1.5.2-1.5.3) fail (i.e. in responding to Hawthorne's arguments 1.5.1). It shows this by emulating principles with respect to evidence that Hawthorne appeals to with respect to knowledge. Not only does it show that these principles that are weaker than evidence closure fail, but it can also be used to explain why despite appearances to the contrary their epistemic counterparts (the principles that are weaker than knowledge closure) are invalid. Why, that is, despite their falsity they seem extremely compelling. Just as it might seem initially inevitable that one's evidence for *p-and-q* must at least count either in favor of *p* or in favor of *q*, this thought is mistaken. Also, one's evidence for *p* might not count in favor of *p-or-q* despite the initial thought that it must. It is these (as well as similar) mistaken inclinations that an open knowledge advocate might use to explain our mistaken thought to endorse closure either within or across contexts of knowledge ascription (or practical environment).

For both fallibilist and infallibilist closure advocates a significant challenge stems from more basic assumptions than closure. Moreover open knowledge has significant advantages. Not only can it meet the theoretical challenge and avoid the problematic implications of Epistemic Ascent and Scepticisms, the Lottery Paradox, The Preface Paradox, etc, some of its intuitive costs can be mitigated by a distinction between epistemic and doxastic justification. Compatible with open knowledge is the idea that doxastic justification is (single-premise) closed. If one has justification for believing that *p* and properly infers *q* from *p*, one is justified in believing *q*. And yet, one will not always have a justification for believing that *q* (at least not an evidential justification).

I conclude, then, first, that it can be shown on both probabilistic and non-probabilistic assumptions that evidence is in fact open. Second, the dependence of knowledge (or at least some knowledge) on evidence (as expressed by (NED) (p. 126) is a compelling defensible idea. Third, that it can be shown that for every fallibly known propositions that are many propositions that follow from it that are not supported by its evidence. Fourth, that this last claim together with the dependence relation of knowledge on evidence gives reason – in light of the latter's openness – to view knowledge as open as well. Fifth, that open knowledge resolves many central epistemological problems that do not seem to be solved by any other single account. Sixth, that it is immune to the kind of arguments launched against other forms of open knowledge. Seventh, that open knowledge is compatible with doxastic justification closure. Eighth, that some of its unintuitive costs can be explained by the intuitive costs that we have to pay anyhow with regard to evidence and its un-intuitiveness can help explain why the problems it re-



solves have such a hold on us. Finally, then, evidence openness provides a good (perhaps even the best) motivation for, and explanation of, knowledge openness.

# Bibliography

- Audi, Robert (1991). "Justification, Deductive Closure and Reasons to Believe," *Dialogue* 30: 77-84.
- Audi, Robert (1995). "Deductive Closure, Defeasibility and Scepticism: A Reply to Feldman," *Philosophical Quarterly* 45: 494-499.
- Austin, John L. (1946). "Other Minds," *Proceedings of the Aristotelian Society*, Supplementary Volume 20: 148-87.
- Bar-Hillel, M. & Ruma Falk, R. (1982). "Some teasers concerning conditional probabilities," *Cognition*, 11: 109-122.
- Bar-Hillel, Maya & Margalit, Avishai (1985). "Gideon's Paradox — A Paradox of Rationality," *Synthese* 63 (2).
- Bovens, Luc and Hartmann, Stephan (2003). *Bayesian Epistemology*, Oxford: OUP.
- Brown, Jessica (2004). *Anti-Individualism and Knowledge*, MIT, Cambridge.
- Brueckner, Anthony (1985a), "Losing Track of the Sceptic," *Analysis* 45: 103-104.
- Brueckner, Anthony (1985b). "Skepticism and Epistemic Closure," *Philosophical Topics* 13: 89-117.
- Brueckner, Anthony (1985c). "Transmission for Knowledge Not Established," *Philosophical Quarterly* 35: 193-196.
- Brueckner, Anthony (1994). "The Structure of the Skeptical Argument," *Philosophy and Phenomenological Research*, 54: 827-35.
- Brueckner, Anthony (2004). "Strategies for Refuting Closure," *Analysis*, 64: 333-335.
- Brueckner, Anthony (2005). "Fallibilism, Underdetermination, and Skepticism," *Philosophy and Phenomenological Research*, 71(2): 384-391.
- Cargile, James (1995). "Justification and Misleading Defeaters", *Analysis*, 55, 3: 216-220.
- Carnap, R. (1950). *Logical Foundations of Probability*, Chicago: Chicago University Press.
- Casullo, Albert (2003). *A Priori Justification*, Oxford: OUP.
- Cohen, Stewart (1988). "How to be a fallibilist," *Philosophical Perspectives*, 2, Epistemology: 91-123.
- Cohen, Stewart (1998). "Two Kinds of Skeptical Argument," *Philosophy and Phenomenological Research*, Vol. 58, No. 1: 143-159.
- Cohen, Stewart (1999). "Contextualism, Skepticism, and The Structure of Reasons," *Philosophical Perspectives* 13: Epistemology: 57-89.
- Cohen, Stewart (2002). "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65: 309-328.
- Cohen, Stewart (2005). "Why Basic Knowledge is Easy Knowledge," *Philosophy and Phenomenological Research*, 75.2: 417-30.
- Conee, Earl and Feldman, Richard (2004). *Evidentialism: Essays in Epistemology*, Oxford: Clarendon Press.

- Cross, Charles B. (2001). "The Paradox of the Knower Without Epistemic Closure," *Mind* 110 (438).
- Davidson, Donald (1990). "A Coherence Theory of Truth and Knowledge," in A.R. Malachowski (ed.) *Reading Rorty. Critical Responses to Philosophy and the Mirror of Nature (and Beyond)*, Oxford: Blackwell Publishers.
- Davies, M. (1998). "Externalism, Architecturalism, and Epistemic Warrant," in *Knowing Our Own Minds: Essays on Self-Knowledge*, (eds.) C. Macdonald, B. Smith & C. J. G. Wright, England: OUP, Oxford, 321-61.
- DeRose, Keith (1995). "Solving the Sceptical Puzzle," *Philosophical Review* 104: 1-52.
- Dretske, Fred (1970). "Epistemic Operators," *Journal of Philosophy* 67: 1007-1023.
- Dretske, Fred (2005a). "Is Knowledge Closed Under Known Entailment?" in M. Steup & E. Sosa (eds.), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA: 13-26.
- Dretske, Fred (2005b). "Reply to Hawthorne," in M. Steup & E. Sosa (eds.), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA, : 43-46.
- Dretske, Fred (2006). "Information and Closure," *Erkenntnis* 64: 409-413.
- Earman, John (1992), *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge, MA: MIT Press.
- Eells, Ellery, and Branden Fitelson. 2000. "Measuring Confirmation and Evidence", *Journal of Philosophy* 97: 663-672.
- Egan, Andy (2007). "Epistemic Modals, Relativism, and Assertion," *Philosophical Studies*, 133: 1-22.
- Engel, Mylan (1992). "Personal and Doxastic Justification in Epistemology," *Philosophical Studies* 67: 133-50.
- Feldman, Richard (1981). "Fallibilism and Knowing That One Knows," *Philosophical Review*, 90: 266-82.
- Feldman, Richard (1995). "In Defense of Closure," *The Philosophical Quarterly* 45: 487-494.
- Fitelson, Branden (1999). "The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity," *Philosophy of Science* 66 (Proceedings Supplement): S362-378.
- Fitelson, Branden (2006). "The Paradox of Confirmation," *Philosophy Compass* 1/1: 95-113.
- Frances, Bryan (1999). "Contradictory belief and epistemic closure principles." *Mind and Language* 14 (2): 203-226.
- Frances, Bryan (2004). *Scepticism Comes Alive*, Oxford: OUP.
- Frank Arntzenius (2003). "Some Problems for Conditionalization and Reflection," *Journal of Philosophy*, 100(7): 356-371.
- Fred Dretske (2000). *Perception, Knowledge and Belief*, Cambridge University Press.
- Fumerton, Richard (2006). *Epistemology*, Oxford and Cambridge: Blackwell.
- Gaifman, Haim and Snir, Marc (1982). "Probabilities Over Rich Languages, Testing, and Randomness." *Journal of Symbolic Logic* 47:495-548.
- Gettier, Edmund L. (1963). "Is Justified True Belief, Knowledge?" *Analysis* 23: 121-123.
- Gilbert Harman (1973). *Thought*, Princeton: Princeton UP.
- Goldman, Alvin (1976). "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73: 771-791.
- Goodman, Nelson (1955). *Fact, Fiction, and Forcst*, Cambridge, Mass: Harvard.

- Goodman, Nelson (1972). *Problems and Projects*, New York: The Bobbs-Merrill Company, Inc.
- Harman, Gilbert (1973). *Thought*, Princeton: Princeton University Press.
- Harman, Gilbert and Sherman, Brett (2004). "Knowledge, Assumptions, Lotteries," *Philosophical Issues*, 14, Epistemology.
- Hawthorne, John (2002). "Lewis, the Lottery and the Preface," *Analysis* 62 (3): 242–251.
- Hawthorne, John (2002). "Deeply Contingent A Priori Knowledge," *Philosophy and Phenomenological Research* 65 (2):247-269.
- Hawthorne, John (2004a). *Knowledge and Lotteries*, Oxford: OUP.
- Hawthorne, John (2004b). "Précis" (pp. 476-481) and "Replies to Cohen, Harman and Sherman, and Vogel" (pp. 510-523), Symposium on *Knowledge and Lotteries*, *Philosophical Issues* 14.
- Hawthorne, John (2005). "The Case for Closure," in M. Steup, E. Sosa, eds., *Contemporary Debates in Epistemology*, Blackwell: 26-43.
- Hawthorne, John and Lasonen-Aarnio, Maria (forthcoming). "Knowledge and Objective Chance," in P. Greenough and D. Pritchard, eds., *Williamson on Knowledge*, OUP.
- Hawthorne, John & Stanley, Jason (2008). "Knowledge and Action," *Journal of Philosophy* 105 (10).
- Heller, Mark (1999). "Relevant Alternatives and Closure" *Australasian Journal of Philosophy*: 196-208.
- Hempel, C. G (1965). "Studies in the Logic of Confirmation," in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, Free Press, New York: 3-53.
- Hintikka, J. (1962). *Knowledge and Belief: an introduction to the logic of the two notions*, Cornell University Press.
- Horwich, Paul (1982). *Probability and Evidence*, Cambridge: Cambridge University Press.
- Howson C. and Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach*, Chicago: Open Court Publishing Co.
- Howson C. (2000). *Hume's Problem*, New York: OUP.
- Huber F. (2005). "What is the Point of Confirmation?," *Philosophy of Science*, Vol. 72, No. 5: 1146-1159.
- Hughes G.E. and Cresswell M. J. (1984). *A Companion to Modal Logic*, New York: Methuen.
- Jackson, F. (1979). "On Assertion and Indicative Conditionals," *Philosophical Review* 88: 565-589.
- Jackson F. and Smith M. (2004). *The Oxford Handbook of Analytic Philosophy*, OUP.
- Janvid, Mikael. (2006). "Contextualism and the Structure of Skeptical Arguments," *Dialectica*, 58: 143–157.
- Jeffrey, Richard (1965). *The Logic of Decision*, Chicago: University of Chicago Press.
- Jeffrey, Richard (2004). *Subjective Probability*, Cambridge: Cambridge University Press.
- Joyce, James (2005). "How Probabilities Reflect Evidence," in John Hawthorne (ed.) *Philosophical Perspectives*, vol.19: Epistemology, Oxford: Blackwell.
- Kaplan, David and Richard Montague (1960). "A Paradox Regained," *Notre Dame Journal of Formal Logic*, 1: 79-90.
- Kaplan, Mark (1996). *Decision Theory as Philosophy*, Cambridge Mass: Cambridge University Press.

- Kaplan, Mark (forthcoming). "Tales of the Unknown," in Martin Gustafsson and Richard Sørli, eds., *New Essays on the Philosophy of J. L. Austin*, OUP.
- Karlander, Karl and Spectre, Levi (forthcoming). "Sleeping Beauty Meets Monday," *Synthese*.
- Karlander, Karl and Spectre, Levi (unpublished manuscript). "Actual World Skepticism."
- Kelly, Thomas (forthcoming). "Disagreement, Dogmatism, and Belief Polarization," *Journal of Philosophy*.
- Klein, Peter (2004). "Closure Matters: Academic Skepticism and Easy Knowledge," *Philosophical Issues*, 14, Epistemology: 165-184.
- Kripke, Saul (1979). "A Puzzle about Belief," in *Meaning and Use*, edited by Margalit, Avishai, Dordrecht and Boston: Reidel.
- Kripke, Saul (forthcoming). "On Two Paradoxes of Knowledge," unpublished lecture delivered to the Cambridge Moral Sciences Club.
- Kvanvig, Jonathan (2006). *The Knowability Paradox*, Oxford: OUP.
- Kyburg, Henry (1961). *Probability and the Logic of Rational Belief*, Dordrecht: Kluwer.
- Kyburg, Henry (1970). "Conjunctivitis," in *Induction, Acceptance, and Rational Belief*, Dordrecht: Reidel.
- Lewis, David (1983). *Philosophical Papers*, Volume I, Oxford: OUP.
- Lewis, David (1986). *Philosophical Papers*, Volume II, Oxford: OUP.
- Lewis, David (1996). "Elusive knowledge," *Australasian Journal of Philosophy* 74 (4): 549 – 567. Page references are to the reprinted version in Lewis (1999).
- Lewis, David (1999). *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press.
- MacFarlane, John (2005). "The Assessment Sensitivity of Knowledge Attributions," *Oxford Studies in Epistemology* 1: 197–233.
- Maitzen, Steve (1998). "The Knower Paradox and Epistemic Closure," *Synthese* 114: 337–54.
- Makinson, D. C. (1965). "The Paradox of the Preface," *Analysis* 25: 205–7.
- McGee, Vann (1985). "A Counterexample to Modus Ponens," *Journal of Philosophy* 82 (9):462-471.
- McGinn, Colin (1984). "The Concept of Knowledge," *Midwest Studies in Philosophy* IX 1984: Causation and Causal Theories, Peter French, Theodore E. Uehling, Jr, and Howard K. Wettstein, eds., Minneapolis: 529–554.
- Mellor, D. H. (2005). *Probability: A Philosophical Introduction*, New York: Routledge.
- Moore, G. E. (1959). "Proof of an External World," in *Philosophical Papers*, London: George Allen & Unwin, Ltd.
- Moore, G. E. (1959). "Certainty," in *Philosophical Papers*, London: George Allen & Unwin, Ltd.
- Murphy, Peter (2006). "A Strategy for Assessing Closure," *Erkenntnis* Volume 65, 3: 365-383.
- Nozick, Robert (1981). *Philosophical Explanations*, Cambridge: Harvard University Press.
- Pagin, Peter (1990). "Review of Roy Sorenson, *Blindspots*, Clarendon Press, Oxford 1988," *History and Philosophy of Logic* 11: 243-5.
- Parfit, Derek (1986). *Reasons and Persons*, Oxford: OUP.
- Prawitz, Dag (2009). "Inference and Knowledge," in *The Logical Yearbook 2008*, M. Pelis (ed), London: College Publications, King's College London:175-192.
- Pritchard, D. (2005). "The Structure of Skeptical Arguments", *The Philosophical Quarterly*, Vol. 55, No. 218: 37-52.

- Pryor, James (2000). "The skeptic and the dogmatist," *Noûs* 34 (4):517–549.
- Quine, W.V.O. (1953). "On a so-called Paradox," *Mind*, 62, 245: 65–67.
- Ramsey, F.P. (1960). *The Foundations of Metaphysics*, Paterson, NJ: Littlefield, Adams & Co.
- Rosenkrantz, Roger (1981). *Foundations and Applications of Inductive Probability*, New York: Ridgeview Publishing Company.
- Russell, Bertrand (1912). *The Problems of Philosophy*, Oxford: Oxford.
- Russell, Bertrand (1948). *Human Knowledge Its Scope and Limits*, New York: Simon and Schuster.
- Schaffer, Jonathan (2004). "From Contextualism to Contrastivism," *Philosophical Studies* 119: 73-103.
- Scheffler, Israel (1963). *The Anatomy of Inquiry*, New York: Knopf.
- Scheffler, Israel and Goodman, Nelson (1972). "Selective Confirmation and the Ravens: A Reply to Foster", *Journal of Philosophy*, Vol. 69, No. 3: 78-83.
- Sharon, Assaf and Spectre, Levi (2008). "Mr. Magoo's Mistake," *Philosophical Studies*, 139: 289-306.
- Sharon, Assaf and Spectre, Levi (forthcoming). "Dogmatism Repuzzled," *Philosophical Studies*.
- Sharon, Assaf and Spectre, Levi (unpublished manuscript a). "Three Problems for Williamsonian Epistemology."
- Sharon, Assaf and Spectre, Levi (unpublished manuscript b). "Evidence and Knowledge Openness."
- Shimony, Abner (1988). "An Adamite Derivation of the Calculus of Probability," in J.H. Fetzer, ed., *Probability and Causality*, Dordrecht: Reidel.
- Silins, Nicholas (2005). "Transmission Failure Failure," *Philosophical Studies*, 126: 71-102.
- Silins, Nicholas (2005). "Deception and Evidence," *Philosophical Perspectives* 19 (1): 375–404.
- Sorensen, Roy (1988a). "Dogmatism, Junk Knowledge, and Conditionals," *The Philosophical Quarterly*, 38: 433– 454.
- Sorensen, Roy (1988b). *Blindspots*, Oxford: OUP.
- Spectre, Levi (unpublished manuscript a), "Compartmentalization and Closure."
- Spectre, Levi (unpublished manuscript b), "Epistemic Ascent."
- Spectre, Levi (unpublished manuscript c), "A reply to Veber - "What Do You Do with Misleading Evidence?""
- Stanley, Jason (2005). "Fallibilism and Concessive Knowledge Attributions," *Analysis*, 65 (286): 126-131.
- Stanley, Jason (2005). *Knowledge and Practical Interests*, Oxford: OUP.
- Steup, Matthias and Sosa, Ernest (eds), (2005). *Contemporary Debates in Epistemology*, Malden, MA: Blackwell Publishing.
- Stine, Gail (1971), "Dretske on Knowing the Logical Consequences," *Journal of Philosophy* 68: 296-299.
- Strevens, Michael (2004). "Bayesian Confirmation Theory: Inductive Logic, or Mere Inductive Framework?," *Synthese* 141: 365-379.
- Strevens, Michael (2006). *Notes on Bayesian Confirmation Theory*, published online: <http://www.nyu.edu/gsas/dept/philo/user/strevens/Classes/Conf06/BCT.pdf>
- Thalberg Jr., Irving (1969). "In Defense of Justified True Belief," *Journal of Philosophy*, 66, no. 22: 794–803.
- Uzquiano, Gabriel (2004). "The Paradox of the Knower Without Epistemic Closure?," *Mind* 113 (449).
- Van Fraassen, Bas C. (1984). "Belief and the Will," *Journal of Philosophy* 81:235-256.

- Van Fraassen, Bas C. (1995). "Belief and the Problem of Ulysses and the Sirens," *Philosophical Studies* 77: 7-37.
- Veber, M. (2004). "What Do You Do with Misleading Evidence?," *The Philosophical Quarterly* 54 (217): 557-569.
- Vogel, Jonathan (1990). "Are There Counterexamples to the Closure Principle?," in *Doubting: Contemporary Perspectives on Skepticism*, M. Roth and G. Ross (eds.), Dordrecht: Kluwer Academic Publishers.
- Vogel, Jonathan (2000). "Reliabilism Leveled," *Journal of Philosophy*, Vol. 97, No. 11: 602-623.
- Vogel, Jonathan (2004). "Varieties of Skepticism," *Philosophy and Phenomenological Research* 68: 1-37.
- Vogel, Jonathan (2007). "Subjunctivitis," *Philosophical Studies* 134: 73-88.
- Warfield, T. A. (2004). "When Epistemic Closure Does and Does Not Fail: a Lesson From the History of Epistemology," *Analysis* 64: 35-41.
- White, Roger (2006). "Problems for dogmatism," *Philosophical Studies* 131 (3): 525-57.
- Williamson, Timothy (1986). "The Contingent A Priori: Has It Anything to Do With Indexicals?" *Analysis*: 113-117.
- Williamson, Timothy (2000). *Knowledge and Its Limits*, Oxford: OUP.
- Williamson, Timothy (forthcoming). "Reply to John Hawthorne and Maria Lasonen-Aarnio," in P. Greenough and D. Pritchard, eds., *Williamson on Knowledge*, OUP.
- Wright, Crispin (2000). "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof," *Philosophical Issues* 10: 140-163.
- Wright, Crispin (2003). "Some Reflections on the Acquisition of Warrant by Inference," In: Nuccetelli S. (ed.), *New Essays on Semantic Externalism and Self-Knowledge*, MIT Press, Cambridge, MA: 57-77.
- Yablo, Stephen (Manuscript). "The Truth and Something But the Truth," first Carl Hempel lecture.
- Yalçın, Ü. (1992). "Skeptical Arguments from Underdetermination," *Philosophical Studies* 68: 1-34.