A SOLUTION TO SKEPTICAL PUZZLES

_____

A Dissertation

presented to

the Faculty of the Graduate School

at the University of Missouri-Columbia

_____

In Partial Fulfillment

of the Requirements for the Degree

Doctor of Philosophy

_____

by

KOK YONG LEE

Dr. Matthew McGrath, Dissertation Supervisor

MAY 2014

The undersigned, appointed by the dean of the Graduate School, have examined the dissertation entitled

A SOLUTION TO SKEPTICAL PUZZLES

presented by Kok Yong Lee,

a candidate for the degree of doctor of philosophy,

and hereby certify that, in their opinion, it is worthy of acceptance.

Professor Matthew McGrath

Professor Peter Markie

Professor Paul Weirich

Professor Todd Schachtman

To my father

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

A Solution to Skeptical Puzzles

Kok Yong Lee

Dr. Matthew McGrath, Dissertation Supervisor

ABSTRACT

In this dissertation, I aim at resolving an epistemic puzzle that I call 'the skeptical puzzle'. An instance of this puzzle is as follows:

(1) I know that I have hands.

(2) I don't know that I am not a brain in a vat (being stimulated to experience having hands).

(3) If I know that I have hands, I know that I am not a brain in a vat.

(1)-(3) are puzzling since, while each is individually plausible, they are jointly inconsistent.

In this dissertation, I argue that the truth values of (1)-(3) vary with contexts. More precisely, (1), not-(2), and (3) are true in the *ordinary* context, while not-(1), (2), and (3) are true in the *skeptical* context. This explains why we initially find (1)-(3) plausible: (1)-(3) are initially plausible since they are indeed true (in some contexts). Moreover, (1)-(3) do not constitute a genuine paradox since, in no context, (1)-(3) are true altogether.

My approach to the skeptical puzzle is in line with the contextualist solution to the puzzle, but differs in important aspects. Contextualists claim that the variability of the truth values of (1)-(3) across the ordinary and skeptical contexts consists in the variability of the *standards for knowledge*, which specify how strong one's *epistemic position* with

respect to p must be in order for one to know p, while I argue that the very variability of truth values consists in the variability of S's epistemic position with respect to p across such contexts. That is, my view, but not contextualism, implies that S's epistemic position is context-sensitive in some sense.

I argue that the contextualist solution (and any account that also appeals to the variability of the standards for knowledge) fails to offer a satisfactory answer to the skeptical puzzle. In Chapter Two, I discuss the main reason (or one of the main reasons) for the thesis of the variability of the standards for knowledge, and find it untenable. This shows that the thesis is unmotivated. In Chapter Three, I further argue that the contextualist solution is theoretically uninteresting since it deals only with a kind of skeptical arguments that is theoretically insignificant. This shows that we should not resolve the skeptical puzzle in terms of the variability of the standards for knowledge. Instead, we should adopt the thesis of the variability of S's epistemic position.

I propose that one's epistemic position with respect to p should be characterized by the *sensitivity* of one's reason, where S's reason for p is sensitive just in case S would not have the same reason if p had not been the case. Sensitivity captures our intuitive idea that knowing that p requires one's reason to track the truth of p, and it also explains the persuasiveness of (2): Typically, a person tends to think that she does not know that she is not a brain in a vat, when she realizes that if she were a brain in a vat, she would still have all her current experiences, memories, etc., that suggest that she is not a brain in a vat.

I formulate sensitivity in terms of counterfactual conditionals. How should the truth conditions of such conditionals be understood? The common practice is to appeal to

the possible-worlds semantics of counterfactual conditionals. However, the possible-worlds semantics has suffered from some serious problems. In Chapter Six, I suggest that epistemologists are better off to adopt the causal modeling semantics for counterfactual conditionals. I point out that causal-modeling semantics is particularly helpful in distinguishing two distinct types of counterfactualization—the *interventional* and *inferential counterfactualization*—which are crucial to the characterization of the sensitivity account.

In Chapter Seven, I argue that the assessment of the sensitivity of one's reason for p is relative to the *epistemic inquiry* the assessor is engaging in. There are two basic kinds of epistemic inquiry, defining by two distinct epistemic goals: the *liberal goal* (i.e., to accept a hypothesis unless one has (sufficiently good) reasons against it) and the *conservative goal* (i.e., to accept a hypothesis only if one has (sufficiently good) reasons for it). I argue that (1) and not-(2) are true relative to the liberal inquiry, while not-(1) and (2) are true relative to the conservative inquiry. This explains why (1) and not-(2) are true in the ordinary context, while not-(1) and (2), the skeptical one.

Chapter One

Introducing the Puzzle

1.1 Introduction

This dissertation concentrates on an epistemic puzzle that has received increasing attention recently. The puzzle consists in three individually plausible statements, which are jointly inconsistent. To illustrate, consider the following example.

First, commonsense has it that we know (or we take ourselves to know) a lot about the external world. Specifically, one gains (or tends to gain) knowledge about the world via one's current, attentive, perceptual knowledge. To use G. E. Moore's famous example, by attentively looking at and making certain gestures with one's hands, one is entitled to claim:

(A)   I know that I have hands.

Second, a modest amount of critical reflection shows that we are not always in a position to distinguish veridical experiences from non-veridical ones. When a hypothesis that we are having certain non-veridical experiences is constructed appropriately, we are inclined to admit that we do not know that the hypothesis is false. Consider the *skeptical* hypothesis that one is a handless brain in a vat being fed with stimuli that make one have the experience of looking at and making certain gestures with one's hands. One is inclined to admit that:

(B)   I don't know that I'm not a handless brain in a vat (being fed with stimuli that make me have the experience of looking at and making certain gestures with my hands).

Third, it seems obvious that if one knows that one has hands, then one knows that it is not the case that one does not have hands. Knowing that it is not the case that one does not have hands, in turn, entails knowing the falsity of most (if not all) propositions that are incompatible with the proposition that one has hands.[1] For instance, it is very odd, to say the least, to claim both that one knows that one has hands and that one does not know that one did not lose one's hands ten years ago. By the same token, it is odd to claim both that one knows that one has hands and that one does not know that one is not a handless brain in a vat (being fed with stimuli that make one experience having hands). In fact, many will naturally find the following very plausible, if prompted appropriately:

(C)   I know that I have hands only if I know that I'm not a brain in a vat (being fed with stimuli that make me have the experience of looking at and making certain gesture with my hands).

Ordinary, untrained, intuitions aside, (C) is often claimed to be supported by a very plausible epistemic closure principle:

(Closure)   If S knows that p and S knows that p implies q, then S knows that q.

---

[1] The requirement that knowing that p implies knowing the falsity of all alternatives to p seems to lead to the skeptical conclusion that we hardly possess any knowledge at all. Some have suggested that knowledge requires only the capacity to rule out *relevant* alternatives (see, e.g., Dretske (1970)). At this stage, I do not want to make a judgment as to whether knowledge requires only the elimination of relevant alternatives or the elimination of all alternatives. But I think it is clear that our naive, untrained intuitions seem to regard that knowledge requires the elimination of all alternatives—at least, such intuitions are not hard to prompt.

It is harmless to assume, at least in cases where (A) and (B) are at issue, that S knows that the proposition *I have hands*[2] entails the proposition *I'm not a brain in a vat (being fed with stimuli that make me have the experience of looking at and making certain gesture with my hands)*. As a result, (C) can be taken as followed directly from (Closure).

(A), (B), and (C) are individually very plausible. Yet, any two of them imply the denial of the third: (A) and (B) imply the denial of (C); (A) and (C) imply the denial of (B); and (B) and (C) imply the denial of (A). (A), (B), and (C) thus exhibit a sense of paradox: They are individually very plausible and yet jointly inconsistent.

It is easy to see that the problem of (A) to (C) generalizes. Let 'S' refers to a normal cognizer, 'op' refers to *o*rdinary *p*ropositions about the external world that we normally take ourselves to know such as *I have hands*, *I am sitting in front of my laptop*, etc., 'sh' refers to *s*keptical *h*ypotheses that are incompatible with op such as *I am a handless brain in vat being manipulated to have a certain sensory experience*, *I am deceived by an evil demon to have a certain (non-veridical) sensory experience*, etc. By the same token, we can show that instances of the following propositional schemata are individually plausible:

1.1   S knows that op.

1.2   S does not know that not-sh.

1.3   S knows that op only if S knows that not-sh.

---

[2] Throughout this dissertation, I adopt the convention of using italics to stand for propositions.

Likewise, 1.1-1.3 are jointly inconsistent.[3] The problem of 1.1-1.3 has a great impact on ordinary knowledge claims about the world; it indicates that (almost) all our knowledge claims about the world are in a head-on clash with some very, if not equally, plausible statements. Call this problem 'the skeptical puzzle'.

The purpose of this dissertation is to tackle the skeptical puzzle. In this chapter, I will first discuss the most prominent approaches to the puzzle, casting doubt on some of them (Section 1.2-1-5). The most promising approach is what I will call 'the pluralistic solution' (Section 1.6-1.7). There are two main kinds of the pluralistic solution. The goal of this dissertation is to argue that one of them is the most promising solution to the skeptical puzzle (Section 1.8).

1.2 Approaches to the Skeptical Puzzle

How should we react to this puzzle? When facing a seemingly paradoxical set of statements, two general approaches suggest themselves. On the one hand, one may adopt the *concessionary approach*, admitting that the set of statements generates a genuine paradox—our ordinary intuitions are inherently incoherent, giving rise to the paradox. On the other hand, one may adopt the *resolving approach*, claiming that the set of statements does not actually constitute a genuine paradox—our ordinary intuitions are not inherently incoherent. The task of the resolving approach s to explain why, while our intuitions being coherent, the puzzle arises in the first place.[4]

---

[3] Strictly speaking, 1.1, 1.2, and 1.3 are propositional schemes, not propositions (or statements). Propositional schemes cannot be true or false; only their instances have truth-values. For simplicity's sake, however, I will continue to treat 1.1-1.3 as if they are propositions (or statements).

[4] Schiffer (1996) calls the concessionary and resolving approach the 'unhappy-face' and 'happy-face' solution respectively. He opts for the unhappy-face solution.

For the skeptical puzzle, the concessionary approach comes at a great cost. The concessionary approach indicates that the ordinary notion of knowledge gives rise to a genuine paradox. However, since, generally, any notion giving rise to genuine paradoxes should not be employed in the serious study of the subject matter (at least not without drastic revision), our ordinary notion of knowledge should thus be discarded in serious epistemological studies. Presumably, this would be unacceptable to traditional epistemologists who believe, in one way or another, that the notion of knowledge is central to epistemology. Moreover, if Timothy Williamson (2000) is right that the notion of knowledge plays a crucial and indispensable role in explicating other epistemic notions such as belief, justification, and evidence, the concessionary approach threatens the study of epistemic notions in general. From the viewpoint of traditional epistemology, the concessionary approach is thus a radical move that should not be adopted unless necessary. Luckily, we are not forced to endorse the concessionary approach, since it is far from being clear that the resolving approach to the skeptical puzzle, in general, is doomed to failure. At any rate, I will ignore the concessionary approach in this dissertation, and focus on the resolving approach instead.

The resolving approach preserves the ordinary notion of knowledge. Its task is to explain why the skeptical puzzle arises in the first place. Different accounts of the resolving approach offer different answers. Depending on whether the developed account denies (at least) one of our intuitions of 1.1-1.3, we can further classify different versions of the resolving approach into *conciliatory* and *repudiating accounts*. The conciliatory account takes our intuitions with respect to 1.1-1.3 as unproblematic; the initial plausibility of 1.1-1.3 is by no means mistaken. The task is to explain why, while our

intuitions with respect to 1.1-1.3 are both coherent and correct, the skeptical puzzle arises in the first place. The repudiating account rejects (at least) one of our intuitions with respect to 1.1-1.3; it takes the rejected intuitions(s) as mistaken. The task is to articulate which of our intuitions with respect to 1.1-1.3 should be rejected, and why (see Figure 1.1).

Is the skeptical puzzle a
genuine paradox?

YES

NO

The Concessionary
Approach

The Resolving
Approach

Are our intuitions
regarding (at least one
of) 1.1-1.3 mistaken?

YES

NO

The Repudiating
Account

The Conciliatory
Account
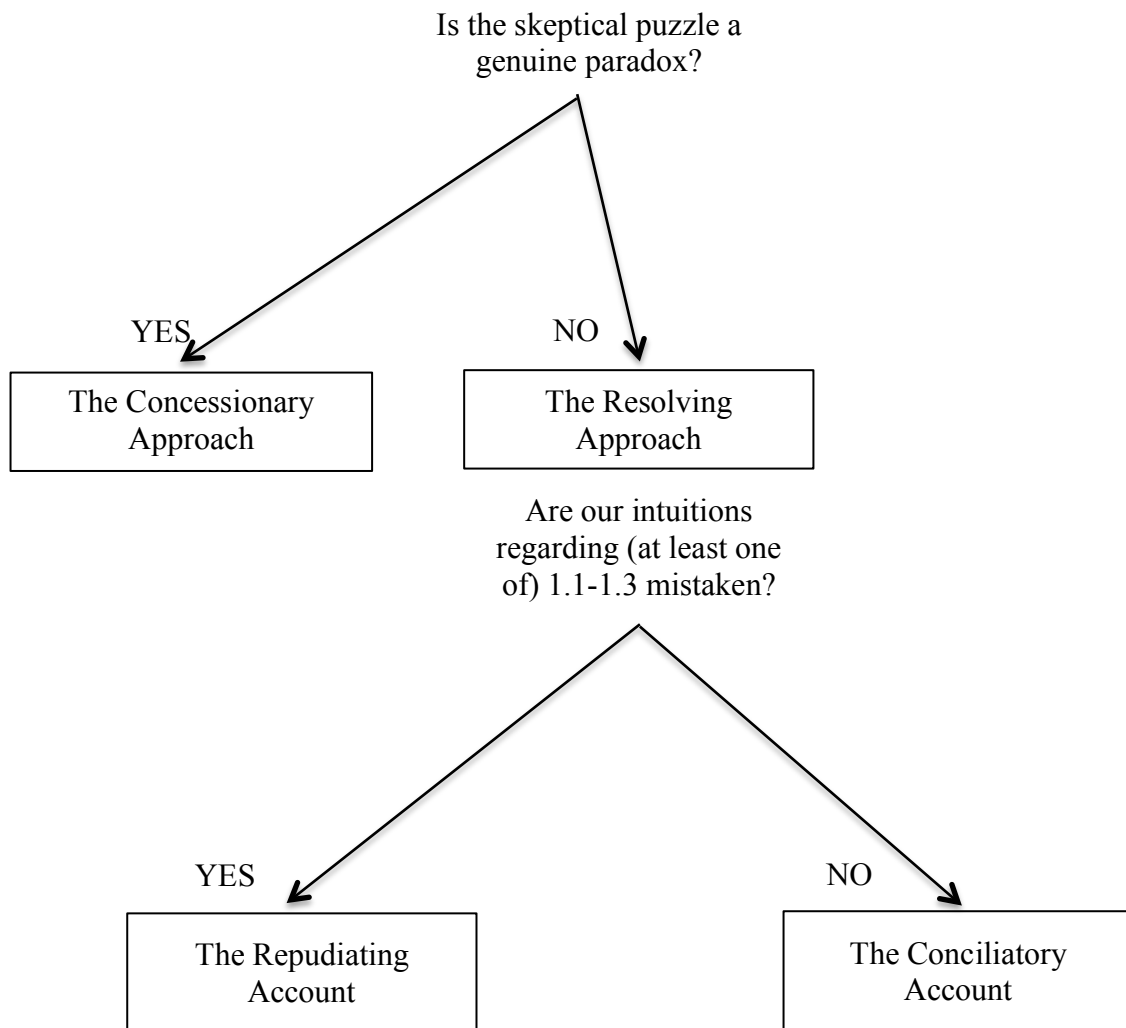
Figure 1.1
Approaches to the Skeptical
Puzzle

Both the repudiating and conciliatory accounts have been pursued. The goal of this dissertation is to develop a specific conciliatory account. Hence, I will briefly mention my qualms about the repudiating account.

1.3 The Repudiating Account

The simplest, though perhaps too naive, repudiating account is to deny our intuitions with respect to one of 1.1-1.3 solely on the basis of endorsing the rest. Call such an account 'the simple account'. There are three types of simple account: First, endorsing 1.1 and 1.3, one might argue that 1.2 was indeed false (by *Modus Ponens*), and thus that our intuitions with respect to 1.2 were mistaken. Second, endorsing 1.2 and 1.3, one might argue that 1.1 was indeed false (by *Modus Tollens*), and thus that our intuitions with respect to 1.1 were mistaken. Third, endorsing 1.1 and 1.2, one might argue that 1.3 was indeed false (by the definition of *Material Implication*), and thus that our intuitions with respect to 1.3 were mistaken.

The simple account is unsuccessful, and its failure illuminating. The simple account, on its bare form, claims that (at least) one of our intuitions with respect to 1.1-1.3 was mistaken since the corresponding statement is false. But this account fails to explain why the false statement appears intuitively plausible in the first place. To see why the simple account fails to offer the desired explanation, notice that the three types of simple account listed above are on a theoretical par with each other: Other things being equal, one is not reasonable to choose any one of them over the others; in the absence of further rationale, to prefer any one of them to the others is, theoretically speaking, arbitrary. The moral from the simple account is that proponents of the repudiating

account must explain not only (a) which of our intuitions with respect to 1.1-1.3 should be rejected and why, but also (b) why the corresponding statement appears initially plausible in the first place. Call the repudiating account that establishes (a) and (b) 'the sophisticated account'.

There are many sophisticated accounts on the market. Proponents of the sophisticated account disagree as to which of our intuitions with respect to 1.1-1.3 should be rejected. However, when it comes to answering (b), most (prominent) sophisticated accounts agree that the initial plausibility of the rejected intuition (or intuitions) of 1.1-1.3 is generated by our commitments to cognitive error of some sort. To discuss all sophisticated accounts is beyond the scope of the present study. In what follows, I will concentrate on the three prominent sophisticated accounts (Section 1.4). I will point out two general worries against them (Section 1.5).

1.4 Three Prominent Sophisticated Accounts

Peter Unger (1975) has proposed a sophisticated *skeptical* account, rejecting 1.1, i.e., S knows that op. Unger argues that knowledge requires (absolute) certainty—S knows that p only if S is (absolutely) certain of p (1975, 33ff.). For S, something is absolutely certain just in case "[n]o matter what any experience may *seem* to show or suggest as to whether or not something is so, [S] will now *reject as misleading* any experience which may seem to show or suggest that the thing is *not* so" (1975, 30). However, it is obvious that one cannot be absolutely certain of the falsity of sh. "One cannot help but think that," Unger writes, "for all [a subject] really can *know*, he might have all his experience artificially induced by electrodes, these being operated by a terribly evil scientist who, having an

idea of what his 'protégé' is saying to himself, chuckles accordingly" (1975, 25). Hence, 1.2 must be true.

Given 1.2 and 1.3, one must deny 1.1 (1975, 14-5). Unger takes the skeptical argument from 1.2 and 1.3 to not-1.1 to be "exceedingly compelling" (1975, 9). He then offers an explanation for why we normally endorse 1.1 (and reject 1.2): "we don't connect the [skeptical argument] with enough other things of interest to hold our prolonged attention, not even our attention in pursuing our philosophical interests. Lacking any larger connected discussion which favours scepticism, our sceptical inclinations do not remain active and we easily fall back on the comfortable, habitual thinking of 'common sense'. We easily return to think confidently that we know all sorts of things [about the external world]" (*op. cit.*). For Unger, however, falling back on such a habitual thinking of commonsense is nothing but returning to "ways of continued error" (*op. cit.*). In brief, according to this view, the skeptical puzzle can be resolved by rejecting 1.1, since it is denied by 1.2 and 1.3. More importantly, the initial plausibility of 1.1 consisted in the 'fact' that one failed (or was inclined to fail) to keep track on the compelling skeptical argument, and fell (or was inclined to fall) back to erroneous commonsensical views.

Ernest Sosa (2000) has argued for a sophisticated *Moorean* account that rejects 1.2, i.e., S does not know that not-sh. He argues that the initial plausibility of 1.2 arises from our confusion over the necessary condition of knowledge attribution. On Sosa's view, roughly, *safety* is a necessary condition for knowledge attribution. That is, one is legitimate in attributing knowledge that p to an agent S only if S's belief that p is *safe*.

S's belief that p is safe if and only if for all nearby possible worlds[5] (or similar circumstances) in which S believes that p, p is true in such worlds. On this view, one's belief in not-sh is safe (i.e. not-1.2 passes the safety condition for knowledge attribution). For instance, in all nearby worlds in which S believe that S is not a brain in a vat, S is not a brain in a vat. Hence, the safety condition for knowledge attribution does not support the truth of 1.2.[6] On this line of thought, one might as well endorse the denial of 1.2 on the basis of 1.1 and 1.3. Sosa suggests that the initial plausibility of 1.2 arises from our confusion between *sensitivity* and safety. One's belief that p is *sensitive* if and only if in all nearby worlds (or similar circumstances) in which p is false, S does not believe that p in such worlds. Notice that one's belief in not-sh is not sensitive. For instance, in all close worlds in which I am a brain in a vat, I will still believe that I am not a brain in a vat. If sensitivity is a necessary condition for knowledge attribution, 1.2 will be true. Sosa suggests that the initial plausibility of 1.2, for one, might consist in the confusion in one's mind between sensitivity and safety. "Sensitivity being so similar to safety, so easy to confuse," Sosa remarks, "it is not surprise that one would find sensitivity so plausible, enough to mislead one into assessing as correct affirmations of [1.2]" (2000, 16). On this view, the skeptical puzzle arises from one's confusion over the necessary condition of knowledge attribution. Once we get things straight, the skeptical puzzle should be resolved.

---

[5] The notion of 'nearby possible worlds' is not Lewis (1979b) but Nozick (1981). See also Williamson (2000) and Sosa (2000). For the possible-worlds semantics of counterfactual conditional, see Chapter Six.

[6] Taking safety as a necessary condition for knowledge attribution does not imply that 1.2 is false. The point, rather, is that, since not-1.2 passes the safety condition (i.e., S's belief in not-sh is safe), this should give us reason to reject 1.2.

Fred Dretske (1970) has argued for a sophisticated account to reject 1.3, i.e., if S knows that op, then S knows that not-sh. For Dretske, 1.3 seems initially plausible since it is sanctioned by (Closure), which is a seemingly plausible epistemic principle:

(Closure)   If S knows that p and S knows that p implies q, then S knows that q.

Dretske endorses the so-called *relevant alternatives* account of knowledge. According to Dretske, "[t]o know that *x* is *A* is to know that *x* is *A* within a framework of relevant alternatives, *B, C,* and *D*" (1970, 1022). Dretske defines a relevant alternative as "an alternative that might have been realized in the existing circumstances if the actual state of affairs had not materialized" (1970, 1021). On Dretske's view, both 1.1 and 1.2 are true under the present account of knowledge. For instance, S knows that S has hands since S knows that S has *hands* as opposed to *prostheses*, *claws*, etc. Similarly, S fails to know that S is not a brain in a vat (being fed with stimuli that make S experience having hands), since S fails to know that S is *not a brain a vat (being fed with stimuli that make S experience having hands)* as opposed to S is *a brain in a vat (being fed with stimuli that make S experience having hands)*, etc.[7] Dretske thus argues that (Closure) must fail to hold provided that 1.1 and 1.2 are both true and (assuming) that S knows that op implies not-sh. On this view, rejecting 1.1 resolves the skeptical puzzle. The initial plausibility of 1.3 arises from our mistaken endorsement of (Closure), which was in fact incorrect.

---

[7] Dretske does not say much about how S fails to know that S is *not a brain a vat (being fed with stimuli that make S experience having hands)* as opposed to S is *a brain in a vat (being fed with stimuli that make S experience having hands)*, etc. At some point of his discussion of skepticism, Dretske asks us to think what reasons we have, what evidence we can produce in favor of not-sh (cf. 1970, 1016). He seems to think that we lack evidence in favor of not-sh since, when considering whether sh or not, our evidence we *had* for op has been effectively neutralized (*op. cit.*). He admits that sh may not be plausible. Nevertheless, he claims that "the question here is not whether this alternative [i.e., sh] is plausible … but whether *you know* that this alternative hypothesis is false. I don't think you do" (*op. cit.*).

1.5 Two General Worries

The aforementioned sophisticated accounts are highly contentious and are by no means exhaustive. However, I will not get into the subtle discussion and evaluation of the pros and cons of each of them. Nor will I deal with other sophisticated accounts that also rely on the idea that the skeptical puzzle arises from our commitment to cognitive error of some sort.

Instead, I want to mention two general worries about the very idea that the skeptical puzzle arises from our commitment to cognitive error of some sort. First, it is hard to pin down what cognitive error we have committed ourselves to (if at all). One should not expect the skeptical puzzle to consist in an obvious cognitive error such that general agreement can simply be reached by pointing out the error in play. In fact, for many, the skeptical puzzle contains no obvious mistake (cf. Greco 2000). Hence, if the skeptical puzzle resides in cognitive error at all, the cognitive error in play must be theoretical—errors based on certain theoretical contentions. The discussion of the aforementioned sophisticated accounts bears witness to this point: Unger's rejection of the plausibility of 1.1 presupposes knowledge as requiring certainty, Sosa's rejection of the plausibility of 1.2 presupposes the safety condition of knowledge attribution, and Dretske's rejection of the plausibility of 1.3 presupposes the relevant-alternatives account of knowledge. Bringing in more theoretical contentions seems to further muddy the water. We are now facing a large amount of disagreement about how the skeptical puzzle should be diagnosed, what reasons we have against some proposals, what reasons to reject some rejections of a certain proposal, etc.

Some might reply that theoretical disagreement is inevitable in theoretical investigation in general. But the situation here is slightly different since all theoretical debates as to which cognitive error we are committed to seem to be on the wrong track in the first place. The point is that the initial plausibility of 1.1-1.3 is very compelling and remains to be so even *after* careful reflection. As Moore famously points out, we are very certain that, say, this is a pencil, and, more importantly, we are more certain that this is a pencil than any (possible) theoretical reasons that suggest the opposite (Moore 1962b, 30).

By contrast, we cannot have been certain that, say, we are not a brain in a vat inasmuch the same way as we are certain about *This is a pencil*, at least not in the sense of certainty defined by Unger (Unger 1975, 26). Is it not *obvious* that we are, in some sense, not in a position to know whether we are a brain in a vat? How can one know such kind of thing in the first place? It seems that when understood in an appropriate way, skeptical hypotheses are simply beyond the realm of knowable. It is thus very suspicious that some theoretical gimmicks are able to bring us knowledge of them.

Likewise for 1.3. Even after reflection, people still feel that the denial of 1.3 (or (Closure)), very implausible. Richard Feldman, for instance, has claimed that "the idea that no version of [closure] is true strikes [him], and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years" (Feldman 1995, 487).

To sum up, when it comes to the skeptical puzzle, it is implausible to simply reject (one of) our intuitions regarding 1.1-1.3 as fallacious solely on the basis of some theoretical considerations, since their initial plausibility arguably sustains (serious)

reflection.[8] This brings us to my second worry about the idea that the skeptical puzzle arises from our commitment to cognitive error of some sort.

That is, it seems that the initial plausibility of 1.1-1.3 can hardly be explained away by appealing to the idea that we commit ourselves to certain cognitive errors. To my knowledge, we still lack a satisfactory account that can *fully* explain away the initial plausibility of (at least) one of 1.1-1.3. Again, the discussion of the aforementioned sophisticated accounts bears witness to this point. Consider, firstly, Unger's contention that 1.1 is plausible since we have failed to keep track on the skeptical argument (the argument from 1.2 and 1.3 to not-1.1), and thus fallen back to the mistaken commonsensical view. Against Unger, it seems that 1.1 is still plausible even if we consciously focus on the skeptical argument. Indeed, it is a common reaction to the skeptical argument that it must be unsound since its conclusion is *obviously* false (cf. Pollock and Cruz 1999, 8), or at least, 1.1 appears to be more plausible than any of the premises of the skeptical argument does (cf. Moore 1962b, 30). Considering the skeptical argument may raise puzzle in our minds, but it surely does not render 1.1 implausible.

Secondly, Sosa's contention that the initial plausibility of 1.2 arises from the confusion between sensitivity and safety does not seem to deprive 1.2 of its plausibility either. On the one hand, even if we grant that the safety condition is a necessary condition for knowledge attribution, still, it does not automatically follow that 1.2 is false—the safety condition as a necessary condition for knowledge attribution does not imply not-1.2, i.e., S knows that not-sh. On the other hand, it seems that the plausibility of the idea

---

[8] My personal opinion is that our intuitions regarding 1.3 are the weaker in comparison to others. Hence, if forced, I am willing to give up 1.3 at the cost of preserving 1.1 and 1.2. However, in this dissertation, I will pursue a different route, a route that allows us to keep all of them.

that knowledge attribution requires sensitivity is derived directly from considering some ordinary cases of knowledge attribution. As Robert Nozick (1981) points out, S's knowing that p seems to require S's belief in p to track the truth of p—for Nozick, S's belief in p tracks the truth of p if and only if S would believe in p if and only if p were true.[9] Hence, since the sensitivity condition implies 1.2, it seems that the plausibility of 1.2 is derived directly from some ordinary cases of knowledge attribution, rather than from confusing safety with sensitivity, *contra* Sosa.

Thirdly, Dretske's denial of the plausibility of 1.3 fares no better than Unger's and Sosa's accounts. As many have argued that the denial of (Closure) is implausible and at any rate not necessary. On the one hand, many have found that (Closure), or the idea that deduction can extend one's knowledge, is intuitively plausible (cf. Feldman 1995; Williamson 2000). Moreover, as Keith DeRose (1995) points out, the denial of 1.3 is implausible in its own right; the following conjunction, which is an instance the denial of (Closure), is 'abominable': "while you don't know you're not a bodiless (and handless!) [brain in a vat], still, you know you have hands" (DeRose 1995, 28). On the other hand, as Gail Stine (1976) correctly points out, it is not necessary for someone endorsing the relevant alternatives account of knowledge to deny (Closure), since (Closure) is indeed compatible with the relevant alternatives account of knowledge.

The aforementioned accounts are not exhaustive. Even if what has been said is correct, we still lack conclusive reasons against the repudiating account in general.

---

[9] Nozick's tracking account of knowledge has encountered serious objections. However, I think that the tracking account of knowledge is worth developing, since those objections to Nozick's account can be properly replied by a better tracking account. For powerful recent defenses of the tracking account of knowledge, see Roush (2005) and Zalabardo (2012). In this dissertation, I will also argue for a tracking account of knowledge. See Chapter Four, Five and Six.

Nevertheless, the fact that these prominent accounts have failed to explain the initial plausibility of 1.1-1.3 as residing in our commitment of cognitive error of some sort should give us some reasons to doubt the prospect for such an account. Hence, lacking a theoretically tenable repudiating account, I think it is reasonable for us to move on to the conciliatory account.

In my opinion, the conciliatory account—the idea that all of our intuitions with respect to 1.1-1.3 are unproblematic—is a more promising approach. In the next section, I discuss a specific way of developing the conciliatory account. This discussion will serve as the basis for my own resolution to the skeptical puzzle, which is a version of the conciliatory account.

1.6 The Contextualist Solution to the Skeptical Puzzle

As noted, the conciliatory account consists of two theses: First, the skeptical puzzle does not constitute a genuine paradox, and second, the initial plausibility of 1.1-1.3 is unproblematic.

There are ways to further develop these two theses. In this section, I will first consider the most influential conciliatory account on the market, i.e., *(epistemic) contextualism*. My purpose is to isolate the genetic structure of the type of conciliatory account of which contextualism is an instance (see Section 1.7). The merit of isolating such a structure is that it allows us to see clearly the main variants of this type of conciliatory account. This is of particular importance for the present study since I believe that among those variants lies the most promising solution to the skeptical puzzle, which,

unfortunately, has been widely ignored, if not due to the fact that it has seldom been explicitly formulated.

On contextualism, the ordinary nation of knowledge is not inherently incoherent. Moreover, and more importantly, our initial intuitions regarding 1.1-1.3 are unproblematic since whenever we intuitively think them to be true, they are indeed true. Contextualists' justification for this last claim consists of three parts.

First, contextualists have argued for an innovative (some may consider it radical) *semantic* thesis about knowledge attributions. They argue that the term 'knowledge' (and its cognates) is a context-sensitive term just like indexicals such as 'I', 'here', etc. or gradable adjectives, such as 'tall', 'flat', etc. Just like the content of 'I' is sensitive to the context the term is used, the content of 'knowledge' is sensitive to the context the term is used. The variation of the content of 'knowledge' may give rise to the variation of the proposition expressed by a knowledge attributing sentence, say, 'S knows that p', leading eventually to the variation of the truth value of the knowledge attribution "S knows that p".

Second, contextualists typically interpret the context-sensitivity of 'knowledge' as indicating that the *standards for knowledge* are sensitive to the speaker context.[10] The

---

[10] What makes the standard of application of 'knowledge' shift from one context to another? Contextualists generally take it that the speaker's practical or conversational considerations, or the pragmatic factors in the context of use of 'knowledge', alone are sufficient to bring out the shift of the standard. For instance, Stewart Cohen once suggests that the standard of application of 'knowledge' is determined by "some complicated function of speaker intentions, listener expectations, presuppositions of the conversation, salience relations, etc., —by what David Lewis calls the conversational score" (1999, 61).

Contextualists are aware of the complexity and difficulty of explaining how exactly such considerations or factors shift the standard of application of 'knowledge' but they do not pretend to have a satisfactory account on that score. Some thought that this was not be a serious problem for contextualism, at least not a problem specifically of the theory, since, arguably, we lack a satisfactory account of exactly how pragmatic factors affect the extensions of context-sensitive terms in general (cf. Feldman 1999). But I think the issue is much complicated here. True, contextualists may not need to offer an account of context-sensitivity in general. But this does not mean that they need to not offer an account of how the standards for

idea is that, given that one believes truly that p, one counts as knowing that p only if *one's standing on knowing dimensions* with respect to p is *good* enough. The standards for knowledge determine how good one's standing must be in order to be good enough for knowledge, and they are, according to contextualism, sensitive to the speaker context.

Third, contextualists argue that the context-sensitivity of the standards for knowledge, so explicated, provide us with necessary resources to resolve the skeptical puzzle. Contextualists distinguish what we may call 'the ordinary context' from 'the skeptical context'. The ordinary context is the daily, normal circumstance in which philosophical skepticism is not seriously pondered (or is ignored altogether). Nor are skeptical hypotheses salient in this context. The skeptical context, by contrast, is where people are pondering over philosophical skepticism, and, as a result, skeptical hypotheses become salient.

Contextualism argues that the difference in pragmatic factors such as whether the skeptical hypothesis is salient between the ordinary and the skeptical contexts gives rise to the variability of the standards for knowledge: The standards for knowledge in the ordinary context are *moderate* such that S's standing on knowing dimensions with respect to p does not need to be very strong in order for S to count as knowing that p. By contrast, the standards for knowledge are extremely *stringent* in the skeptical context such that the strength of S's standing on knowing dimensions with regard to p has to be very strong in order for S to know that p.

Contextualists argue that our standing on knowing dimensions with regard to op (and not-sh), is generally good enough for knowing op (and not-sh) in the ordinary

knowledge are shifted in different contexts. To the contrary, lacking such an account, we cannot even try to evaluate whether 'knowledge' is context-sensitive or not (see Chapter Two).

context. That is, both 1.1 and not-1.2 are true in the ordinary context. By contrast, our moderate standing on knowing dimensions with regard to op (and not-sh) is not good enough for knowing op (and not-sh) in the skeptical context. As a result, both 1.1 and not-1.2 are false in the skeptical context.

Contextualists have typically proposed a conciliatory account of the skeptical puzzle. Contextualists claim that while our intuitions regarding 1.1-1.3 are unproblematic, they do not render 1.1-1.3 genuinely paradoxical. The idea is that, firstly, we initially find 1.1 to be plausible since we normally evaluate the truth-value of 1.1 in the ordinary context, and in this context, 1.1 is indeed true. 1.1-1.3 do not constitute genuine a paradox in the ordinary context since 1.2 is false in this context. Secondly, we initially find 1.2 to be plausible since considering skeptical hypotheses tend to put us in the skeptical context, and 1.2 is indeed true in such a context. Likewise, 1.1-1.3 do not constitute a genuine paradox in the skeptical context since 1.1 is false in this context. Thirdly, we initially find 1.3 to be plausible since it is in fact true in all contexts.[11]

Put differently, according to contextualism, 1.1 is plausible since it is true in the ordinary context; 1.2 is plausible since it is true in the skeptical context; 1.3 is plausible since it is true in all contexts; 1.1-1.3 do not constitute a genuine paradox since in no context will we find all 1.1-1.3 to be intuitively plausible *and* true.

The sense of paradox of 1.1-1.3 is generated by an unnoticed equivocation of the propositions expressed by the same knowledge-attributing sentence used in the ordinary and skeptical contexts. It is because we are somehow blind to fact that different

---

[11] As noted, many have thought that 1.3 is supported by (Closure). Although many contextualists accept (Closure), the denial of (Closure) is not incompatible with contextualism. Heller (Heller 1999a; 1999b), for instance, has argued for a version of contextualism that denies (Closure). For simplicity's sake, however, I will assume that contextualism accepts (Closure).

propositions have been expressed by 1.1 (and 1.3) in the ordinary and skeptical contexts that we fail to realize that 1.1-1.3 are not plausible in the same context.[12]

The cogency of the contextualist solution to the skeptical puzzle is debatable (see Chapter Three). My present goal is to isolate a generic structure of the type of conciliatory account of which the contextualist solution is an instance. Such an endeavor will be rewarding since it will greatly clarify the conceptual territory and, more importantly, crystalize a widely disregarded alternative account to the skeptical puzzle, which, I believe, is most promising.

To facilitate the following discussion, let me formulate the contextualist solution's two theses:

> *The contextualist theory (of knowledge attribution)*. The variability of the standards for knowledge across contexts, due to the difference in pragmatic factors, gives rise to the variability of content of 'know' (and its cognates), and thus results in the variability of the truth-values of the same knowledge attribution in such contexts.

> *The contextualist solution*. The ordinary and skeptical contexts differ in some pragmatic factors. Given contextualism, it follows that the standards for knowledge vary in such a way that S's standing on knowing dimensions with respect to op, and not-sh, satisfies the standards for knowledge in play in the ordinary context but not the one in the skeptical context. As a result, (1) and not-(2) are true in the ordinary context but false in the skeptical one. (3), by contrast, is true in both contexts.

---

[12] The contextualist solution to the skeptical puzzle (and contextualism in general) has many theoretical merits. Firstly, it adopts the resolving approach to the puzzle that avoids taking the notion of knowledge as inherently deficit. Second, it is a conciliatory account, which, unlike the repudiating account, does not claim that the skeptical puzzle arises from our commitment of cognitive error of some sort—the initial plausibility of 1.1-1.3 consists in 1.1-1.3 being true (at least in some contexts). Thirdly, since 1.3 is widely regarded as supported by (Closure), which many have found plausible, the contextualist solution, by admitting the truth of 1.3 in all contexts, is in a position to respect (Closure). Fourthly, it is claimed that the contextualist thesis of the context-sensitivity of the standard of application of 'knowledge' can account for some puzzling cases according to which the same knowledge-attributing sentence seems to have different truth values in different contexts, even though the subject's standing on knowing dimensions appears to remain constant across such contexts—such cases are claimed to be incomprehensible without assuming the context-sensitivity of the standards for knowledge (cf. Chapter Two)

1.7 The Pluralistic Solution to the Skeptical Puzzle

The aforementioned is only the bare bone of the contextualist solution to the skeptical puzzle. More needs to be said.[13] For the present purposes, I will ignore such discussions, and focus solely on the general structure of the contextualist solution in order to get at the generic approach to the skeptical puzzle of which the contextualist solution is an instance.

In order to clarify the conceptual territory, it is useful to introduce some terminology. Throughout this dissertation, I will adopt the orthodox picture of how to determine the truth value of an utterance, which is derived primarily from David Kaplan (1989).

Utterances are taken to be sentences uttered in the *context of use*. In general, the truth value of an utterance "p" in a context C is determined by two factors: the *intension* of the expression 'p' in C and a *circumstance of evaluation*, which is determined or specified by the context of use. Traditionally, the circumstance of evaluation is regarded as consisting of a possible world (and a time). The intension of an expression is the *semantic content* expressed by the expression. The intension of a complex expression is determined by the intensions of its parts and certain *rules of combination*. The intension of an utterance can be regarded as its truth-condition—roughly, the condition that determines whether the utterance is true or false. More generally, intensions can be regarded as functions from circumstances of evaluation to *extensions*. The extension of

---

[13] For instance, it is crucial for contextualists to specify, and justify, the (semantic) mechanism for the variability of the content of 'knowledge' across contexts. Also, contextualists need to explain why and how the standard of application of 'knowledge' varies with the ordinary and skeptical contexts. Contextualists also need to justify their claim that the initial sense of paradox of 1.1-1.3 consists in our ignorance of the variability of the standard of application of 'knowledge' (or the variability of the content of 'knowledge') across the ordinary and skeptical contexts. Many proposals have been offered with regard to each of these questions (and to many others), and more discussions and objections ensue.

an utterance is its truth value. The extension of a proper name is an individual, and the extension of a predicate is a set, whose members may be individuals, ordered pairs, or sets. The extension of a complex expression is determined by the extensions of its parts and certain rule of combination.[14]

According to the orthodox view, to say that a knowledge attribution, say, "S knows that p", has different truth values across two contexts C1 and C2 is tantamount to saying that the extension of at least one of the expressions 'S', 'knows' and 'p' varies across C1 and C2, since the extension of a sentence depends on the extensions of its parts. On contextualism, even if the extensions of 'S' and 'p' remain constant across C1 and C2, the extension of 'knows' may still vary.

The extension of 'knowledge' depends on its intension and the circumstance of evaluation. Contextualists typically claim that the intension of 'knowledge' may vary with contexts even holding fixed the circumstance of evaluation. In other words, contextualists contend the variability of the intention of 'knowledge' across contexts.

How should we characterize the intention of 'knowledge'? Intuitively speaking, the intention of 'knowledge' specifies the condition for knowledge. According to contextualists, one's true belief that p counts as knowledge only if one's standing on knowing dimensions with respect to p is good enough. The standards for knowledge determine how good one's standing must be in order for one possesses knowledge.

More precisely, we may characterize the intension of 'knowledge' as follows (cf., e.g., MacFarlane 2009, 236):

---

[14] The orthodox picture is supposed to help us to clarify some views. There are some contentious claims related to this picture that I do not endorse. In particular, it has been argued that the intension of an utterance is what is said by the utterance in a context. Many people disagree, and I think we don't need to take a stance in this debate. All we need is to use this framework to illustrate some positions.

(K) "S knows that p" is true in C1 if and only if

    (i)   S believes that p in C2[15],

    (ii)   p is true in C2, and

    (iii)   S's standing on knowing dimensions with respect to p in C2 satisfies the standards for knowledge in play.

(Note: C1 and C2 does not have to be identical)

Some clarifications are called for. First, siding with the orthodox picture, (K) implies that knowledge implies true belief.[16]

Second, 'S's standing on knowing dimensions' is simply used as a *placeholder* here, standing for any condition that turns true belief into knowledge. [17] As it turns out, many think that one's standing on knowing dimensions with respect to p should be determined by one's standing on truth-relevant dimensions with respect to p (call it 'one's *epistemic position* with respect to p'). However, this notion of one's standing on knowing dimensions is too *narrow*. Whether one's true belief that p counts as knowledge or not does not only depend on one's epistemic position with respect to p. It also depends on whether one is epistemically responsible for believing that p (cf. Fogelin 1994; Goldman 1988; Williams 2001; also see Section 4.4). Since one's standing on doxastic-responsibility dimensions with respect to p does not, strictly, counts as one's epistemic position with respect to p, it is strictly false to say that one's standing on knowing dimensions with respect to p should be determined solely by one's epistemic position

---

[15] Strictly speaking, (K)(i) should specify that S believes that p in C1 on the basis of (iii) (this is sometimes called 'the basing condition').

[16] Some may claim that believing is not the appropriate propositional attitude for knowledge. Be that as it may, the following discussion does not hinge on the claim that knowing requiring believing. In principle, one can replace believing with other propositional attitudes without significantly affecting the following discussion.

[17] I am here adhering to another orthodox idea that knowledge is more than truth belief.

with respect to p. However, throughout this dissertation, when talking about the variation of one's standing on knowing dimensions with respect to p across contexts, I mean the variation of one's epistemic position with respect to p (unless specially specified). The reason is that most of the cases we are going to discuss have held fixed the subject's standing on doxastic-responsibility dimensions with respect to p.

Third, I take one's standing on knowing dimensions (or one's epistemic position) with respect to p to come in degrees. We can talk about the strength of one's standing on knowing dimensions (or one's epistemic position) with respect to p and one's standing on knowing dimensions (or epistemic position) being stronger or weaker than another's. I leave it open, at least at this stage, how the strength of one's standing on knowing dimensions (or epistemic position) with respect to p is to be measured (see Chapter Four).

Fourth, the standards for knowledge determine how strong one's standing on knowing dimension with respect to p must be in order for one to be truthfully described as "knows p" (or to fall into the extension of knowing p). *Standards variantists* argue that the standards for knowledge vary with contexts while *standards invariantists* deny it. To remain neutral to the issue, I use the expression 'the standards for knowledge *in play*'.

Let us get back to (K). According to (K), to say that "S knows that p" is true in C1 but false in C2 amounts to saying that at least one of (K)(i) to (K)(iii) must fail to obtain in C2, holding fixed the extensions of 'S' and 'p'. It should be obvious that (K)(i) and (K)(ii) are not relevant to the skeptical puzzle: Even if S believes truly that not-sh, 1.2 (i.e., S does not know that not-sh) remains plausible. Hence, as I see it, contextualism is correct in claiming that it is (K)(iii) that is pertinent to the skeptical puzzle. On contextualism, "S knows that op" (or "S knows that not-sh") is true in the ordinary

24

context but not in the skeptical one, since (K)(iii) obtains—S's standing on knowing dimensions with respect to op (or not-sh) satisfies the standards for knowledge in play—in the ordinary context but not in the skeptical one. In other words, the contextualists' insight is that the extension of 'knowledge' (and its cognates) varies with contexts due to the difference in pragmatic factors in such contexts, and that the variability of the extension of 'knowledge' consists in the fact that S's standing on knowing dimensions satisfies the standards for knowledge in play in some contexts but not in the others (i.e., (K)(iii) obtains in some contexts but not in the others).

If S's standing on knowing dimensions with respect to p satisfies the standards for knowledge in play in C1 but not in C2, then one of the following cases must hold:

The standards-variantist thesis. The standards for knowledge in play vary with C1 and C2.

The standing-variantist thesis. S's standing on knowing dimensions with respect to p varies with C1 and C2

The general-variantist thesis. Both the standards for knowledge in play and S's standing on knowing dimensions with respect to p vary with C1 and C2.

For simplicity's sake, I will not deal with the general-variantist thesis, since this thesis is theoretically less simple than the standards- and standing-variantist theses, and arguably, other things being equal, we should prefer the simpler account.

Contextualism opts for the standards-variantist thesis; it claims that pragmatic factors alone are sometimes enough to shift the standards for knowledge in play. The contextualist theory is thus an instance of the following general approach:

Standards variantism. The variability of the standards for knowledge across contexts, due to the difference in pragmatic factors, gives rise to the variability of the extension of 'knowledge' (and its cognates), and thus the variability of the truth values of the same knowledge-attributing sentence in such contexts, other things being equal.

Standards variantism is the foundation of the contextualist solution to the skeptical puzzle. On contextualism, the ordinary and skeptical contexts differ only in some pragmatic aspects such as some skeptical hypotheses are salient in the skeptical context but not in the ordinary one. Contextualists argue that the pragmatic difference in these two contexts gives rise to the variability of the standards for knowledge in play. In the ordinary context, the standards for knowledge in play are moderate such that they are satisfied by S's moderate strength of standing on knowing dimensions with respect to op (and not-sh). In the skeptical context, the standards for knowledge in play are very stringent such that they are not satisfied by S's moderate strength of standing on knowing dimensions with respect to op (and not-sh). In other words, the contextualist solution of the skeptical puzzle is a special version of what we may call the *standards-variantist solution (to the skeptical puzzle)*:

> *The standards-variantist solution.* The ordinary and skeptical contexts differ in some pragmatic factors. Given standards variantism, it follows that the standards for knowledge in play in the ordinary and skeptical contexts vary in such a way that S's standing on knowing dimensions with respect to op, and not-sh, satisfies the standards for knowledge in play in the ordinary context but not the one in the skeptical context. As a result, 1.1 and not-1.2 are true in the ordinary context but not in the skeptical one; 1.3, by contrast, is true in both contexts.

In sum, the contextualist theory is a version of standards variantism, and the contextualist solution is a version of the standards-variantist solution. The contextualist solution presupposes the contextualist theory (but not the other way around); the standards-variantist solution presupposes standards variantism (but not the other way around).

We can formulate a even more general account of which standards variantism is an instance:

> *Epistemic pluralism.* The extension of 'knowledge' (and its cognates) varies with contexts, due to the difference in pragmatic factors, and the variability of the

extension of 'knowledge' gives rise to the variability of the truth-value of the same knowledge-attributing sentence in such contexts, other things being equal.

In the contextualist spirit, one can thus argue for an epistemic-pluralistic solution to the skeptical puzzle:

> *The pluralistic solution*. The ordinary and skeptical contexts differ in some pragmatic factors. Given epistemic pluralism, it follows that the extension of 'knowledge' (and its cognates), in the ordinary and skeptical contexts varies in such a way that S's standing on knowing dimensions with respect to op, and not-sh, satisfies the standards for knowledge in play in the ordinary context but not the one in the skeptical context. As a result, 1.1 and not-1.2 are true in the ordinary context but not in the skeptical one; 1.3, by contrast, is true in both contexts.

Obviously, standards variantism (and hence contextualism) is a version of epistemic pluralism, and the standards-variantist (and hence the contextualist) solution is a version of the pluralistic solution. Just like the standards-variantist solution presupposes standards variantism, the pluralistic solution presupposes epistemic pluralism.

Similarly, just like contextualism does not presuppose the contextualist solution to the skeptical puzzle, epistemic pluralism does not presuppose the pluralistic solution. Whether epistemic pluralism is correct—whether the extension of 'knowledge' (and its cognates) varies with contexts due to the difference in pragmatic factors in such contexts—is itself of great theoretical importance. A brief survey of the current literatures reveals that a heated debate has been concerned with epistemic pluralism in general (and standards variantism in particular).[18] But it is not mandatory for advocates of epistemic pluralism to endorse the pluralistic solution to the skeptical puzzle. Nor is it necessary that they have to take a stance on the skeptical puzzle.

---

[18] See the debate between indexical contextualism (e.g. Cohen 1999; DeRose 2009; Lewis 1996; Blome-Tillmann 2009a), nonindexical contextualism (e.g. Brogaard 2008), and relativism (MacFarlane 2005; 2009). There are also invariantists who deny the variability of the standards for knowledge due merely to the variation of pragmatic factors of some sort (cf., e.g., Williamson 2000; Stanley 2005; Hawthorne 2004).

The articulation of the pluralistic solution has great clarificatory merits. It allows us to see clearly an alternative solution to the skeptical puzzle, which unfortunately has been widely ignored or disregarded. Recall that there are three ways the extension of 'knowledge' (and its cognates) may vary with contexts. That is, if S's standing on knowing dimensions with respect to p satisfies the standards for knowledge in play in C1 but not in C2, then either the standards for knowledge vary across C1 and C2, or S's standing on knowing dimensions with respect to p varies across C1 and C2, or both. Standards variantism opts for the first option. An alternative is to opt for the second:

> *Standing variantism*. The variability of S's standing on knowing dimensions with respect to p across contexts, due to the difference in pragmatic factors, gives rise to the variantist of the extension of 'knowledge' (and its cognates), and thus the variability of the truth value of the same knowledge-attributing sentence in such contexts, other things being equal.

By the same token, we have the standing-variantist solution to the skeptical puzzle:

> *The standing-variantist solution*. The ordinary and skeptical contexts differ in some pragmatic factors. Given standing variantism, S's standing on knowing dimensions with respect to op, and not-sh, varies in such a way that S's standing on knowing dimensions with respect to op, and not-sh, satisfies the standards for knowledge in play in the ordinary context but not the one in the skeptical context. As a result, 1.1 and not-1.2 are true in the ordinary context but not in the skeptical one; 1.3, by contrast, is true in both contexts.

Likewise, although proponents of the standing-variantist solution endorse standing variantism, proponents of standing variantism need not endorse the standing-variantist solution.[19]

---

[19] The so-called *subject sensitive invariantism* (SSI) is also a form of standing variabilism (cf. Hawthorne 2004; Stanley 2005; Fantl and McGrath 2009a). According to this view, the mere variation of one's practical considerations such as stakes may give rise to the variation of one's standing on knowing dimensions. But, to my knowledge, no proponent of SSI has explicitly adopted the standing-variabilist solution to the skeptical puzzle.

To sum up, both the standards- and standing-variantist solutions share the same general structure, i.e., the pluralistic solution. The difference between them consists in what factors are taken as being responsible for the variability of the extension of 'knowledge' (and its cognates) across the ordinary and skeptical contexts. Proponents of the standards-variantist solution opt for the variability of the standards for knowledge, while advocates of the standing-variantist solution, the variability of S's standing on knowing dimensions.

The main contention of this dissertation is that a version of the standing-variantist solution offers the best solution to the skeptical puzzle. Since the standing-variantist solution presupposes standing variantism, I also argue for a version of standing variantism. Few proponents of standing variantism have explicitly argued for the standing-variantist solution to the skeptical puzzle.

1.8 Sensitivity and Standing Variantism

The standing-variantist solution presupposes standing variantism. In this thesis, I argue that one's standing on knowing dimensions—more precisely, one's epistemic position—with respect to p is characterized in terms of the *sensitivity* of one's (sole) reason R for p. S's reason for p is some experiential state of S, referring to S's "evidence, grounds, or reasons" for knowing that p (Dretske 1971a, 4). S's reason R for p is *sensitive* if and only if S would have lacked R if p were not the case.[20]

---

[20] The present formulation takes sensitivity to be a property of one's reason. This, however, is not the dominant formulation. It is common to take sensitivity to be a property of one's belief: S's belief that p is sensitive if and only if S would not have believed that p if p were false. In Chapter Four, I argue that taking sensitivity as a property of belief conflates the question concerning epistemic position and the question concerning doxastic responsibility (Section 4.4).

Some have argued that sensitivity is a necessary condition for knowledge (e.g., Nozick 1981; Dretske 1970). The underlying idea is that knowing that p implies that having reason that indicates the truth of p, which could not be easily mistaken. If I know that the cat is on the mat, I must believe that the cat is on the mat on the basis of some reason sufficiently good for indicating that the cat is on the mat (and nothing else). Sensitivity captures this sense of good reason. If one's reason R for p is sensitive, then R indicates the truth of p in the sense that if p were false, one would not have the same R for p. In other words, a sensitive reason for p is a reason sensitive to the truth (falsity) of p.

Given that sensitivity is a necessary condition of knowledge, it is not hard to see that the skeptic's is right in claiming that S does not know that not-sh. The idea, roughly, is that if sh is the case, say, if S is a brain in a vat, then S would have exactly the same experiences as she does in the actual situation. Hence, no matter what S's reason for not-sh is, S would still have the same reason if she were a brain in a vat. Hence, S's reason for not-sh is not sensitive. As a result, S does not know that not-sh; our intuitions regarding 1.2 are correct. By the same token, S's reason for op, say, *S has hands*, is generally sensitive. Ordinarily, S's reason for *S has hands* includes S's (perceptual) experiences of having hands. If S did not have hands—for instance, if S's hands were chopped off—S would not have the same experiences of having hands. Hence, S's reason for *S has hands* is generally sensitive. This explains why the ordinary intuition that S knows that op is also correct (i.e., 1.1).

My account of knowledge preserves epistemic closure. Following Sherrilyn Roush (2005), I take sensitivity to be the basic feature for the recursive definition of

knowledge. Roughly, if one's reason for p is sensitive, S knows that p provided other conditions for knowledge are also satisfied. If one knows that p, then one also knows that q if one competently infers q from p. The set of known propositions are recursively defined by the set of propositions the reason for which is sensitive. This explains our intuitions regarding 1.3.

Although the skeptic is right in claiming that S does not know that not-sh and the ordinary people are also right in claiming that S knows that op, they are not right in the same context. The skeptic and the ordinary people are engaging in different kinds of *epistemic inquiries* (i.e., roughly, inquiries that aim at pursuing truth and avoiding error). The ordinary people engage in what we may call the *ordinary-world evaluation* which has presupposed the falsity of sh. The skeptic, by contrast, engages in the *non-ordinary-world evaluation* which has focused on what would have happened if sh is the case. What is important is that the measurement of the sensitivity of one's reason depends on the inquiry one—as an evaluator—is engaging in. More precisely, the statement "S would have lacked R if p were false" might have different truth values provided that different inquiries are in play.

1.9 Taking Stock

We have seen that the skeptical puzzle can be solved or resolved by many approaches. The concessionary approach is not particularly attractive since it seems to stand in tension with traditional epistemology. The resolving approach is more promising on that score. One way to develop the resolving approach is to reject one of our intuitions concerning 1.1-1.3, i.e., the repudiating account. The simple repudiating account is

unpromising since it fails to explain why the rejected intuition appears plausible in the first place. The sophisticated repudiating account argues that we initially uphold the rejected intuition concerning 1.1-1.3 since we have committed to cognitive error of some sort. The sophisticated account faces two general worries. On the one hand, it implies that we have made some theoretical error in holding the rejected intuition concerning 1.1-1.3. Arguably, this implication leads to more controversies. On the other hand, our survey of three prominent sophisticated accounts suggests that our initial intuitions regarding 1.1-1.3 cannot be explained satisfactorily in terms of theoretical error of some sort.

If what has been said is correct, we are left with the conciliatory account. As noted, one way to develop the conciliatory account is to argue for the pluralistic solution (to the skeptical puzzle). The pluralistic solution can be further classified into the standards- and standing-variantist solutions. Which one should we choose? Some might argue that this question is trivial since there has been a well-developed standards-variantist solution to the skeptical puzzle on the table, i.e., the contextualist solution. Contextualists have argued that there are independent reasons to accept contextualism—the thesis that 'knowledge' (and its cognates) is context-sensitive—and the contextualist solution to the skeptical puzzle—the thesis that the context-sensitivity of 'knowledge' gives rise to a satisfactory solution to the skeptical puzzle. Why not be satisfied with the contextualist solution? Why bother with the standing-variantist solution? The answer, in my opinion, is straightforward: Because both contextualist theses are problematic. The next two chapters are dedicated to this claim.

In Chapter Two, I discuss contextualists' main reason for the context-sensitivity of 'knowledge' (or its cognates), namely, the idea that some cases of ordinary knowledge

attributions directly supported the contextualist (and the standards-variantist) interpretation of 'knowledge'. As I will argue, contra contextualists (and standards-variantists), those cases in fact support standing-variantism—the thesis that one's standing on knowing dimensions varies with contexts due solely to pragmatic factors of some sort (see Chapter Seven).

In Chapter Three, I consider the contextualist solution (and standards-variantist solution in general) to the skeptical puzzle. I argue that the contextualist solution fails to account for the most promising skepticism, which is intended to be captured by the skeptical puzzle. By contrast, such kind of skepticism can be easily handled by the standing-variantist solution.

My main focus in Chapter Two and Three is contextualism, but my conclusion can be easily extended to standards variantism in general. Hence, it seems that we have a good reason not to be satisfied with the contextualist solution (and, in general, the standards-variantist solution). However, some might have qualms about the standing-variantist solution since they thought that standing variantism, which is assumed by the standing-variantist solution, was problematic on its own.

Chapter Two and Three serve two purposes. On the one hand, they explain why the standards-variantist solution (and standards variantism in general) is not promising in solving the skeptical puzzle. Hence, they help to motivate those sympathetic to the pluralistic solution to turn their heads to the standing-variantist solution (and standing variantism in general). On the other hand, the discussion of (the problems) of the standing-variantist solution also helps to set the stage for the standing-variantist solution.

By the end of the first three chapters, we should get some ideas of what a successful the standing-variantist solution should look like.

In Chapter Four, I start to characterize one's epistemic position with respect to p (or one's truth-relevant dimensions with respect to p), and construct an account of knowledge called 'the tracking-based account'. With respect to the former, I begin by investigating the more general notion of tracking. I argue that knowledge requires tracking. I then further argue that sensitivity is the kind of tracking required by knowledge. With respect to the latter, the notion of recursive definition is adopted. Sensitive beliefs are taken to be the base of a definition of knowledge.

The sensitivity account has suffered from serious objections in the literatures. In Chapter Five (and Chapter Six), I deal with the most prominent objections (or counterexamples) to the sensitivity account and find all of them wanting—at least, none of them seem to have a serious punch on the tracking-based account.

The sensitivity account I adopt characterizes sensitivity in terms of counterfactual conditionals, such as, "If p were false, S would not have had R". This prompts us to dig into the nature of counterfactual conditionals. In Chapter Six, I argue that epistemologists should adopt Judea Pearl's (2000) causal-modeling semantics of counterfactual conditionals as opposed to David Lewis's (1973) orthdox possible-worlds semantics.

In Chapter Seven, I apply the tracking-based account of knowledge to solve the skeptical puzzle. The basic idea is that the skeptic and the ordinary people are engaging in different kinds of epistemic inquiries. Since the evaluation of one's epistemic position depends on the kind of inquiry the evaluator is engaging, it follows that both the skeptic's

denials of knoweldge and the ordinary people's attributions of knowledge are correct in their corresponding contexts.

In Appendix 1, I then use the same idea to explain ordinary cases of knowledge attiributions. More precisely, I will account for the kind of cases (which will be called 'stakes-shifting cases') discussed in Chapter Two.

Chapter Two

Contextualism and Stakes-Shifting Cases

2.1 Introduction

The standards-variantist solution argues that the same knowledge attribution say "S knows that p" can be true in the ordinary context but fail to be the case in the skeptical one, due to the variation of the standards for knowledge in such contexts. It is important that proponents of the standards-variantist solution must also argue that the standards for knowledge, in general, vary with contexts, i.e., standards variantism, since if the standards for knowledge, in general, do not vary with contexts (in such a way that gives rise to the standards-variantist solution), the standards-variantist solution is simply *ad hoc* (cf. DeRose 2009). It is thus not surprising to see proponents of the standards-variantist solution are also proponents of standards variantism. My goal here is to consider and deny the strongest support (or one of the strongest supports) for standards variantism.

Central to standards variantism is the idea that the mere variation of pragmatic factors of some sort is sufficient for giving rise to the variation of the standards for knowledge. Many have argued that standards variantism (more precisely, contextualism) was directly supported by ordinary cases of knowledge attributions of the following sort:

> *Airport.*[21] Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in

---

[21] Throughout this dissertation, I will give titles to all the cases I discuss.

Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes, I know—it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen 1999, 58)

Suppose that in both cases, Smith believes truly that the flight stops in Chicago. Intuitively, Smith's knowledge attribution "I know the flight stops in Chicago" seems true and Mary and John's attribution "Smith does not know that the flight stops in Chicago" seems true as well. This is puzzling, since Smith's belief that the flight stops in Chicago is held fixed. The difference between Smith and Mary and John is that the stakes in whether the flight stops in Chicago are high for Mary and John and low for Smith. Stakes are pragmatic (non-truth-relevant) factors, and it seems that pragmatic factors should have nothing to do with whether one knows or not.

Similar cases are abundant in current literatures (e.g., Keith DeRose's *Bank Case* (1992), Jeremy Fantl and Matthew McGrath's *Train Case* (2002), etc.). What is common in all these cases is that they all share (or are designed to share) the following two features: (a) the same knowledge-attributing sentence, say, 'S knows that p', seems true when uttered in context C1 but fails to be so when uttered in context C2, and (b) C1 and C2 are similar in all significant aspects except that someone's[22] stakes in whether p are high in C2 but low in C1. Following Fantl and McGrath (2009b), let us call any pair of cases satisfying (a) and (b) 'stakes-shifting cases'.

---

[22] There are variants here. In some cases, the subject's stakes are high. In other cases, the speaker's stakes are high. See Stanley (2005) for a more complete list of cases.

In this chapter, I consider the question whether stakes-shifting cases offer *prima facie* support for the contextualist theory (of knowledge attributions). I conclude that they do not. My objection to the contextualist theory is different from what is currently on the table on two scores. Firstly, one of the main objections to contextualism has focused on the linguistic basis for contextualism (cf. MacFarlane 2005). On this objection, the contextualist theory has failed to predict some ordinary linguistic behaviors regarding knowledge attributions such as disagreement and retraction. But few have challenged the contextualist explanation of stakes-shifting cases. Even if stakes-shifting cases are the focus, cases like *Airport* are still considered to be the refuge of the contextualist theory.

My main thesis here, however, is that the contextualist theory even fails to explain our linguistic disposition concerning knowledge attributions in stakes-shifting cases such as *Airport*. Secondly, many objections have targeted the contextualist theory. They do not necessarily apply to other forms of standards variantism. My objection, by contrast, can easily be extended to apply to standards variantism in general, since my objection targets all (prominent) accounts of the standards for knowledge.

2.2 The Contextualist Assumption

It has been widely accepted or argued that stakes-shifting cases provide *prima facie* support for the contextualist theory. For instance, Keith DeRose, a contextualist, has claimed that stakes-shifting cases are "the best grounds" (DeRose 2009, 47) for the contextualist theory (of knowledge attributions):

> [Stakes-shifting cases] provide us with the best possible type of evidence you could ask for that 'know(s)' is context-sensitive in at least roughly the way contextualists claim it is. (DeRose 2009, 67)

However, DeRose's claim is misleading, at best. Even if stakes-shifting cases do indicate the variations of the standards for knowledge, they still, to some extent, remain neutral to different forms of standards variantism; specifically, they do not necessarily support DeRose's favorite type of standards variantism (i.e., what MacFarlane (2009) calls 'indexical contextualism') over other types of standards variantism such as nonindexical contextualism or relativism.[23]

Still, its proponents regard stakes-shifting cases as directly supporting standards variantism. Here are remarks from John MacFarlane:

> If I was speaking literally both times [in stakes-shifting cases] and didn't make a mistake, then presumably the standards I must meet in order to count as "knowing" must have changed. I met the laxer standards that were in play [in the first case], but not the stricter ones that come into play [in the second case]. (MacFarlane 2005, 201; my italics)

It appears that standard variantists often assume that stakes-shifting cases directly support standards variantism, so much so that few have bothered to justify this assumption—most effort has been dedicated rather to show that stakes-shifting cases, when elaborated, support one type of standards variantism over another.

Interestingly, many opponents of contextualism (or at least opponents of the contextualist solution) are also sympathetic to the contextualist theory. For instance, Ernst Sosa, an opponent of the contextualist solution, agrees that contextualism is semantically or linguistically plausible:

> The main thesis of epistemic contextualism (EC) has considerable plausibility as a thesis in linguistics or in philosophy of language. (Sosa 2000, 3)

---

[23] For nonindexical contextualism, see Brogaard (2008). For relativism, see MacFarlane (2005) and Richard (2004).

Richard Feldman, an opponent of the contextualist solution, has also shown his sympathy towards the assumption that stakes-shifting cases seemingly support contextualism:

> It may be that knowledge attributions are context dependent. Perhaps the ordinary standards for knowledge are somehow flexible. Perhaps, setting aside the typical skeptical problems for a moment, it is sometimes true to say that a person knows a proposition and sometimes true to deny that the person knows that same proposition. Thus, for example, maybe the standards for knowledge shift in such a way that in casual conversation just prior to an election for which there are reliable polls indicating a clear winner, it is correct to say that we know what the outcome will be. Maybe in other contexts stricter standards apply and it is not correct to say that. That makes contextualism correct. (Feldman 1999, 111)

In brief, it seems to be a widespread assumption among theorists (standard variantists or not) that stakes-shifting cases offer *prima facie* support for standards variantism. Call this *the standards-variantist assumption*. Of course, standards *in*variantists—theorists reject the variability of the standards for knowledge (e.g., Stanley 2005; Hawthorne 2004)—want to reject the standards-variantist assumption. But even invariantists may well feel the persuasiveness or naturalness of the standards-variantist assumption.

In this chapter, I want to argue against this assumption. More precisely, I will focus on rejecting a specific version of this assumption, i.e., *the contextualist assumption*—the idea that stakes-shifting cases offer *prima facie* support for the contextualist theory. There are three reasons for focusing on contextualism. First, many people (myself included) have found the contextualist explanation very plausible when encountering stakes-shifting cases the first time. Indeed, stakes-shifting cases are originally designed to bring support to the contextualist theory (cf. DeRose 1992; Cohen 1999). Second, contextualists have been the only standard variantists who explicitly endorse the standards-variantist solution to the skeptical puzzle (in their version, the contextualist solution), and since the skeptical puzzle has been a main motivation of the

present study, it is not inappropriate if we pay attention to contextualism. Third, the following discussion also applies to other forms of standards variantism. My argument consists in the contention that no plausible account of the standards for knowledge is readily available in accounting for stakes-shifting cases. This should have an impact on all forms of standards variantism.

Let us first explicate the contextualist explanation of our intuitions regarding knowing attributions in stakes-shifting cases, focusing on *Airport*.

2.3 The Contextualist Rational Reconstruction of *Airport*

Recall that stakes-shifting cases are (usually) a pair of cases such that (a) the same knowledge-attributing sentence, say, 'S knows that p', seems true when uttered in context C1 but not in context C2, and that (b) C1 and C2 are similar in all significant aspects except that S's stakes (or the speaker's, or the evaluator's) in whether p are high in C2 but low in C1. Let us call C1 and C2 'LOW' and 'HIGH' respectively (indicating low-stakes and high-stakes circumstances respectively).

It seems that our intuitions concerning knowledge attributions in *Airport* are that:

2.1   Smith's utterance "I (Smith) know that the flight stops in Chicago" is true in LOW and Mary and John's utterance "Smith does not know that the flight stops in Chicago" is true in HIGH.

Contextualists take the contextualist theory to offer the best explanation for 2.1.

Let us consider contextualists' *rational reconstruction* of how 2.1 seems true in the first place. Rational reconstruction of this sort is intended to capture the structure of our mental procedure when encountering stakes-shifting cases; it need not describe the

41

conscious, explicit mental process. What is important is that it explains how our intuitions regarding stakes-shifting cases are generated in the first place.

First of all, contextualists take the standards for knowledge to be determined by the context of use (i.e., the speaker context). They formulate the truth condition of knowledge attribution as follows (see Section 1.7):

(K_C)   "S knows that p" is true in C2 if and only if

    (i)   S believes that p in C1,

    (i)   p is true in C1, and

    (iii)   S's standing on knowing dimensions with respect to p in C1 satisfies the standards for knowledge in C2.

The stipulation of *Airport* has it that Smith believes truly that the flight stops in Chicago in LOW and HIGH. Hence, 2.1 and (K_C) imply that at least one of the following theses must hold:

The standards-variantist thesis: The standards for knowledge vary with LOW and HIGH.

The standing-variantist thesis: Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* varies with LOW and HIGH.

In other words, either standards variation or standing variation is responsible for our intuitions concerning knowledge attributions in *Airport*.[24]

According to contextualists, the standards-variantist thesis is the better explanation since the standing-variantist thesis is incompatible with an intuitively plausible epistemic principle:

Intellectualism. S's standing on knowing dimensions with respect to p depends exclusively on S's standing on truth-relevant dimensions with respect to p.[25]

---

[24] Due to the consideration of simplicity, I ignore the thesis that both the standards for knowledge and Smith's standing on knowing dimensions vary with LOW and HIGH. See Section 1.7.

The argument is quite simple. The following seems to be the case:

2.2   Smith's standing on truth-relevant dimensions with respect to *The flight stops in Chicago* remains constant in LOW and HIGH.

After all, Smith's belief in *The flight stops in Chicago* is held fixed in LOW and HIGH. 2.2 and intellectualism rule out the standing-variantist thesis.

As a result, the standards-variantist thesis is the best explanation: In order to account for our intuitions concerning knowledge attributions in stakes-shifting cases (i.e., 2.1), we must assume that the standards for knowledge vary with LOW and HIGH. Hence, stakes-shifting cases offer *prima facie* support for the contextualist theory; the contextualist assumption holds.

2.4 A Specious Dilemma

The contextualist argument implies that the standards-variantist thesis is the only explanation for 2.1 assuming 2.2, (K_C), and intellectualism.

Opponents of the contextualist argument would have to reject at least one of the premises. Which one? 2.1, presumably, is unproblematic. At any rate, I will not dispute 2.1 here. (K_C) presupposes the contextualist theory. Non-contextualist standard variantists will reject (K_C). But the dispute, as far as I can see, is an in-house dispute among standard variantists, a dispute that is not immediately relevant here. The following discussion will not be seriously affected if we replace contextualism by other forms of standard variantism. For the sake of argument, I will stick to (K_C). 2.2 seems a natural description

---

[25] The term 'intellectualism' is from Stanley (2005). The doctrine is also called 'purism (about knowledge)' by Fantl and McGrath (2011).

of how people normally comprehend stakes-shifting cases (I disagree, though. See below). This leaves us intellectualism.

In other words, many have taken stakes-shifting cases as imposing the dilemma between the contextualist theory and the denial of intellectualism (cf. Hawthorne 2004; Stanley 2005). Contextualists, of course, think that contextualism is better than the denial of intellectualism. But there are others who want to take the alternative route as well.

According to *subject-sensitive invariantism* (SSI), we should deny intellectualism (cf. Fantl and McGrath 2002; 2009a; Stanley 2005; Hawthorne 2004)[26]. More precisely, SSI-ists take S's standing on knowing dimension with respect to p to consist in two factors: Firstly, S's standing on truth-relevant factors with respect to p and, secondly, S's stakes in whether p. SSI is a kind of standing variantism. On SSI, since S's stakes are different in LOW and HIGH, S's standing on knowing dimensions with respect to p is different as well.

SSI has been controversial and rightly so. The denial of intellectualism is very counterintuitive. One of the main motivations for SSI has been the dilemma between the contextualist theory and the denial of intellectualism. According to SSI-ists, intellectualism is the right price to pay for rejecting contextualism.

Siding with contextualists, I also do not share SSI-ists' opinion that the denial of intellectualism is better than the acceptance of contextualism. As I see it, contextualism, though implausible in some aspects, is more palatable than the denial of intellectualism.

---

[26] Stanley (2005) calls his own view 'Interest-Relative Invariantism'; Hawthorne (2004) calls his view 'Sensitive Moderate Invariantism'; Fantl and McGrath (2011) call their view 'pragmatic encroachment', a term from Kvanvig. The term 'Subject-Sensitive Invariantism' is from DeRose (2009).

Nevertheless, I do not think we are forced to choose between contextualism and the denial of intellectualism; I reject the alleged dilemma between them. I think it is plausible, and possible, to resist the contextualist theory (and the standards-variantist theory in general) while at the same time preserving intellectualism. The aforementioned contextualist rational reconstruction indeed fails, since 2.2 is not what people have in mind when comprehending stakes-shifting cases:

> 2.2   Smith's standing on truth-relevant dimensions with respect to *The flight stops in Chicago* remains constant in LOW and HIGH.

Here is a brief explanation of why 2.2 is incorrect. According to the contextualist reconstruction, our intuitions regarding the truth value of "Smith knows that the flight stops in Chicago" should vary with LOW and HIGH *regardless of whether we consider LOW and HIGH individually or together*. But this is not the case. Considered individually, we tend to think that the same knowledge attribution is true in LOW but false in HIGH. Considered together, by contrast, we tend to think the same knowledge attribution has the same truth value in both contexts—we tend to think that either Smith knows or Smith does not know in both cases, when considering LOW and HIGH at the same time.[27] In my opinion, this piece of data is best explained by giving up 2.2 (see Appendix A).

This is my first reason against the contextualist assumption—that stakes-shifting cases offer *prima facie* support for the contextualist theory—namely, the contextualist

---

[27] DeRose also admits that:
> If [we] were presented with [HIGH] and [LOW] together, then the pressure to give the same verdict about whether the subject in question knows in the two cases would be great—and greater than is the pressure to rule that one or the other of the claim made within the cases (that the subject 'knows' in LOW, and doesn't 'know' in HIGH) must be false. (DeRose 2009, 49)

Surprisingly, DeRose does not seem to be bothered by the fact that this observation is a counterexample of his contextualist theory.

assumption is based on an incorrect reconstruction of the way we comprehend stakes-shifting cases.

2.5 The Second Reason against the Contextualist Assumption

I suspect most theorists would not be completely persuaded by my denial of 2.2—after all, the aforementioned dilemma between contextualism and the denial of intellectualism is so widely adopted. Another drawback for this objection is that it does not automatically extend to other forms of standards variantism. Some forms of standards variantism can deal with the fact that our intuitions tend to be unified when considering LOW and HIGH together (cf. MacFarlane 2005; Brogaard 2008).

For these reasons, I will focus on my second reason for denying the contextualist assumption, a reason focusing on a core component of the contextualist theory (indeed a core component of all forms of standards variantism). Putting aside whether the aforementioned (contextualist) rational reconstruction is correct or not, I argue that we still should not endorse the contextualist explanation of 2.2, since there is no plausible account of the standards for knowledge that can explain our intuitions regarding knowledge attributions regarding stakes-shifting cases.

Let me elaborate. On the contextualist explanation, the difference in linguistic dispositions between Smith and Mary and John should be explained by the difference in the standards for knowledge in LOW and HIGH. This gives us a basis to test the contextualist explanation, namely, if our intuitions concerning Smith's and Mary and John's knowledge-attribution dispositions manifest the systematic differences predicted

by the variation of the standards for knowledge, contextualism offers (or at least is in a position to offer) a proper explanation for stakes-shifting cases; otherwise, not.

However, a problem immediately arises once we try to evaluate whether our intuitions concerning Mary and John's and Smith's linguistic dispositions reflect the variation of the standards for knowledge. That is, what is the nature of the standards for knowledge? Unfortunately, contextualists lack agreement on this issue; radically different accounts have been proposed. Without an account of the standards for knowledge, however, the contextualist explanation of stakes-shifting cases is simply non-evaluable.

For the sake of argument, I will consider all (prominent) contextualist accounts of the standards for knowledge, evaluate them individually, and see whether or not they can account for our intuitions concerning knowledge attributions in stakes-shifting cases. This is the plan of the remainder of this chapter.

Jonathan Schaffer (2005) has offered a useful framework. Schaffer points out that when contextualists talk about the shift of the standards for knowledge, there are three possible accounts: In his terms, the shift of 'threshold', 'standard', and 'alternative'. Following his framework, I will distinguish two main accounts of the standards for knowledge: I call them the *general* and *particular accounts*. According to the general account, the shift of the standards for knowledge will affect any proposition with a certain general property. That is, the effect of the standards for knowledge over the logical space is, in Jonathan Schaffer's term (2005), "globally encompassing". The particular account contends that the shift of the standards for knowledge affects only a specific set of propositions that does not form a globally encompassing logical space. The effect of the standards on logical space is, in Schaffer's term, "pointlike".

47

I will discuss possible variants of these two accounts, arguing that none of them offers satisfactory explanation for our intuitions concerning knowledge attributions.

2.6 The Linear Account

There are two types of the general account: I will call them the *linear* and *spherical accounts*. The linear account takes the standards for knowledge as specifying the threshold for S's standing on knowing dimensions with respect to p such that S can be truthfully described as "knows that p" only if S's standing on knowing dimensions with respect to p meets the threshold. On the linear account, the logical space extends in the linear, scale-like, manner.

There are two prominent ways to develop the linear account. According to the evidentialist account, the standards for knowledge set the threshold for the strength of evidence such that the strength of S's evidence for p must reach a certain threshold in order for S to be counted as "knows that p". On the reliabilist account, by contrast, the standards for knowledge set the threshold for the degree of reliability of the belief's forming processes such that the degree of reliability of the process that forms S's belief that p must meet the threshold in order for S to be counted as "knows that p". In what follows, I will focus on the reliabilist account, but my objection can be easily extended to the linear account in general.

Suppose that in LOW, the degree of reliability of the process forming Smith's belief that the flight stops in Chicago is .85, or, in short, *the degree of reliability of Smith's belief* is .85. Moreover, suppose that the threshold of the degree of reliability in

HIGH, according to the standards for knowledge, is .90 while the threshold in LOW is .80. Accordingly, the reliabilist account's typical explanation of *Airport* is as follows:

> Smith's utterance "I know that the flight stops in Chicago" is true in LOW since the degree of reliability of Smith's belief that the flight stops in Chicago, i.e., .85, meets the degree of reliability specified by the standards for knowledge in play, i.e., .80. By contrast, Mary and John's utterance "Smith does not know that the flight stops in Chicago" is also true in HIGH since the degree of reliability of Smith's belief that the flight stops in Chicago is held fixed, i.e., .85, and the latter does not meet the standards in play in HIGH, i.e., .90.

To test the reliabilist account's explanation of *Airport*, let us consider some of Smith's beliefs whose degree of reliability is also .85 in LOW. Consider the following follow-up of Airport:

> *Airport′.* Everything is like *Airport* except that Mary and John and Smith are friends and Mary and John knows that Smith is a Lakers fan. Suppose that Mary knows that the Lakers won yesterday's game, and that Smith believes truly that the Lakers won yesterday's game since she saw him reading today *The New York Times*. Suppose that, in LOW, Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* is epistemically equivalent to his standing on knowing dimensions with respect to *The Lakers won yesterday's game*. Now, on their way to find the airline agent to check for the flight schedule, Mary and John are chatting. John asks Mary whether Smith knows that the Lakers have won yesterday's game.

Now, consider two possible situations:

> *Airport(a).*   Mary says to John, "Yes, Smith knows the Lakers won yesterday's game."

> *Airport(b).*   Mary says to John, "No, he doesn't. Smith believes truly that the Lakers won yesterday's game, but he does not know."

According to our intuitions, as I see it, *Airport(a)*, rather than *Airport(b)*, is the natural reply. But the reliabilist account implies the opposite. To elaborate, consider the following rational reconstruction of the contextualist explanation of *Airport′*.

> *Airport′* is stipulated as follows:

2.3 In LOW, Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* is epistemically equivalent to his standing on knowing dimensions with respect to *The Lakers won yesterday's game*.

On the reliabilist account, one's standing on knowing dimensions with respect to p should be characterized by the degree of reliability of one's belief that p. Hence, from 2.3, the reliabilist account implies that:

2.4 In LOW, the degree of reliability of Smith's belief that the flight stops in Chicago and the degree of reliability of Smith's belief that the Lakers won yesterday's game are both .85.

Moreover, contextualists assume that:

2.2 Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* remains constant across LOW and HIGH.

2.2 and 2.4 imply that:

2.5 In HIGH, the degree of reliability of Smith's belief that the flight stops in Chicago and the degree of reliability of Smith's belief that the Lakers won yesterday's game are both .85.

Arguably, 2.5 further entails that:

2.6 In HIGH, Smith knows that the flight stops in Chicago if and only if Smith knows that the Lakers won yesterday's game.

But, according to the specification of *Airport'*,

2.7 In HIGH, Smith does not know that the flight stops in Chicago.

2.6 and 2.7 thus imply that:

2.8 In HIGH, Smith does not know that the Lakers won yesterday's game.

The aforementioned argument is valid, and proponents of the reliabilist account, I take it, must accept all the premises. Hence, 2.8 must be true for them. However, 2.8, in turn, indicates that *Airport(b)* rather than *Airport(a)* is the natural follow-up of *Airport'*, *contra* our intuitions.

In sum, if the reliabilist account was correct, we should expect *Airport(b)* as opposed to *Airport(a)* to be the natural follow-up of *Airport'*. But this fails to square well with our intuitions. Since one can construct infinitely many counterexamples along the aforementioned line, the argument thus indicates that the reliabilist account fails thoroughly in explaining cases like *Airport'*. In general, the reliabilist account fails to explain the variations of the attributors' dispositions of knowledge attributions in stakes-shifting cases.

Moreover, since similar arguments can be constructed against other versions of the linear account such as the evidentialist account, the prospect of the linear account is thus dim. Without further ado, let us consider the other type of the general account, i.e., the spherical account.

2.7 The Spherical Account

On the spherical account, the standards for knowledge specify a certain sphere s in the logical space such that the sphere determined by the subject's standing on knowing dimensions with regard to p must include or at least coincide with s. Like the linear account, the spherical account also claims that the standards for knowledge in a context specify a certain globally encompassing, unified area in the logical space. But unlike the linear account, the spherical account does not take the logical space to extend in the linear, scale-like, manner. Rather, the logical space is regarded as extending in a spherical manner (cf. Schaffer 2005). One way to understand the notion of sphere in question is to consider a sphere as a set of possible worlds, whereas the distance of a world is to be

defined by how far away the world is from the center of the sphere (i.e. the actual world).[28]

There are two prominent ways one may develop the spherical account. On the *sensitivity account*, S counts as "knows that p" only if S's belief that p is *sensitive* with regard to the set of possible worlds specified by the standards for knowledge. S's belief that p is sensitive if and only if in the nearest possible world w in which p is false, S will not believe that p in w (cf. Nozick 1981). The shift of the standard is thus the shift of the *range* of possible worlds that S's belief that p remains sensitive (cf. DeRose 1995). On the *safety account*, S counts as "knows that p" only if S's belief that p is *safe* with regard to the set of possible worlds specified by the standards for knowledge. S's belief that p is safe if and only if in the nearest possible world w in which S believes that p, p is true in w (cf. Sosa 1999). The shift of the standard is thus the shift of the range of possible worlds that S's belief that p remains safe.[29] In what follows, I will focus on the sensitivity account.[30]

First, consider how the sensitivity account will explain the attributors' linguistic dispositions in *Airport*:

> The standards for knowledge in LOW specify the set of possible worlds w such that *the flight stops in Chicago* is true in each w, whereas the standard in HIGH specifies the set of possible world v such that *the flight stops in Chicago* is false in some v (suppose that the flight itinerary contains a misprint in such v). Moreover, in *Airport*, Smith will believe falsely that the flight stops in Chicago if the flight itinerary contains a misprint. As a result, Smith's belief that the flight stops in Chicago is sensitive in LOW, since in the set of possible worlds specified by the standards for knowledge in LOW, there is no possible world w in which *the flight*

---

[28] Normally, the distance between a possible and the actual worlds is the function of the similarity between these two worlds. The more similar a possible is to the actual one, the closer the former is to the latter.
[29] I should note that no one, to my knowledge, has officially defended the safety account.

[30] A similar argument can easily be constructed against the safety account.

stops in Chicago is false and Smith believes it. Smith's belief is not sensitive in HIGH, since in the closest possible world v in which *the flight stops in Chicago* is false, Smith still believes that the flight stops in Chicago in v.

The sensitivity account suffers from the similar problem as the reliability account does, however.

To illustrate, let us reconstruct the contextualist explanation as follows. Notice that *Airport′* is stipulated as that:

> 2.3   In LOW, Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* is epistemically equivalent to his standing on knowing dimensions with respect to *The Lakers won yesterday's game*.

Now, the sensitivity account takes it that one's standing on knowing dimensions with respect to p is to be characterized by the sensitivity of one's belief that p. As a result, the following is a direct consequence of 2.3 for the sensitivity account:

> 2.9   The distance of the closest possible world in which the itinerary contain a misprint to the actual world is the same as the distance of the closest possible world in which the *New York Times* contain a misprint to the actual world.

Now, recall that contextualists assume that:

> 2.2   Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* remains constant across LOW and HIGH.

2.2 and 2.9 further imply that:

> 2.10   The distance of the closest possible world in which the itinerary contain a misprint to the actual world is the same as the distance of the closest possible world in which the *New York Times* contain a misprint to the actual world.

Now, on the sensitivity account, the following holds:

> 2.11   In HIGH, Smith's belief that the flight stops in Chicago is insensitive.

2.10 and 2.11 then imply that:

2.12   In HIGH, Smith's belief that the Lakers won yesterday's game is insensitive.

Now, the sensitivity account assumes that S knows that p only if S's belief that p is sensitive, 2.12 thus implies that:

2.8   In HIGH, Smith does not know that the Lakers won yesterday's game.

The aforementioned argument is valid and proponents of the sensitivity account, I take it, must accept each premise. However, 2.8, in turn, indicates that *Airport(b)* rather than *Airport(a)* is the natural follow-up of *Airport'*, *contra* our intuitions.

Let us take stock. We have seen that the general account fails to properly predict the variations of the attributors' linguistic disposition in stakes-shifting cases. Both the linear and spherical accounts have suffered from a similar problem, namely, they mistakenly take *Airport(b)* as opposed to *Airport(a)* as the natural follow-up of *Airport'*. The moral, I take it, is that in stakes-shifting cases, our intuitively think that the attributors' linguistic dispositions do not undergo a global variation, namely, when the attributor counts the subject as "know that p" in LOW, but "does not know that p" in HIGH, the attributor does not, at least not always, also take the subject as "does not know that q" in HIGH for all q with respect to which the subject's standing on knowing dimensions in LOW is epistemically equivalent to the subject's standing on knowing dimensions with respect to p in HIGH. The basic flaw of the general account, in brief, consists in its implication that the attributors exhibit a certain global variation in their linguistic dispositions in stakes-shifting cases.

Of cause, proponents of the general account may avoid this problem, by giving 2.2 (i.e., that Smith's standing on knowing dimensions remains constant across LOW and HIGH). Contextualism with the denial of 2.2 thus contends that both the standards of

knowledge and S's standing on knowing dimensions vary with LOW and HIGH. Such modification of contextualism makes the theory less attractive than standing variantism, which upholds only the variation of S's standing on knowing dimensions. The reason is that if standing variantism (at least some version of it) can account for our intuitions concerning knowledge attributions in stakes-shifting cases, then the variation of the standards for knowledge is theoretically redundant—it plays no crucial role in explaining stakes-shifting cases (exactly the same pattern of argument will be used to argue against the contextualist solution (to the skeptical puzzle), see Section 3.7).

2.8 The Particular View

Realizing the aforementioned problem of the general account, some contextualists thus prefer the particular account according to which the variations of the standards for knowledge will not have a global impact on the attributor's linguistic dispositions in stakes-shifting cases (cf. Schaffer 2005). Instead, the standards for knowledge will only have a "pointlike" effect on the logical space.

The most prominent particular account is the *relevant alternative account (of the standards for knowledge)* (hereafter 'the RA account') which characterizes the variations of the standards for knowledge in terms of the variations of the set of relevant alternatives. The idea is that the subject S counts as "knows that p" only if S's standing on knowing dimensions with regard to p enables S to rule out all propositions in the set of relevant alternatives, specified by the standards for knowledge in context.[31, 32]

---

[31] The RA account of *knowledge* has been endorsed by many philosophers. See, e.g. Dretske (1970). But proponents of the RA account of knowledge may not endorse the RA account of *the standards for knowledge*. For the latter, see, e.g., Lewis (1996), Blome-Tillmann (2009) (also see Section 2.8-2.13).

The RA account's typical explanation of *Airport* is as follows:

In *Airport*, Smith's evidence is unable to rule out the proposition that *The itinerary contains a misprint*. In LOW, however, the standards for knowledge do not specify the proposition to be a relevant alternative. As a result, Smith counts as "knows that the flight stops in Chicago" even though he is unable to rule out *The itinerary contains a misprint*. By contrast, the standards for knowledge in HIGH do specify *The itinerary contains a misprint* to be a relevant alternative. Accordingly, Mary and John's utterance "Smith does not know that the flight stops in Chicago" is true in HIGH.

Now, some might offer a similar RA account's explanation for *Airport'*:

In *Airport'*, Smith's evidence is unable to rule out the proposition that *The New York Times contains a misprint*. In HIGH, however, the standards for knowledge do not specify the proposition to be a relevant alternative. As a result, Smith counts as "knows that the Lakers won yesterday's game" in HIGH.

A closer scrutiny reveals that the core of the RA account's explanation of *Airport* and *Airport'* consists of two claims:

RA1: The alternative *The flight itinerary contains a misprint* is relevant in HIGH.

RA2: The alternative *The New York Times contains a misprint* is irrelevant in HIGH.

Obviously it is not enough for proponents of the RA account to merely point out RA1 and RA2. They need to motivate them; they need to explain *why* the alternative *The flight itinerary contains a misprint* is relevant in HIGH while the alternative *The New York Times contains a misprint* is irrelevant in HIGH. Otherwise, the RA account is just an ad hoc story tailor-made to account for whatever our intuitions are concerning stakes-shifting cases. What is needed here are principles that determine whether an alternative is relevant or not. Call them 'principles of relevance'. The RA account needs to show that RA1 and RA2 are the case based on those principles.

---

[32] For simplicity's sake, I will not analyze the notion of *ruling out an alternative* here.

Not all principles of relevance will do the job, though. Particularly, principles that characterize the notion of relevance in terms of the notion of knowledge will not do. I have in mind something like Cohen's definition of relevance in one of his early papers (Cohen 1988, 101):

(Def$_R$) An alternative q to p is relevant (for S) =df. S's epistemic position with respect to q precludes S from knowing p.

Putting some details aside, the central approach of (Def$_R$) is to explicate relevance in terms of knowledge. Roughly, an alternative q to p, according to (Def$_R$), is always irrelevant for S if S knows that p and S is unable to rule out p. And if q is relevant, then either S does not know p or S is able to rule out q. However, principles of relevance of this sort cannot be useful in explaining RA1 and RA2.[33] To explicate the notion of relevance in terms of the notion of knowledge is tantamount to explaining why RA1 and RA2 are the case by appealing to our intuitions that, in HIGH, Smith knows the Lakers won yesterday's game but does not know the flight stops in Chicago. To do so is to beg the question—using what should be the explanandum as the explanans. We need principles of relevance whose formulations do not (at least not explicitly) appeal to the notion of knowledge.

The specification of the principles of relevance has been a serious problem for the RA account (cf. Vogel 1999). Proponents of the RA account have yet had agreement on this issue, and, worse, in many cases, only vague descriptions are provided.

David Lewis and Michael Blome-Tillmann are two exceptions. Lewis (1995), I believe, is the first philosopher to offer a complex, systematic account of the principles of

---

[33] Notice that Cohen's original goal in proposing (Def$_R$) is not to explain RA1 and RA2. His target is skepticism.

relevance. Blome-Tillmann (2009), a follower of Lewis, endorses all Lewis's principles except one. The following discussion is mainly based on Lewis's and Blome-Tillmann's works.

Let us note two general requirements for the satisfactory explanation of RA1 and RA2. Firstly, the alleged principles of relevance must be readily available to predict RA1 and RA2. Secondly, the cited principles must be individually plausible. The first requirement should be uncontroversial. The second requirement is also necessary since implausible principles would have no theoretical merit in the first place. Against Lewis and Blome-Tillmann, I want to argue that none of Lewis' or Blome-Tillmann's principles has satisfied these two requirements. I will start with RA2 (Section 2.9), and then move on to RA1 (Section 2.10-2.13).

2.9 The Reliability Principle

How do we know that an alternative is irrelevant in a context? The principles of relevance that determine the irrelevance of an alternative are *permissive*—they allow some alternatives as being properly ignored. Lewis' has listed four permissive principles (what he calls "the Rule of Reliability", "(two) Rules of Method", and "the Rule of Conservatisism").

With respect to RA2, the Rule of Reliability is particularly relevant. The rule is a presumptive rule, as Lewis writes:

> Within limits, we are entitled to take [perception, memory, and testimony] for granted. We may properly presuppose that they work without a glitch in the case under consideration. Defeasibly—very defeasibly!—a possibility in which they fail may properly be ignored. (Lewis 1996, 558)

The rule can thus be reformulated as follows:

The reliability principle. If p is incompatible with the assumption that perception, memory, and testimony do not fail, then p is not relevant.

On the reliability principle, the alternative *The New York times contains a misprint* is irrelevant in HIGH since it is, arguably, incompatible with the assumption that Smith's testimony (*The New York Times*) does not fail in HIGH. Is this enough to explain RA2?

Not really. Note that the reliability principle is defeasible; its verdict can be defeated if there are some other Lewisian principles of relevance showing that *The New York Times contains a misprint* is relevant. Hence, a satisfactory explanation requires that the reliability principle is not overthrown by other Lewisian principles of relevance. Unfortunately, a close scrutiny of Lewis's other principles of relevance suggests the opposite.

"The possibility that actually obtains," according to Lewis, "is never properly ignored" (1996, 554). Hence, the following is also a principle of relevance:

The actuality principle. If p is true in the actual world[34], then p is always relevant. Moreover, Lewis also thinks that an alternative that resemble a relevant alternative is also relevant:

> Suppose one possibility saliently resembles another. Then if one of them may not be properly ignored, neither may the other. (Or rather, we should say that if one of them may not properly be ignored in virtue of rules other than this rule, then neither may the other. Else nothing could be properly ignored; because little steps of resemblance can take us from anywhere to anywhere.) (Lewis 1996, 556)

This gives us the following principle:

The resemblance principle. If p saliently resembles q, and p is relevant (in virtue of principles other than this one), then q is relevant.

---

[34] One may wonder whose actuality Lewis means there: the speaker's or the subject's? Lewis argues that it is the subject's actuality that always counts (1996, 554).

The combination of the actuality and resemblance principles, however, implies that *The New York Times contains a misprint* is relevant. The idea is that the actual world in which *The New York Times* does not contain a misprint saliently resembles the counterfactual world in which *The New York Times* contains a misprint—after all, the only difference between these two worlds is that *The New York Times* contains a misprint in one but not in the other. Now, by the actuality principle, *The New York Times does not contain a misprint* is relevant. Moreover, since *The New York Times does not contain a misprint* is relevant, the alternative *The New York Times contains a misprint* is also relevant here, by the resemblance principle. In brief, the actuality and resemblance principles overthrow the verdict of the reliability principle.

One might reply to this objection by claiming that the resemblance principle was not applicable here since the actual world in which *The New York Times* does not contain a misprint did not *saliently* resemble the counterfactual world in which *The New York Times* contains a misprint. This reply does not work, since we can easily modify *Airport′* such that the difference between these two worlds becomes salient. For instance, we can stipulate that in *Airport′* Mary and John just discuss that *The New York Times* has a .0001 rate of misprinting the result of games in general. Or we can stipulate that when John asks Mary whether Smith knows the Lakers won yesterday's game, Mary says to John, "Yes, Smith knows the Lakers won yesterday's game. *The New York Times* rarely misprints the result of basketball games." In either case, the possibility of *The New York Times* contains a misprint has been made salient, and yet, I take it, we still intuitively think that Mary's knowledge attribution "Smith knows that the Lakers won yesterday's game" is true.

I think the moral generalizes. As Lewis formulates them, all permissive principles—principles that determine what can be properly ignored—are defeasible. In principle, all of them can be overthrown by the actuality and resemblance principles. This shows that Lewis' original proposal of principles of relevance indeed leaves little room for RA2.[35]

2.10 The Actuality and Resemblance Principles

Let us consider RA1:

RA1: The alternative *The flight itinerary contains a misprint* is relevant in HIGH.

Notice that the actuality and resemblance principles discussed above are readily available in predicting RA1.

On the actuality principle, the proposition *The flight itinerary does not contain a misprint* is relevant since it is true in the actual world. Since the world in which the flight itinerary does not contain a misprint saliently resembles the world in which the flight

---

[35] Some might have noticed that Blome-Tillmann, in one of his (relatively) early papers, has suggested a permissive principle which seems to be able to account for RA2. He calls it 'the Rule of Conversational presupposition':

      (RCP)   If w is incompatible with the speaker's conversational presuppositions in C, then w can be properly ignored in C (unless there is a rule other than RCP states that w cannot be properly ignored in C). (Blome-Tillmann 2007, 393)

(RCP) predicts that *The New York Times contains a misprint* is irrelevant if Mary and John's conversational presupposition in *Airport'* has it that *The New York Times* do not contain a misprint (for Blome-Tillmann's characterization of conversational presupposition, see Section 2.13).

      (RCP) has suffered from two problems. First, being defeasible, its verdict is overthrown by the verdict of the actuality and resemblance principles, just like the reliability principle. Second, (RCP) has suffered from the following counterexample (which is later discussed by Blome-Tillmann himself):

      Suppose that you tell Lazy Johnny that you will either be at home or at work. Assuming that you are at home, can Lazy Johnny in his context of deliberation, come to 'know' that you are at home merely by presupposing that you are not at work? Can he properly ignore the possibility that you are at work simply in virtue of presupposing that you are not at work? (Blome-Tillmann 2009a, 257)

To these questions, (RCP) says "Yes", and this is implausible (Blome-Tillmann appears to abandon (RCP) in his later works).

itinerary contains a misprint, the alternative *The flight itinerary contains a misprint* is also relevant, by the resemblance principle.

Is this a satisfactory answer? Not if we show that the actuality and resemblance principles are plausible principles of relevance in the first place. But these two principles are indeed problematic. Lewis himself has noticed a serious problem of these two principles:

> We must apply the Rule of Resemblance with care. Actuality is a possibility uneliminated by the subject's evidence. Any other possibility that is likewise uneliminated by the subject's evidence thereby resembles W actuality in one salient respect: namely, in respect of the subject's evidence. That will be so even if W is in other respects very dissimilar to actuality—even if, for instance, it is a possibility in which the subject is radically deceived by a demon. (Lewis 1996, 556)

If these two principles were proper principles of relevance, skepticism easily followed. Lewis' own reply to this problem is very unsatisfactory, to say the least:

> Plainly, we dare not apply the Rules of Actuality and Resemblance to conclude that any such W is a relevant alternative—that would be capitulation to scepticism. The Rule of Resemblance was never meant to apply to *this* resemblance! We seem to have an *ad hoc* exception to the Rule, though one that makes good sense in view of the function of attributions of knowledge. What would be better, though, would be to find a way to reformulate the Rule so as to get the needed exception without *ad hoc*ery. I do not know how to do this. (Lewis 1996, 556–7)

One should feel deeply unsatisfied with Lewis' reply, and is rightly so, since Lewis' reply is essentially, borrowing Robert Fogelin's comments, "a defense through a candid admission of failure" (Fogelin 2000, 52).

Hence, on pain of skepticism, we should give up either the actuality or the resemblance principles (I leave it open which one we should give up). Perhaps, proponents the RA account also have motivation for doing so—as mentioned last section, these two principles also debar the RA account from accounting for RA2.

2.11 The Attention Principle

In addition to the principles discussed above, it seems that what Lewis calls 'the Rule of

Attention' can also predict RA1. Here are Lewis's remarks on this rule:

> [A] possibility not ignored is *ipso facto* not properly ignored. What is and what is not being ignored is a feature of the particular conversational context. No matter how far-fetched a certain possibility may by, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative. (Lewis 1996, 559)

Let us reformulate the principle as such:

> *The attention principle.* If the subjects in C are attending to p, then p is a relevant alternative in C.

On the face of it, the attention principle seems to explain RA2 pretty well. The idea is

that both Mary and John are attending to possibility of the flight itinerary containing a

misprint in HIGH, and hence, by the attention principle, the alternative *The flight*

itinerary contains a misprint is a relevant alternative.

The attention principle, however, is implausible in its own right. One problem is

that it is too strong—mere attending an alternative does not necessarily make it relevant.

Consider Blome-Tillmann's example:

> Imagine you saw your teenage son sneaking away through the window of his room late at night. When you confront him the next morning he replies somewhat desperately: 'How do you know I left the house? I mean, for all you know you might have dreamt it. It was late at night, wasn't it?' (Blome-Tillmann 2009a, 246–7)

If the attention principle was right, then the alternative *The parent is merely dreaming of*

the teenage sneaking out last night would be relevant in the imagine context, and given

the assumption that the parent cannot rule out this alternative, the parent has to admit that

he or she does *not* know that the son was sneaking out last night. But this conclusion, presumably, is absurd.

To reply, one may modify the attention principle such that mere attendance is not enough to render p relevant. What is needed, one might suggest, is that the subjects *seriously* attending to p (cf. Hawthorne 2004, 64). In other words:

> *The attention principle′*. If the subjects in C seriously attend to p, then p is a relevant alternative in C.

Like the attention principle, the attention principle′ is readily available to explain RA1. Moreover, the attention principle' can avoid Blome-Tillmann's aforementioned case: the parent in that context is (presumably) not seriously attending to the possibility of him being dreaming. Therefore, the alternative that *The parent is merely dreaming of the teenage sneaking out last night* is not relevant in that context.

Still, the attention principle′ is implausible. First, it seems that even serious attention should not be a sufficient condition for the relevance of alternatives—a paranoiac patient may take many remote alternatives seriously, but this does not automatically make those alternatives relevant. Second, it is unclear, according to this proposal, how the standards for knowledge are to be determined by serious attention. Recall that in explaining Blome-Tillmann's case, the attention principle′ has it that if one of the subjects (i.e., the parent) does not seriously attend to p, then the antecedent of the principle is false. How should we interpret this application of the attention principle″? Should we say that as long as there is one subject in a context who does not seriously attend to p, then the antecedent of the attention principle′ is false? What if we are in a situation where a group of normal people seriously attend to a reasonable alternative except *silly* Smith? Should we say that the subjects that are not seriously attending to p

must not be less than a half of all the subjects? What if we are in a situation where one third of the subjects seriously attend to a reasonable alternative while the rest of the subjects stubbornly ignore the reasonable alternative? The condition of application of the attention principle′ is simply not clear.

In brief, the attention principle and the principle′ are both problematic. They fail to satisfy the second requirement for a satisfactory explanation of RA2. Hence, proponents of the RA account are better off appealing to other principles of relevance.

2.12　The High-Stakes Principle

When considering stake-shifting cases, one natural thought is to think that whether something is a relevant alternative to p in C is to be determined by whether the stakes in whether p are high or not. Consider Lewis's Rule of High Stakes:

> [W]hen error would be especially disastrous, few possibilities are properly ignored. (Lewis 1996, 556, footnote 12)

Accordingly, we might construct the following principle:

> *The high-stakes principle*. If the stakes of p are high in C, then few alternative q to p are irrelevant in C.

The high-stakes principle can explain RA1: Since the stakes of *The flight stops in Chicago* are very high in HIGH, *The flight itinerary contains a misprint* is thus relevant in HIGH.

Is the high-stakes principle what the RA account needs? I do not think so. One obvious problem of the high-stakes principle is that it is *vague*. It requires few alternatives to be irrelevant in high-stakes situations. But how few are few here? Another, more serious, problem is that while plausible in explaining *Airport*, the high-stakes

principle does not seem to do well in other ordinary cases. In particular, the principle also has difficulty explaining Blome-Tillmann's aforementioned case.

Suppose that in that context, the stakes in whether the teenage has sneaked out last night are high. But does that imply that the alternative *The parent is merely dreaming of the teenage sneaking out last night* is a relevant alternative? An affirmative answer is surely counterintuitive.

To reply to the aforementioned objections, one might modify the high-stakes principle as follows:

> *The high-stakes principle′.* If the stakes of p are high in C, and the subjects have some reason to believe in an alternative q to p, then q is a relevant alternative in C.

The high-stakes principle′ is not vague and it is readily available in taking care Blome-Tillmann's case. The idea is that since both the parent and the teenage have no reason to believe that the parent is merely dreaming that the teenage sneaked out last night, we thus lack reason to take the alternative as relevant.

But the high-stakes principle′ is still implausible; it incorporates all alternatives q to p as relevant so long as the subjects have some reason to believe in q. In some cases, this literally renders knowledge impossible. Consider *Airport* again.

Suppose that Mary and John go to check the flight schedule with the airline agent and learn that the flight stops in Chicago. Now, given that the stakes are high, it seems that by the high-stakes principle, the hypothesis that the airline agent has misread the flight information on the screen is also a relevant alternative in HIGH, since they have at least some reason to believe that human beings make mistake sometimes. Since Mary and

John are, let us assume, unable to rule out this alternative, they do not count as "knows that the flight stops in Chicago" even *after* consulting the airline agent. Counterintuitive.

Notice that one could not reply to this counterexample by pointing out that even Mary and John did not know that the flight stops in Chicago after consulting the airline agent, they did know if they asked the airline agent to double-check (or triple-check) for the flight information, since they had no reason to believe that the agent would make such a trivial mistake after double-checking (or even triple-checking). This reply does not work since even if we granted that Mary and John had no reason to believe that the airline agent would make such a trivial mistake after double-checking (or triple-checking), still, they would have some reason to believe that the airline agent has been impatient to them and does not double-check (or triple-check) properly. After all, not everybody is nice and patient.

In general, Mary and John would always have some reason to believe in some alternatives to *The flight stops in Chicago* in HIGH, which they are incapable of ruling out. The high-stakes principle thus renders them relevant in HIGH, and this implies that Mary and John are in principle incapable of knowing *The flight stops in Chicago* in HIGH. The result, presumably, is unacceptable.

2.13 The Presupposition Principle

In a recent paper, Blome-Tillmann (2009) offers a form of the RA account that seems to be able to account for ordinary knowledge attributions (stakes-shifting or not). I will consider his account in this section.

Blome-Tillmann endorses all Lewis' principles except the attention principle. He suggests that we should replace this principle by what we may call 'the presupposition principle'. This principle appeals to the notion of pragmatic presupposition to determine whether or not an alternative is relevant in a context:

> *The presupposition principle*. If an alternative q to p is compatible with the speaker's pragmatic presupposition in C, then q is relevant in C.

Blome-Tillmann further argues that S pragmatically presupposes that p in C if and only if S "is disposed to behave, in her use of language, as if she believed p to be common ground in C" (2009a, 253).

The flight itinerary contains a misprint is compatible with Mary and John's presupposition in HIGH—for all they know, that might be the case. As a result, the presupposition principle implies that *The flight contains a misprint* is relevant in HIGH.

The presupposition principle, unfortunately, is still implausible in its own right; it implies that one can be correctly counted as "does not know" while one in fact knows. Consider the following case:

> *Trial*. Adam is testifying in the court. Adam remembers, correctly, that he saw Jenny murdered Allen the other night. However, for some reason, Adam has decided not to testify against Jenny. Suppose that Adam cannot rule out the alternative *An alien hypnotized him to hallucinate that Jenny murdered Allen* to *Jenny murdered Allen*. Moreover, he is disposed to behave, in his use of language, as if he believes *An alien may hypnotize him to hallucinate that Jenney murdered Allen* to be common ground. Now, when asked whether he knows that Jenny murdered Allen the other night, Adam answers, "I don't' know".

Adam's knowledge attribution "I don't know that Jenny murdered Allen the other night" seems obviously false. But the presupposition principle says the opposite. That is, since, by stipulation, Adam pragmatically presupposes *An alien may hypnotize him to hallucinate that Jenny murdered Allen* and he is incapable of ruling out the alternative *An*

alien hypnotized him to hallucinate that Jenny murdered Allen, it follows that Adam can be truthfully described as "not knows that Jenny murdered Allen".

Some might argue that the problem arose from the fact that *Trial* is stipulated in such a way that Adam is somehow *insincere* in his presupposition, he presupposes something that he does not seem to believe, i.e., that an alien may hypnotize him in hallucinating that Jenny murdered Allen. Some might think that the problem could be avoided if we modified the presupposition principle so that it requires what is presupposed be believed as well. But this modification does not really save the principle from the objection. We can modify *Trial* such that Adam indeed believes that it possible (although unlikely) for an alien to hypnotize him the other night. Adam can just be a fallibilist.

The problem of the presupposition principle, in my opinion, is that it gives too much freedom for the speaker to determine the content of 'knowledge'. Presumably, the speaker has quite a freedom to decide whether he wants to pragmatically presuppose something or not. The problem is that the speaker does not have the same amount of freedom in determining the content of 'knowledge'. Blome-Tillmann seems to think that the move of identifying pragmatic presupposition with the content of 'knowledge' is a move with "advantages" (2009, 248). He thinks that this move avoids the aforementioned counterexample to the attention principle:

> …since the speaker can, to a certain extent, voluntarily decide what they take seriously and which propositions they presuppose, they have, to a certain extent, voluntary control over the content of 'know' in their contexts. For instance, as long as you make clear to your son that the possibility that you dreamt seeing him sneaking out of his window is not a 'live option' in your conversation, you remain in a context in which you satisfy 'know', even though your son has drawn attention to the possibility that you might have dreamt the relevant episode. (Blome-Tillmann 2009a, 248)

In my opinion, taking pragmatic presupposition as an essential factor in determining the content of 'knowledge' is a deeply mistaken view. The speaker should not enjoy the freedom of determining the content of 'knowledge' inasmuch the same way as she is free to determine what proposition to be pragmatically presupposed. *Trial* shows that the speaker has no right to voluntarily 'stupefy' herself, to deprive herself from knowledge on purpose. Whether someone knows something or not should not consists in this person's pragmatic presupposition.

2.14 Taking Stock

To sum up, in this chapter, I argue against the standards-variantist assumption, which states that standards variantism is *prima facie* supported by stakes-shifting cases. I focus on a specific version of this assumption, i.e., the contextualist assumption. But my objection against the contextualist assumption can easily be extended to the standards-variantist assumption since my target is the core idea of all forms of standards variantism—the account of the standards for knowledge.

Contextualists have argued that stakes-shifting cases present us a dilemma between contextualism and the denial of intellectualism. This dilemma, if true, supports the contextualist assumption since it seems better contextualism than the denial of intellectualism. However, the dilemma is specious, consisting in an incorrect reconstruction of how we comprehend stakes-shifting cases.

My main objection to the contextualist assumption consists in my attack on the prominent accounts of the standards for knowledge. A closer scrutiny on all such accounts suggests that the contextualist theory fails to account for our intuitions

concerning knowledge attributions in stakes-shifting cases. Adopting Schaffer's framework of the prominent accounts of the standards for knowledge, I distinguish such accounts into the general and particular accounts.

I first consider the general account, whose tenet is that the shift of the standards for knowledge will have a global impact on the subject's linguistic disposition. I consider two types of general accounts, the linear and spherical accounts. The linear account takes the shift of the standard to be the shift of threshold that determines how strong one's standing on knowing dimensions with respect to p must be in order to be truthfully ascribed as "knows". The spherical account takes the shift of the standard to be the shift of the range of possible worlds that S's belief remains, say, sensitive. Unfortunately, both the linear and spherical accounts fail to account for our intuitions concerning knowledge attributions in *Airport'*. The general moral is that in stakes-shifting cases we do not think that the attributor's knowledge-attributing disposition has undergone a global change so that the attributor would assert "S does not know q" in HIGH for all q with respect to which S's standing on knowing dimensions in LOW is equivalent to S's standing on knowing dimensions with respect to p in HIGH.

I then move on to discuss the particular account, particularly the RA account. According to the RA account, the shift of the standards for knowledge is the shift of the set of relevant alternatives. I consider Lewis's and Blome-Tillmann's accounts of the principles of relevance, including the reliability principle, the actuality principle, the resemblance principle, the attention principle, the high-stakes principle, and the presupposition principle. In one way or another, these principles all fail to account for our

intuitions concerning knowledge attributions in stakes-shifting cases. Hence, the RA account does not seem to explain stakes-shifting cases as well.

All prominent accounts of the standards for knowledge have failed to account for stakes-shifting cases. Some might argue that I did not provide conclusive argument against contextualism here. I agree. Still, the previous discussion rejects, quite conclusively, the contextualist assumption: stakes-shifting cases apparently do not offer *prima facie* support for the contextualist theory.

Things get worse for contextualism (and standards variantism in general) if standing variantism (at least some version of it) is able to explain our intuitions of knowing attributions concerning stakes-shifting cases. I believe such version of standing variantism is indeed available (see Appendix A).

Finally, it is worth mentioning that SSI, which is a type of standing variantism, does not fare better than contextualism in explaining stakes-shifting cases. Recall that according to SSI, S's standing on knowing dimensions with respect to p is determined by two factors: Firstly, S's standing on truth-relevant dimensions with respect to p, and, secondly, S's practical interests, such as stakes, in whether or not p. Consider *Airport*. Smith's stakes in whether *The flight stops in Chicago* is held fixed across LOW and HIGH. Moreover, like contextualists, SSI-ists usually assume that Smith's standing on truth-relevant dimensions with respect to *The flight stops in Chicago* remains constant across LOW and HIGH. As a result, SSI implies that Smith knows that the flight stops in Chicago in LOW if and only if he knows the same proposition in HIGH, *contra* our intuitions concerning *Airport*. Of cause, proponents of SSI have tried to reply to this problem (cf. Hawthorne 2004).

For simplicity's sake, I will not go into the detail of their defense. For the present purposes, it is enough to point out that proponents of SSI are facing two difficult problems: on the one hand, they need to explain why the denial of intellectualism is acceptable. More precisely, since the denial of intellectualism is not necessary in preserving our intuitions concerning stake-shifting cases, what motive do we have in rejecting intellectualism, which is intuitively plausible?[36] On the other hand, proponents of SSI have to explain why our intuitions concerning *Airport* should not be respected.

My suggestion to standing variantists is this: If you do not want to answer these two questions, you should consider the alternative account offered in this dissertation.

---

[36] Stanley (2005) has an answer to this question. For simplicity's sake, however, I will not talk about Stanley's account here.

Chapter Three

Against the Contextualist Solution

3.1 Introduction

The standards-variantist solution (to the skeptical puzzle) is mainly pursued by epistemic contextualists (in their version, the contextualist solution). In fact, the rise of contextualism is mainly motivated by the skeptical puzzle (or philosophical skepticism in general). "Contextualist theories of knowledge attributions," as Keith DeRose points out, "have almost invariably been developed with an eye toward providing some kind of answer to philosophical skepticism." (1995, 4)

Recall that the skeptical puzzle consists of three individually plausible yet jointly inconsistent propositions:

1.1   S knows that op.

1.2   S does not know that not-sh.

1.3   S knows that op only if S knows that not-sh.

'op' stands for ordinary propositions we normally take ourselves to know, while 'sh' stands for skeptical hypotheses incompatible with op.

The contextualist solution presupposes the contextualist theory (of knowledge attributions), i.e., the view that the mere variation of some pragmatic factors in the

speaker context may suffice to shift the standards for knowledge. Contextualists claim that the contextualist theory makes rooms for the contextualist solution.

The core is that the standards for knowledge vary with the ordinary and skeptical contexts such that S's standing on knowing dimensions with respect to op (and to not-sh) satisfies the standard in play in the ordinary context but fails to do so in the skeptical context. As a result, 1.1 and not-1.2 are true in the ordinary context but false in the skeptical one. This further explains why the skeptical puzzle arises in the first place: We intuitively find 1.1 plausible since we normally evaluate 1.1 in the ordinary context and it is indeed true in such context; we intuitively find 1.2 plausible since 1.2 is normally evaluated in the skeptical context, and it is indeed true in such context. In no context, however, 1.1-1.3 are both plausible and true. That explains why the skeptical puzzle is not a genuine paradox.

We have seen, in the last chapter, that the contextualist theory is indeed not well supported; the alleged strongest evidence for the view—stakes-shifting cases—is in fact specious evidence. This greatly undermines the contextualist solution, which presupposes the contextualist theory.

In this chapter, I further consider the question whether the contextualist solution, by itself, is a satisfactory solution to the skeptical puzzle. My conclusion is that it is not.

The failure of the contextualist solution as a solution to the skeptical puzzle has two impacts on the present study. On the one hand, it undermines the prospect of the standards-variantist solution in general, since my argument against the contextualist solution can be extended to cover the standards-variantist solution in general. On the other hand, it motivates proponents of the conciliatory account to investigate into the

standing-variantist solution. As I shall argue in Chapter Seven, a version of the standing-variantist solution is very promising in solving the skeptical puzzle.

3.2 The Main Argument

In this section, I will formulate my main argument against the contextualist solution. I will then use the rest of this chapter to argue for two of its premises.

As a start, notice that contextualists contended that the variation of the standards for knowledge across the ordinary and skeptical contexts was responsible for the variation of the truth-values of the knowledge attribution "S knows that op (and not-sh)". Contextualists generally assume the *invariability* of S's standing on knowing dimensions with respect to op (and not-sh) across the ordinary and skeptical contexts. Here are Stewart Cohen's remarks:

> On the contextualist view, we explain our confidence in the truth of our everyday knowledge ascriptions (the appeal of [1.1]) by supposing that our reasons are sufficient for us to know, relative to the standards of everyday context. When confronted with skeptical arguments however, the chance of error becomes salient and the standards can shift. …In this new context, the standards are stricter and knowledge ascriptions true in everyday contexts are false. So *while the strength of our reasons remains fixed*, the strictness of the standards for how strong those reason have to be varied across contexts. (Cohen 1999, 65–6, my italics)

Similarly, Mark Heller wrote that:

> We start with ordinary, every day [sic.] standards for knowledge. But when confronted with the skeptical challenge our desire to be cooperative conversational partners, or the salience of the possibility of an evil genius, or some combination thereof, changes the standards, pushing the boundary line out to the point at which we can no longer satisfy the standards. *S's epistemic condition is unchanged*, but it is no longer true to say of S that she knows p. (Heller 1999b, 121, my italics)

Let us call the contextualist solution that presupposes the invariability of S's standing on knowing dimensions with respect to op (and not-sh) across the ordinary and

skeptical contexts 'the standing-invariability contextualist solution' ('the SI-contextualist solution' for short). The contextualist solution that presupposes the correspondent *variability* of S's standing on knowing dimensions will be called 'the standing-variantist contextualist solution' ('the SV-contextualist solution' for short).

Most (if not all) traditional contextualists have argued for the SI-contextualist solution. Yet, the SV-contextualist solution is a totally consistent position. Let us not debar contextualists from accepting the SV-contextualist solution. Let us claim that:

> 3.1   As a solution to the skeptical puzzle, either the contextualist solution takes the form of the SI-contextualist solution or it takes the form of the SV-contextualist solution.

3.1, as should be clear, is a tautology.[37]

> Now, I want to argue for the following two premises:

> 3.2   As a solution to the skeptical puzzle, the SI-contextualist solution is epistemically uninteresting.

> 3.3   As a solution to the skeptical puzzle, the SV-contextualist solution is either epistemically uninteresting or theoretically not motivated.

From 3.1-3.3, we can conclude:

> 3.4   As a solution to the skeptical puzzle, the contextualist solution is either epistemically uninteresting or theoretically not motivated.

3.1-3.4 is my main argument against the contextualist solution. This argument, if sound, shows that the contextualist solution is not a satisfactory solution to the skeptical puzzle.

Since 3.1-3.4 is valid and 3.1 is a tautology, the main task of this chapter is to show that both 3.2 and 3.3 are true.

---

[37] The contextualist solution cannot accept the invariability of the standards for knowledge across the ordinary and skeptical contexts.

3.3 An Argument for 3.2

As it turns out, there are different forms of skepticism, and they are not created equal. Some are theoretically more significant than others. A solution to a certain skepticism is *epistemically uninteresting* if the targeted skepticism is not epistemically significant.

For the present purposes, two forms of skepticism are particularly relevant. Recall that if S believes truly that p, whether S can be truthfully described as "knows that p" depends on two factors: One's standing on knowing dimensions with respect to p and the standards for knowledge in play (see Section 1.7). As a result, there are two prominent ways the skeptic and the non-skeptic may disagree with each other. As Richard Feldman observes:

> The different conclusions drawn by the skeptics and non-skeptics could come either from differences in their views about the standards or conditions we had to satisfy in order to have knowledge or from differences in their assessments of the quality or character of reasons we have. (Feldman 1999, 91)

More precisely, the skeptic may challenge the normal assessment of S's standing on knowing dimensions with respect to a certain set of p. The skeptic may argue that, *contra* what the non-skeptic believes, S's standing on knowing dimensions with respect to p is in fact very weak, so weak that it generally fails to satisfy the standards for knowledge. The skeptic may also argue that our standing on knowing dimensions with respect to a certain set of p collapses such that no one is in a better standing on knowing dimensions with respect to p than others do. The skeptic who takes this route may grant that the standards for knowledge are very moderate—S's standing on knowing dimensions need not be very strong in order to satisfy it. Call this form of skepticism 'full-blooded skepticism'.

On the other hand, the skeptic may challenge the commonsensical view about the standards for knowledge. The skeptic may argue that, contrary to what we ordinarily believe, the standards for knowledge are in fact very stringent—so stringent that S's standing on knowing dimensions with respect to a certain set of p generally fails to satisfy it. The skeptic who takes this route can even grant that S's standing on knowing dimensions with respect to p is generally strong. Call this form of skepticism 'high-standards skepticism'[38].

With respect to these two forms of skepticism, we can propose two forms of pluralistic solution. First, one may argue that the full-blooded skeptic is right in claiming 1.2 and not-1.1 in the skeptical context and that the non-skeptic is also right in claiming that 1.1 and not-1.2 in the ordinary context, since S's standing on knowing dimensions with respect to op (and not-sh) varies with such contexts. Second, one may argue that the high-standards skeptic is right in claiming 1.2 and not-1.1 in the skeptical context and that the non-skeptic is right in claiming 1.1 and not-1.2 in the ordinary context, since the standards for knowledge vary with such contexts. Call the former 'the full-blooded solution' and the latter 'the high-standards solution'.

Since the SI-contextualist solution does not allow the variability of S's standing on knowing dimensions across the ordinary and skeptical contexts, the following is obvious:

3.2.1   The SI-contextualist solution is a high-standards solution.

Now, I want to further argue that:

---

[38] The terms 'full-blooded skepticism' and 'high-standards skepticism' are from Kornblith (2000). My definition of full-blooded skepticism, however, is different from Kornblith's.

3.2.2   As a solution to skepticism, the high-standards solution is epistemically uninteresting.

3.2.1 and 3.2.2 imply the desired claim:

3.2   As a solution to skepticism, the SI-contextualist solution is epistemically uninteresting.

Since 3.2.1 is trivial and the derivation from 3.2.1 and 3.2.2 to 3.2 is valid, 3.2.2 is the crucial premise here. By definition, whether 3.2.2 is the case depends on whether high-standards skepticism is epistemically significant.

My (and many others') objection to contextualism is that high-standards skepticism is not epistemically significant. Ironically, contextualists, so long as they uphold the SI-contextualist solution, need to defend on behalf of high-standards skepticism on this score (contextualists do not have to grant that high-standards skepticism is always true; all they need is to show that such kind of skepticism is epistemically significant).

In the next section, I consider philosophers' arguments for the epistemic insignificance of high-standards skepticism and contextualists' replies to them. Some arguments are off the mark, and can be taken care of easily. But we will discover, through our dialectical investigation, the real problem of this form of skepticism. In Section 3.5, a contextualist reply is discussed. I argue that the reply eventually fails.

3.4 The Problem of High-Standards Skepticism

Is high-standards skepticism epistemically insignificant? Hilary Kornblith thinks so. As Kornblith sees it, the high-standards skeptic, who argues that the standards for knowledge

are extremely stringent, has "a deeply deviant view" about "how the term 'knowledge' should be used" (2000, 26). Kornblith then offers an illuminating analogy:

> [The high-standards skeptic is like] the Vermonter who insists that he won't say that it is cold outside unless it is at least 25 degrees below zero Farenheit [sic.]. If he recognizes that there are important distinctions to be made in temperatures above minus 25, and that these distinctions have a bearing on how one should interact with the world, then the only difference between him and us is a bit of charming linguistic eccentricity. (Kornblith 2000, 26)

Kornblith thinks that the skeptic's claim for the high-standards of 'know', like the Vermonter's claim for the high-standards of 'cold', is "a wholly trivial and uninteresting position" (2000, 27). Epistemically significant skepticism should claim that one generally fails to satisfy the ordinary standards for knowledge, not the extremely stringent (non-ordinal) one.

Obviously, if the skeptic is analogous to the eccentric Vermonter, then high-standards skepticism is indeed trivial and uninteresting. However, the skeptic does have a good reply to this objection. Notice that in order for the analogy to hold, the standards for knowledge must not be stringent in ordinary usage, just like the standard of 'cold' is not stringent in ordinary usage. The skeptic may argue that high-standards skepticism is different from the eccentric Vermonter's position, since the standards for knowledge are indeed very stringent in *ordinary* usage.

Consider Peter Unger's (1975) famous discussion of *absolute* terms, whose standards of application are very stringent. Unger's paradigmatic example of an absolute term is 'flat'. To describe something as "flat" is to say that this thing is *absolutely* flat (1975, 54). The same goes for 'certain'. To describe someone as "is certain of p" means, roughly, this person is absolutely certain—*all* doubt is absent in this person's mind (1975, 64). For Unger, knowledge requires certainty. As a result, the standards for knowledge

should be very stringent as well—to describe S as "knows that p" implies that it is not at all doubtful (for S) that p. The skeptic thus can conclude that we normally fail to satisfy the ordinary standards for knowledge, since the ordinary standard is stringent in the first place (this is what DeRose (1995) calls 'bald skepticism').

Unger's account of 'knowledge' breaks Kornblith's analogy between the high-standards skeptic and the eccentric Vermonter—while the Vermonter is linguistically eccentric, the skeptic is not. Kornblith's argument from analogy can be resisted.

It is not to say that one cannot try to resume the analogy, to argue for, say, a moderate ordinary standards for knowledge. Of cause, one can. But doing so is engaging in a dialectical discussion with the skeptic, and this indicates, to some extent, that high-standards skepticism is not a wholly trivial and uninteresting position after all.

Kornblith, however, has another objection to high-standards skepticism. Kornblith claims that the high-standards skeptic is not arguing for something practically significant:

> [I]n practice, if confronted with such a skeptic, it would probably be wise simply to capitulate. "Let us use the term 'knowledge' as you do," I would say. Nothing much hangs on this. Since the skeptic agrees that we can make meaningful and important distinctions about how well justified we are in various claims, and agrees with us about which claims we should believe and act upon, nothing much turns on it. (Kornblith 2000, 26)

The idea here seems to be that high-standards skepticism is insignificant because it has no practical bearing. For the purpose of daily practices, we lose nothing by simply capitulating to the skeptic.

Michael Williams also proposes a similar objection. Arguing against Unger's high-standards skepticism, Williams claims that this kind of skepticism is not "a very serious problem":

This is because it would not point to any deep clash between philosophy and common life, or to any clash at all. Clearly, the discovery that nothing is absolutely flat is not disturbing. This is because, when we claim that a billiard table is flat, we mean that the surface has no irregularities that could disturb the path of the billiard ball; and if we claim that it is *really* flat, we mean flat enough to satisfy the standards of the most exacting player… Accordingly, if skepticism rested on no more than the analogy between "know" and "flat," skepticism would not point to a deep and unsettling truth about the human condition. To preserve the analogy, ordinary knowledge claims would have to stand to justification as ordinary flatness claims stand to flatness, demanding whatever measure is appropriate to the purpose in hand. Absolute certainty, like absolute flatness, would be an idealized condition that we do not expect to find realized in the actual world. (Williams 1996, 49)

What is important is what is "appropriate to the purpose in hand". And since high-standards skepticism does not challenge this, it does not reveal "a deep and unsettling truth about human condition".

While I am sympathetic to Kornblith's and Williams' denial of the practical interest of high-standards skepticism, I do not think this denial renders high-standards skepticism epistemically insignificant. At any rate, skepticism is a matter of *theoretical* consideration. Not only does high-standards skepticism lack practical interests, so does full-blooded skepticism. The investigation of skeptical arguments, as Descartes correctly emphasizes, is a project that "does not involve action but merely the acquisition of knowledge" (1984, 2:15). One should not be surprised if nothing practically significant hangs on our capitulation to the skeptic (high-standards or not).

Perhaps, the problem is that high-standards skepticism is not a kind of epistemic skepticism, since, while epistemology is concerned with knowledge, high-standards skepticism is concerned with 'knowledge'. What the high-standards skeptic argues—that the standards for knowledge are very stringent—is a linguistic claim, and it does not necessarily have any bearing on issues epistemologists interested in. The *contextualist*

fallacy, according to Ernest Sosa, is "the fallacious inference of an answer to a question from information about the correct use of the words in its formulation" (2000, 2). So perhaps high-standards skepticism is only linguistically, or semantically, relevant; to regard it as epistemically relevant is to commit the contextualist fallacy.

This cannot be quite right, however. The discussion of (the usage of) epistemic terms (e.g., 'knowledge') surely is related to the discussion of epistemic properties (e.g., knowledge). The specification of the intension of 'knowledge' naturally requires the specification of the property of knowledge, as shown in Section 1.7. William P. Alston even goes so far as taking them as basically identical:

> Of late a number of theorists have been driving a wedge between what it is to *be* P or what *property* P is, on the one hand, and what belongs to the concept of P or what is the meaning of 'P' on the other. … I shall take it that no such distinction is applicable to epistemic justification, that here the only reasonable interpretation to be given to 'what it is' is 'what is involved in the concept' or 'what the terms means'. (Alston 1985, 83, endnote 1)

At least this much seems clear: The discussion of 'knowledge' is not necessarily irrelevant to the discussion of knowledge. It is not clear that high-standards skepticism commits the contextualist fallacy.[39]

So, what is the real problem of high-standards skepticism? The real problem, as I see it, is that high-standards skepticism is in complete disregard of S's standing on knowing dimensions. Traditionally, skepticism is a challenge of one's reasons, or evidence, for believing in p (i.e., one's standing on knowing dimensions with respect to p). As Feldman writes in the beginning of his (1999):

---

[39] Blome-Tillmann (2009a) argues that it is possible for contextualists to use, as opposed to mention, 'knowledge' when dealing with epistemic issues. He concludes that the fact that contextualism is about 'knowledge' does not necessary debar the theory from having impact on the discussion of knowledge.

> In the good old days, a large part of the debate about skepticism focused on the quality of the reasons we have for believing propositions of various types. Skeptics about knowledge in a given domain argued that our reasons for believing propositions in that domain were not good enough to give us knowledge; opponents of skepticism argued that they are. (Feldman 1999, 91)

High-standards skepticism claims that we generally cannot be described, at least not truthfully, as "knows that op" since the standards for knowledge are very stringent. The high-standards skeptic, however, can grant that the standing on knowing dimensions with respect to op varies from people to people. The skeptic can even grant that our standing on knowing dimensions with respect to op is generally strong. It seems that high-standards skepticism has granted (or can grant) epistemologists everything except the right to use 'knowledge'.

Some previous objections resonate, though now in a more precise and pungent form. Given that the high-standards skeptic can agree with epistemologists about anything except the usage of 'knowledge', high-standards skepticism is just *a verbal dispute for epistemologists*—nothing epistemologically significant hangs on it. High-standards skepticism is thus not epistemologically disturbing. The skeptic does not, to quote Williams, "point to a deep and unsettling truth about the human condition". Epistemologists hardly lose anything by capitulating to the skeptic. In my opinion, this is *the* problem of high-standards skepticism; this explains why it is epistemically insignificant.

Contextualists sometimes seem to misunderstand the challenge here. Epistemically interesting skepticism such as full-blooded skepticism typically challenges the quality of one's reasons for believing (i.e., one's standing on knowing dimensions). Some theorists take the skeptic as arguing that we do not know that not-sh since we have

no evidence whatsoever for the denial of sh (cf. Kornblith 2000). Sometimes, contextualists react to this objection as if the whole point is to argue that since we have no evidence for not-sh whatsoever, our standing on knowing dimensions with respect to not-sh does not even satisfies the moderate standards for knowledge (in the ordinary context)—the contextualist solution of the skeptical puzzle must be wrong. Cohen, for instance, has contrasted high-standards skepticism with what he calls 'global skepticism' (Cohen 1988; 1999). Global skepticism argues that we lack (empirical) evidence altogether for certain sh. Cohen admits that global skepticism poses problem for the contextualist solution, but claims that the problem can be solved by taking (or assuming) that our belief in not-sh is rational. Focusing on global skepticism does not capture the full force of the real problem of high-standards skepticism. The real problem is the disregard of, say, the quality of S's evidence for p (for all p in a certain domain), not so much as we lack evidence whatsoever for not-sh.

The full-blooded skeptic, for instance, can grant that one has some (even non-empirical) reasons for believing not-sh. He can also grant Cohen that belief in not-sh is generally rational. Still, the skeptic may argue that there is no difference in one's standing on knowing dimensions with respect to p, for all p of a certain domain. This contention is more disturbing than the previous one. Epistemologists think that reasons can be good or bad, or that evidence can be strong or weak. It should not be the case that one's standing on knowing dimensions with respect to p is always on epistemic par with another's.

In sum, I have been arguing that high-standards skepticism is epistemically insignificant since it is in complete disregard of one's standing on knowing dimensions, and that epistemically significant skepticism typically challenges the quality of one's

standing on knowing dimensions such as full-blooded. They are, as I see it, two sides of the same coin.

Contextualists might argue that high-standards skepticism can be epistemically significant even if it disregards one's standing on knowing dimensions. Jonathan Schaffer, for instance, agrees that the high-standards skeptic and the non-skeptic is having a "verbal disagreement", and that the skeptical challenge does not collapse our "epistemic standing" (2004a, 152). Still, he insists on the value of high-standards skepticism as revealing the limit of our cognitive capacities:

> Skeptical scenarios concern possibilities that one cannot eliminate. They reveal the limits of our discriminatory range. But the existence of possibilities outside one's discriminatory range does not imply the absence of any possibilities inside the range. … This is why skeptical doubts do not collapse our epistemic standing. (Schaffer 2004a, 152)

Suppose that the stringent standards for knowledge that the skeptic proposes reveal the limits of our cognitive capacities. It sets the upper bound of the strength of one's standing on knowing dimensions with respect to p, for all p in a certain domain. Is this 'revelation' epistemically significant? It depends. If by revealing our cognitive capacities the skeptic shows that our cognitive capacities are generally more problematic (or vulnerable) than we originally thought, the revelation can indeed be epistemically significant. But if, by contrast, all the skeptic shows is just that we are creatures with finite capacities—we are not perfect or ideal agents—then the 'revelation' could hardly be theoretically significant. It is trivial to note that humans are finite agents.

Obviously, by granting that one's standing on knowing dimensions can be as strong as epistemologists want it to be, so long as it does not satisfy the stringent standards for knowledge, the high-standards skeptic is not trying to challenge our

ordinary appraisal of one's standing on knowing dimensions. The revelation of the skeptic, hence, is the trivial note that humans are not ideal agents. Schaffer is right that high-standards skepticism does not collapse one's epistemic standing. But he is wrong in thinking that there is something epistemically significant left in the skeptic's position.

3.5 Neta's Reply

Ram Neta (2003), a contextualist, takes up Feldman's idea that epistemically significant skepticism is about the quality of one's reasons for believing something. Specifically, he grants that "effective skeptical hypotheses" are such that they "involve calling into question the trustworthiness of our evidence" (2003, 19). However, he thinks that high-standards skepticism (at least a form of it) can be epistemically significant (or effective) in this sense. On Neta, when the skeptic claims that we generally cannot be truthfully described as "knows that op", "the skeptic does not employ unusually high standards for knowledge. Rather, she employs unusually restrictive standards of evidence" (2003, 1).

There is a problem, though. That is, Neta does not seem to bother by the distinction between the standard of 'evidence'—the standard that determines whether a certain mental state can be truthfully described by the term "evidence"—and the standard of evidence—the standard that determines whether a certain mental state falls into the category of evidence. He talks about "the standards of evidence" (2003, 1), "what counts as our evidence" (2003, 2), "what l can truthfully regard as my evidence" (2003, 25), and also one "cannot be truthfully be said to have any evidence that [one is] not a BIV" (2003,

25). These expressions are ambiguous: they equivocate between using and mentioning the term 'evidence'.[40]

Let us focus on Neta's more frequently used expression "what counts as one's evidence". Let us consider both use and mention interpretations of this expression. I want to argue, however, that neither of these interpretations can salvage high-standards skepticism from Feldman's worry.

If 'evidence' is *mentioned* in "what counts as one's evidence", the term is more precisely described as "what can be truthfully described by the term 'evidence'"[41]. On this reading, Neta's defense of high-standards skepticism has the skeptic arguing, not for stringent standards for knowledge, but for a stringent standard of 'evidence'—few mental states can be truthfully described as 'evidence'. The skeptic claims that since our mental states ordinarily counted as 'evidence' do not rule out not-sh, we cannot be truthfully described as "knows that not-sh" as well, assuming that knowing requires having (appropriate) evidence, or at least to be counted as "knows that p" requires to be counted as "having (appropriate) evidence".

Is high-standards skepticism, so construed, epistemically significant? I do not think so. When proposing the very stringent standard of 'evidence', the high-standards skeptic (of 'evidence') can grant that people's standings on knowing dimensions with respect to p are generally strong, and that the standing on knowing dimensions with respect to p varies from person to person. If so, the high-standards skeptic (of 'evidence')

---

[40] At first sight, it seems that both interpretations are coherent. At some point, it seems that Neta is equivocating between the use and mention interpretations.

[41] This is how DeRose uses the term "counts as". As DeRose uses it, "S counts as knowing" is synonymous with "It can be truthfully said that S 'knows'" (DeRose 2009, 187, also see footnote 3 on the same page).

is challenging the ordinary usage of 'evidence', not what evidence is, just like the high-standards skeptic (of 'knowledge') challenges the ordinary usage of 'knowledge', not what knowledge is. For the same reason, high-standards skepticism (of 'evidence') is epistemically insignificant—it is not epistemically disturbing.

Let us consider the second interpretation, the interpretation that 'evidence' is *used* in the expression "what counts as one's evidence". So construed, the skeptic is arguing for a very stringent condition for evidence; he is altering (or shrinking) the extension of 'evidence'. The idea, roughly, is that the skeptic tries to disqualifying many things as evidence in order to render "S does not know that not-sh" true. On behalf of contextualists, Peter Klein (2000) has also suggested the high-standards skeptic to modify his view along this line:

> [The skeptic] not only establishes more or less stringent standards for knowledge or the possession of adequate evidence, but also establishes what is to be count as relevant evidence. For example, in the global skeptical case, the utterer can make [not-sh] relevant evidence for [op] (perhaps by merely mentioning it in a serious tone of voice!). (Klein 2000, 112)

If not-sh is a relevant evidence for op, and knowledge requires relevant evidence, then one cannot be truthfully described as "knows that op", since one's body of evidence generally fails to contain not-sh.[42]

So construed, the skeptic *is* actually challenging S's standing on knowing dimensions with respect to p, for all p in a certain domain. Skepticism of this sort is epistemically significant since it implies that the standing on knowing dimensions does not vary from people to people—since everybody's lack evidence of certain sort,

---

[42] Klein's own objection to high-standards skepticism, so construed, is that the skeptic is not entitled to change what count as evidence in the way he wants it to be. For the present purposes, we may put Klein's objection aside. What is important is whether or not high-standards skepticism is epistemically significant, not so much as true. High-standards skepticism can be significant and yet false.

everybody's standing on knowing dimensions is on epistemic par with each other. "High-standards skepticism", so construed, is essentially a kind of full-blooded skepticism.

Still, this does not salvage high-standards skepticism. Skepticism, so construed, is full-blooded skepticism, not high-standards one. Calling a full-blooded skepticism 'high-standing skepticism' does not render high-standards skepticism significant—calling a tail 'leg' would not make a tail have the properties of a leg.

To sum up, I have been arguing that high-standards skepticism is not epistemically significant. Its main problem is that it is in complete disregard of one's standing on knowing dimensions. If I am right, the high-standards solution, which targets on high-standards skepticism, is epistemically trivial. Thus, I have made my case for:

> 3.2.2   As a solution to skepticism, the high-standards solution is epistemically uninteresting.

What is left, in justifying my argument against the contextualist solution, is the third premise:

> 3.3   As a solution to the skeptical puzzle, the SV-contextualist solution is either epistemically uninteresting or theoretically unmotivated.

Let us move on to 3.3.

3.6 An Argument for 3.3

Most (if not all) traditional contextualists do not propose the SV-contextualist solution, since this solution denies that S's standing on knowing dimensions with respect to op (and not-sh) remains constant across the skeptical and ordinary contexts. However, since the SI-contextualist solution is epistemically uninteresting (or so I claim), could proponents of the contextualist solution switch to the SV-contextualist solution instead? I

shall argue, in what follows, that such a move will not serve contextualists any good, either. Before we proceed, it is perhaps worth mentioning that I do not think many will be attracted by the move of shifting from the SI-contextualist solution to the SV-contextualist solution. My main reason for writing this section is to exhaust the most prominent move for a contextualist to modify his view. If one does not find the SV-contextualist solution promising in the first place, one can simply skip this present section.

First of all, notice that the SV-contextualist solution may aim either at high-standards skepticism or at full-blooded skepticism. That is,

> 3.3.1   As a solution to the skeptical puzzle, the SV-contextualist solution is either a high-standards or full-blooded solution.

We have seen from the previous discussion that the high-standards solution is not epistemically significant. Hence,

> 3.3.2   If the SV-contextualist solution is a high-standards solution, it is epistemically uninteresting.

Now, I want to argue that:

> 3.3.3   If the SV-contextualist solution is a full-blooded solution, it is theoretically unmotivated.

From 3.3.1-3.3.3, we can deduce 3.3, the desired result.

The argument from 3.3.1-3.3.3 to 3.3 is valid. 3.3.1, it seems, is trivial. I suspect most contextualists would want to deny both 3.3.2 and 3.3.3. It seems that most contextualists would reject 3.3.2 because they reject 3.2.2:

> 3.2.2   As a solution to skepticism, high-standards solution is epistemically uninteresting.

I have argued, perhaps painstakingly, that 3.2.2 is indeed the case. The crucial premise, hence, is 3.3.3.

3.7 Arguing for 3.3.3

Full-blooded skepticism argues that, for all p in a certain domain, our standing on knowing dimensions with respect to p is generally very weak, and/or our standing generally collapses such that there is no difference in people's standing on knowing dimensions with respect to p.

The SV-contextualist solution is a hybrid account; it consists of the contextualist (and hence standards-variantist) idea—that the standards for knowledge varies with the ordinary and skeptical contexts—as well as the standing-variantist one—that S's standing on knowing dimensions with respect to op (and not-sh) varies with such contexts. On the SV-contextualist solution, the full-blooded skeptic is right in claiming 1.2 and not-1.1 in the skeptical context, and the non-skeptic is also right in claiming 1.1 and not-1.2 in the ordinary context, since both S's standing on knowing dimensions with respect to op (and not-sh) and the standards for knowledge vary with such contexts.

The SV-contextualist solution, allegedly, is a form of contextualist solution. However, if it turns out that the contextualist component of the SV-contextualist solution plays no significant role in explaining the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts, then the legitimacy of the solution—whether we are justified and motivated in developing such kind of solution in the first place—is undermined (hereafter, the variation of truth values of 1.1 and 1.2 are assumed to be generated by the variation of S's standing on knowing dimensions with respect to op and not-sh, since we are dealing with full-blooded skepticism). Why do we need a 'contextualist' solution that does not even use the contextualist idea to solve the problem?

I will call the SV-contextualist solution 'theoretically unmotivated' if its contextualist part is futile in explaining the variation of the truth values of 1.1 and 1.2. The contextualist part is futile if the explanatory power of the SV-contextualist solution is not affected by assuming the denial of it.

Is the contextualist component of the SV-contextualist solution futile in explaining the variation of truth values of 1.1 and 1.2 across the ordinary and skeptical contexts? I think so. Two points are crucial here. One the one hand, the variation of truth values of 1.1 and 1.2 can be explained solely by the variation of S's standing on knowing dimensions with respect to op (and not-sh). On the other hand, the variation of the standards for knowledge is not required for the full-blooded solution. Let me elaborate them respectively.

A moderate amount of reflection suffices to show that the variation of S's standing on knowing dimensions with respect to op (and not-sh) across the ordinary and skeptical contexts is readily available in accounting for the variation of the truth values of 1.1 and 1.2 in such contexts. The skeptic argues that 1.1 and not-1.2 is false since S's standing on knowing dimensions with respect to op (and not-sh) is too weak to satisfy the (moderate) standard for 'knowledge'. The non-skeptic, by contrast, takes S's (presumably moderate) standing on knowing dimensions with respect to op (and not-sh) to satisfy the standards for knowledge. The standing-variantist solution tries to reconcile the skeptic and the non-skeptic by saying that the skeptic is right in the skeptical context while the non-skeptic is also right in the ordinary context. The crux is that S's standing on knowing dimensions is different in these two contexts such that it satisfies the standards for knowledge in one but fails to satisfy the same standard in the other. The standing-

variantist solution thus shows that the variation of S's standing on knowing dimensions with respect to op (and not-sh) alone is sufficient for explaining the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts.

Now, it is not hard to recognize that the variation of the standards for knowledge is not required in the full-blooded solution. It is not required since full-blooded skepticism, which is the target of the full-blooded solution, does not need to argue for the variation of the standards for knowledge. Full-blooded skepticism attacks, as Feldman emphasizes, the quality of one's reasons for believing something, i.e., one's standing on knowing dimensions with respect to a certain proposition. It is not necessarily for the full-blooded skeptic to claim that the standards for knowledge are high or context-sensitive. More importantly, the skeptic lacks motivation for endorsing the variability of the standards for knowledge. Clearly, if full-blooded skepticism lacks motivation to endorse the variability of the standards for knowledge, the full-blooded solution lacks motivation to appeal to the variation of the standards for knowledge as well.

Put in another way, I have argued for two claims:

3.3.3.1   The variation of S's standing on knowing dimensions with respect to op (and not-sh) alone is sufficient for explaining the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts.

3.3.3.2   The variation of the standards for knowledge is not required in the full-blooded solution's explanation of the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts.

3.3.3.1 and 3.3.3.2 strongly suggest that the variation of the standards for knowledge is not required in explaining the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts. It follows that the explanatory power of the SV-contextualist

solution is not affected (at least not seriously diminished) if the solution presupposes the *in*variability of the standards for knowledge.

An analogy is helpful here. Suppose that people score differently in a certain test. Suppose someone tries to explain this fact by the variation of both genetic makeup and education background. Call this explanation 'E'. Suppose that the difference in people's scores can be explained by the variation of education background alone, say, education background alone is a good enough indicator of someone's score of the test, or education background is not worse an indicator than the combination of education background and genetic makeup is. Suppose, also, that there is no motivation for one to explain the variation of the score in terms of the variation genetic makeup. Now, it seems that the explanatory power of E is not affected even if it is modified to deny the explanatory role of genetic makeup in people's score of the test. It is obvious that the variation of genetic makeup is futile in explaining the variation of the score of this test.

If what has been said so far is correct, the contextualist part of the SV-contextualist solution is clearly futile in explaining the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts. Hence, the SV-contextualist solution is theoretically unmotivated. 3.3.3.1 and 3.3.3.2 thus strongly support:

> 3.3.3 If the SV-contextualist solution is a full-blooded solution, it is theoretically unmotivated.

What would contextualists say about 3.3.3? What would be their (possible) reply? I suspect contextualists can grant 3.3.3.2—as noted above, many contextualists are aware of the fact that they are dealing with (some form of) high-standards skepticism, and most of their effort has been dedicated to showing that high-standards skepticism is epistemically significant.

But contextualists probably would deny 3.3.3.1, since most of them would deny the variability of S's standing on knowing dimensions with respect to op (and not-sh) in the ordinary and skeptical contexts—after all, most contextualists endorse the SI-contextualist solution. In other words, contextualists' denial of 3.3.3.1 is based on their denial of standards variantism in general.

The contextualist argument, thus, can be formulated as follows:

3.5   3.3.3.1 is true only if standing variantism is true.

3.6   Standing variantism is not true.

Therefore,

3.7   3.3.3.1 is not true.

3.5-3.7 is a valid argument. The crucial premise, I take it, is 3.6. Hence, whether or not our objection against the contextualist solution is successful depends (at least partly) on whether or not standing variantism is true.

As it turns out, standing variantism is not very popular among philosophers. The reason, perhaps, is that the main form of standing variantism on the market, i.e., SSI, is a radical claim that denies intellectualism. Philosophers from both the contextualist and non-contextualist camps have criticized SSI.[43] In general, philosophers have qualms about the denial of intellectualism.

I also feel the denial of intellectualism implausible. Perhaps, philosophers were even right in rejecting standing variantism if this was the only way to preserve intellectualism. But I think standing variantism does *not* necessarily require us to give up

---

[43] For contextualists' objections to SSI see, e.g., DeRose (2009) and Blome-Tillmann (2009b). For non-contextualists' objections to SSI see, e.g., Blaauw (2008).

intellectualism. Failing to distinguish SSI from other forms of standing variantism, perhaps, is also responsible for philosophers not being eager for standing variantism.

At any rate, the positive thesis of this dissertation is to develop and defend a form of standing variantism that preserves intellectualism.

3.8 Taking Stock

Let me briefly summarize where we are right now. Chapter Two and the present chapter are dedicated to argue against the standards-variantist solution (to the skeptical puzzle). On the standards-variantist solution, the skeptical puzzle could be resolved by appealing to the variability of the standards for knowledge across the skeptical and ordinary contexts. However, proponents of the standards-variantist solution also need to argue that, in general, the standards for knowledge vary with contexts (i.e., standards variantism), otherwise the standards-variantist solution will simply be *ac hoc* (this requirement has been explicitly pointed out by DeRose (2009)).

In Chapter Two, I argue against standards variantism (in particular, the contextualist theory). More precisely, I argue that stakes-shifting cases, which have been commonly regarded as the strongest support (or one of the strongest supports) for standards variantism, have escaped standards variantism.

In the present chapter, I further argued that the standards-variantist solution (in particular, the contextualist solution) has fallen short of offering a satisfactory reply to skepticism, at least not to the kind of skepticism that is epistemically significant. I thus conclude that the standards-variantist solution (the contextualist solution in particular) is implausible as a solution to the skeptical puzzle.

My main argument against the contextualist solution is as follows:

3.1   As a solution to the skeptical puzzle, either the contextualist solution takes the form of the SI-contextualist solution or it takes the form of the SV-contextualist solution.

3.2   As a solution to the skeptical puzzle, the SI-contextualist solution is epistemically uninteresting.

3.3   As a solution to the skeptical puzzle, the SV-contextualist solution is either epistemically uninteresting or theoretically not motivated.

Therefore,

3.4   As a solution to the skeptical puzzle, the contextualist solution is either epistemically uninteresting or theoretically not motivated.

Obviously, 3.2 and 3.3 are crucial premises here. I offer the following argument for 3.2:

3.2.1   The SI-contextualist solution is a high-standards solution.

3.2.2   As a solution to skepticism, the high-standards solution is epistemically uninteresting.

Therefore,

3.2   As a solution to skepticism, the SI-contextualist solution is epistemically uninteresting.

3.2.1 is trivial. The non-trivial claim is 3.2.2. I argue that high-standards solution is epistemically uninteresting because high-standards skepticism is epistemically insignificant. Contextualists have tried to argue that high-standards skepticism is epistemically significant, but if their arguments seem successful or persuasive at all, it is because we have failed to identify the real problem of high-standards skepticism (or so I argue). The real problem of high-standards skepticism is that it is in complete disregard of S's standing on knowing dimensions. Once we identify this problem, it seems obvious that high-standards skepticism is epistemically insignificant.

My argument for 3.3 is the following:

3.3.1   As a solution to the skeptical puzzle, the SV-contextualist solution is either a high-standards or full-blooded solution.

3.3.2   If the SV-contextualist solution is a high-standards solution, it is epistemically uninteresting.

3.3.3   If the SV-contextualist solution is a full-blooded solution, it is theoretically not motivated.

Therefore,

3.3   As a solution to skepticism, the SV-contextualist solution is either epistemically uninteresting or theoretically not motivated.

3.3.1 is trivial and 3.3.2 is a direct result from 3.2.2. Hence, on Section 3.7, I focus on

3.3.3. 3.3.3 is strongly supported by two theses:

3.3.3.1   The variation of S's standing on knowing dimensions with respect to op (and not-sh) alone is sufficient for explaining the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts.

3.3.3.2   The variation of the standards knowledge is not required in the full-blooded solution's explanation of the variation of the truth values of 1.1 and 1.2 in the ordinary and skeptical contexts.

3.3.3.2 is pretty straightforward—it somehow follows directly from the definition of full-blooded skepticism. 3.3.3.1 presupposes standing variantism.

Hence, my argument depends (at least partly) on the correctness of standing variantism. This leads us to the defense of standing variantism. As noted, standing variantism is not very popular among philosophers. One reason, perhaps, is that the only standing variantist theory on the market, i.e., SSI, has been pretty radical. Other reasons, perhaps, have to do with few philosophers noticing that it is possible to develop a form of standing variantism that preserves intellectualism. This is the task of the dissertation.

After revealing the problems of the standards-variantist solution, I turn now to offer my own solution to the skeptical puzzle. I will construct a kind of standing-variantist solution. As noted, the standing-variantist solution presupposes standing variantism (but not the other way around). Hence, my first task is to discuss in depth about the nature of one's standing on knowing dimensions with respect to p. This project will occupy the next three chapters. In Chapter Seven, I will offer my solution to the skeptical puzzle based on the theory of knowledge constructed in Chapter Four to Chapter Six.

Chapter Four

Epistemic Position and Tracking

4.1 Introduction

Ever since Plato, philosophers have recognized that knowledge is more than true belief (cf. Kvanvig 2003). You can believe something truly by chance, but your belief is not knowledge. What is the main difference between mere true belief and knowledge? As addressed in Chapter One, if one believes truly that p, one knows that p if and only if one's standing on knowing dimensions with respect to p satisfies the standards for knowledge. Hence, the difference between merely believing truly that p and knowing that p consists in the subject's standing on knowing dimensions with respect to p.

In what does one's standing on knowing dimensions with respect to p consist? Consider that the cat is on the mat, and I see it. I thus *know* that the cat is on the mat. By contrast, if I do not see the cat but make a wild guess that the cat is on the mat, my belief is true, but falls short of being knowledge. The comparison indicates that mere true belief can be obtained accidentally, i.e., by luck, while knowledge is incompatible with epistemic luck. Knowledge is non-accidental (anti-luck) true belief.[44]

To prevent epistemic luck, one must have good epistemic ground for one's belief. Let us call it one's *epistemic position*. The strength of one's epistemic position with

---

[44] Perhaps not all luck is bad. As some have argued, some kinds of luck is necessary in order for knowledge (Cf. Fogelin 1994).

respect to p is determined by how well one's epistemic ground for p is. What factors determine one's epistemic position with respect to p? Common candidates include "reliability, strength of evidence, epistemic probability, and the like" (Fantl and McGrath 2011, 562). The list can be distinguished into two classes: those that appeal to *external properties* such as reliability, epistemic probability, and those that appeal to *internal* ones such as evidence, justification.[45]

In this chapter, I opt for the externalist account, more precisely, the reliabilist account.[46] One merit for the reliabilist account is that, unlike the internalist account, it can easily account for the epistemic position of animals that are not as cognitively competent as adult humans (cf. Dretske 1971b; Roush 2005).

There are two prominent reliabilist accounts. On the process reliabilist account, one's truth-relevant dimension with respect to p is the degree of reliability of the belief forming mechanism in terms of which one forms the belief that p. On the tracking reliabilist account, one's truth-relevant dimension with respect to p is how good one tracks the truth of p. The former depends on the notion of the reliability of belief forming mechanism and the latter on the notion of tracking. These two reliabilist accounts need not be the anti-thesis of each other; one can be used to illustrate the other (cf. Goldman (2011); Zalabardo (2012, 41–2)).

In the present chapter, I characterize one's epistemic position with respect to p in terms of the notion of tracking. That is, the strength of one's epistemic position with

---

[45] The main difference between internal and external properties, I take it, is that the former, but not the latter, is in principle (consciously) accessible to the subject.

[46] The evidentialist (justificationist) account has been supported by Feldman and Conee (2008). For simplicity's sake, I will not deal with evidentialism here.

respect to p is characterized as how well one tracks the truth of p. When one's epistemic position with respect to p is strong enough for knowledge, call it one is *in a position to know* that p. On my view, one is in apposition to know that p if and only if one tracks the truth of p.

What is tracking? To illustrate, let us consider a paradigmatic case of tracking. You order a gift to your girlfriend online, say, on Amazon. You want to receive the gift by Friday since Saturday is her birthday. On Thursday, you login your Amazon's account and track the location of your order. The table below is listed on your account:

Date/Time	Location	Activities
March 21, 2013 10:55 pm	KANSAS CITY, MO 64121	Processed through USPS Sort Facility
March 20, 2013 3:53 pm	ELK GROVE VILLAGE, IL 60007	Depart USPS Sort Facility

Table 4.1

Now, you feel relief since you *know* that your package is in Kansas on March 21[th] and you should be able receive your order by March 22[th]. Suppose that your package is in Kansas on March 21[th]. How do you know *that*? It seems obvious that your knowledge consists in your capacity to track the whereabouts of your package, or, put in another way, your capacity to track the truth of *The package is in Kansas on March 21[th]*. One crucial feature of tracking is that the information listed in the website must be *sensitive* to the actual location of the package in an appropriate way. For instance, it is not the case

that Table 1 will show that the package is in Kansas regardless of whether the package is in Kansas or not. Yet, the notion of appropriate sensitivity in play is vague, and a more rigorous formulation of the notion is called for.

Historically, Robert Nozick (1981) has famously and influentially formulated the notion of tracking pertaining to knowledge. As we shall see later, Nozick's account has in many ways set the tone for the subsequent development of the tracking account of knowledge. However, Nozick's account of knowledge has suffered from a variety of counterexamples. Many of those counterexamples have been widely considered as conclusive. Nozick's account, many have claimed, is beyond repair. Together with the denial of Nozick's account is the denial of the tracking account in general, since Nozick's account is often taken to be the representative of the tracking account of knowledge. This, however, is a historically unfortunate development. The tracking account of knowledge has captured some important insight regarding the notion of one's epistemic position that survives (almost) *all* counterexamples to Nozick's account. Dismissing the tracking account merely on the basis of the denial of Nozick's account, epistemologists have thrown out the baby with the bath water.

The goal of this chapter is to resurrect the important insight of the tracking account. In what follows, I will first discuss some old tracking accounts, starting from Nozick's one. I will point out three important flaws in Nozick's original formulation of the notion of tracking which have had a great influence on the subsequent accounts. I then propose a tracking account of knowledge that does not suffer from such flaws.

4.2 Nozick's Tracking Account of Knowledge

Nozick (1981) characterizes the notion of tracking in terms of *counterfactual conditionals*[47]. On Nozick's view, S *tracks the truth (of p)* if and only if the following two counterfactual conditionals hold:

Sensitivity. If p were false, S would not have believed that p.

Adherence. If p were true, S would have believed that p.

Tracking, understood as such, specifies one's epistemic position with respect to p. "To know," Nozick said, "is to have a belief that tracks the truth" (1981, 178).

A modest amount of reflection on some cases should be enough to motivate the idea of sensitivity as a necessary condition of the epistemic-status requirement. I will mention three cases here. The first are cases of mere true belief. Suppose that Mary believes, truly, that the president is in New York by asking a magic 8-ball. Obviously, Mary merely believes truly that the president is in New York without knowing it. Arguably, sensitivity can explain this case: Mary's belief is clearly insensitive—if the president were not in New York, Mary would still believe that he is in New York.

Secondly, sensitivity is able to explain Edmund Getteir's (1963) original examples that justified true belief falls short of being knowledge. Suppose that Smith believes that Jones, who is a co-worker of Smith, owns a Ford because John, who is a reliable informant, told him so. As a result, Smith believes that someone in his office owns a ford. It turns out that Jones does not in fact own a Ford, but Mary, who is also Smith's co-worker, does (Smith has not idea that Mary owns a Ford). Smith is justified in

---

[47] By counterfactual conditionals, I mean statements whose typical form is 'If p were the case, then q would have been the case'. It is also common simply to call a counterfactual conditional a 'counterfactual' (a convention I will apply occasionally). Some philosophers prefer to call such statements 'subjunctive conditionals' (or 'subjunctives' for short). There is a consensus that none of these terminologies are perfect. But for lacking a better term, I will stick to my current usage.

believing truly that someone in his office owns a ford, but he does not know it. Sensitivity explains this case nicely: Smith's belief, though true and justified, is not sensitive—if no one in Smith's office owned a Ford, Smith would still believe that someone in his office owns a Ford.

Thirdly, suppose one possesses a ticket for a fair lottery. Intuitively, one does not know that the ticket is a loser, even if the probability of one's ticket being a winner is extremely low, say, 1 out of 10,000,000. The sensitivity condition of knowledge can easily explain this intuition: One's belief that the ticket is a loser is insensitive—if the ticket were a winner, one would still believe that it is a loser.

Enough for sensitivity. How about adherence? Nozick has proposed two cases to motivate this condition:

> *Tank.* If someone floating in a tank oblivious to everything around him is given (by direct electrical and chemical stimulation of the brain) the belief that he is floating in a tank with his brain being stimulated, then even though that fact is part of the cause of his belief, still he does not know that it is true. (Nozick 1981, 172)

Although the person lacks knowledge, his belief that he is floating in a tank is sensitive—if he were not floating in the tank, he would not believe that he is. His belief, however, is not adherent: if he were to float in the tank, he might not believe that he is floating in the tank—it is not hard to imagine that the person in the tank may not be stimulated to believe that he is floating in a tank.

Nozick also borrows an example from Gilbert Harman to illustrate that adherence is necessary for knowledge:

> *Dictator.* The dictator of a country is killed; in their first edition, newspaper print the story, but later all the country's newspapers and other media deny the story, falsely. Everyone encounters the denial believes it (of does not know what to

believe and so suspends judgment). Only one person in the country fails to hear any denial and he continues to believe the truth. (Nozick 1981, 177)

Siding with Harman, Nozick thinks that the person does not know that the dictator was killed. This person's belief is sensitive: If the dictator were not killed, this person would not have believed that he was killed. Nevertheless, his belief does not satisfy adherence: If the dictator were killed, the person might not believe that he was killed—it is not hard to imagine that the person somehow hears the denial, and thus suspends his belief.

Some of the core ideas of Nozick's account of tracking have already argued by some previous works with slightly different formulations (as we shall see later). Yet, I think it is fair to say that Nozick's formulation of the notion of tracking (sensitivity and adherence) has been most influential in contemporary literatures. In what follows, I focus on three features of Nozick's formulation that are particularly relevant to the goal of this chapter.

First, Nozick argues that the tracking relation consists in sensitivity and adherence. As it turns out, many have endorsed the sensitivity condition but not the adherence condition, e.g., Keith DeRose (1995), Tim Black (2008), José Zalabardo (2012). Nonetheless, some side with Nozick in regarding tracking as consisting in both sensitivity and adherence, e.g., Sherrilyn Roush (2005) (Roush rejects Nozick's counterfactual formulation of these two conditions, and suggests to use conditional probability to formulate them). Both sensitivity and adherence have suffered from various (alleged) counterexamples (cf. Roush 2005) (also see Section 4.3 and Chapter 5).

Second, Nozick's formulation of tracking (sensitivity and adherence) has two distinctive features: (a) to characterize tracking in terms of counterfactual conditionals and (b) to treat tracking as a property of belief. Call (a) and (b) the 'counterfactual-

conditional' and the 'belief' formulation respectively. Some subsequent discussions adhere both of the counterfactual-conditional and the belief formulation in characterizing epistemic position (e.g., DeRose (1995); Sosa (2000)). A few have adopted the counterfactual-conditional formulation but not the belief one (e.g. Dretske (1971a)). Still, some have adopted the belief formulation while rejecting the counterfactual-conditional formulation (e.g., Roush (2005); Zalabardo (2012)).

Third, Nozick (and Dretske) has argued against:

(Closure)   If S knows that p and that p entails q, then S knows that q.

On Nozick's view, (Closure) is incompatible with sensitivity as a necessary condition of knowledge. Let us elaborate.

Nozick claims that sensitivity is not close under known entailment; a proposition known to be entailed by a sensitive belief may not be sensitive. Suppose that you are in a zoo and see that there is a zebra in a pen. Presumably, your belief that there is a zebra in the pen is sensitive: If there were no zebra in the pen, you would not believe that there is a zebra in the pen. Now, suppose that you also know that *There is a zebra in the pen* entails *The creature in the pen is not a mule painted to look like a zebra*. But, presumably, your belief that the creature in the pen is not a mule painted to look like a zebra is insensitive: If the creature in the pen were a mule painted to look like a zebra, you would still believe that there is a zebra in the pen. Hence, sensitivity is not closed under known entailment.

However, according to Nozick's account, sensitivity is a necessary condition for knowledge; the epistemic-status requirement consists partly in one's belief being

sensitive. It follows that knowledge could not be closed under known entailment as well, namely, (Closure) must not hold.[48]

I think that Nozick's idea that knowledge consists in tracking (or, more precisely, sensitivity) captures a very important feature of knowledge. However, as I see it, Nozick's development of the tracking account is unsatisfactory. In particular, the aforementioned Nozick's formulations of tracking are essentially incorrect and/or misleading. My own tracking account is different from Nozick's in the following aspects: (a) I only take sensitivity, but not adherence, as necessary for knowledge. (b) While I adopt the counterfactual-conditional formulation—the view that characterizes tracking (sensitivity) in terms of counterfactual conditionals—I deny the belief formulation—the view that takes tracking (sensitivity) as a property of belief. (c) I also reject the alleged incompatibility between tracking (sensitivity) and (Closure).

In the following three sections, I will discuss the adherence condition, the belief formulation, and (Closure) respectively. During the course of our discussion, we should be able to see why (a)-(c) are required (or desired) by the tracking account of knowledge.

4.3 The Adherence Condition

Nozick's rationale for the adherence condition, as I see it, is to use this condition to capture the intuitive idea that knowing that p requires that the belief that p is true non-

---

[48] Both Nozick and Dretske claim that the failure of (Closure) provides gives room to a nice solution to the skeptical puzzle. The idea, as noted in Section 1.4, is that, on the one hand, S generally knows that op (i.e., 1.1), since (other things being equal) S's belief that op is sensitive: If op were not the case, S would not believe that op. On the other hand, S generally does not know that not-sh (i.e., 1.2), since (other things being equal) S's belief that not-sh is insensitive: If sh were the case, S would still believe that not-sh. The skeptical puzzle does not arise since 1.3, i.e., if S knows that op, then S knows that not-sh, is false. It is false precisely because its alleged rationale, i.e., (Closure), does not hold. In other words, the denial of (Closure) has been taken as the key to the solution to the skeptical puzzle.

accidentally. Recall Nozick's cases for the adherence condition mentioned above. In *Tank*, the fact that the person's belief that he is floating in a tank is true is accidental—Nozick clearly wants to interpret the case as the person could easily have a different (false) belief or no belief at all. The same goes for *Dictator*. In this case, the person believes something truly by sheer luck—he could have easily had a false belief or suspended judgment. The adherence condition has it that if p were true in situations (enough) similar to the actual situation, then S would believe that p in those situations.[49] This guarantees that the truth of p and S's belief that p co-exist in all situations (enough) similar to the actual one. Perhaps the underlying idea is that such a co-existence is enough to block the possibility that one's belief being true by accident.

I agree that non-accidentality is a character of knowledge. But I disagree that the adherence condition captures this character. The reason is simple: The adherence condition is simply not a necessary condition of knowledge. Saul Kripke (2011) has made a very powerful case for this point. But I have come to realize that the point has been pointed out by David Armstrong (1973) long before Nozick published his *Philosophical Explanations*.

Armstrong proposes what he calls "the 'Thermometer' view (of non-inferential knowledge)". The idea behind this view is to compare non-inferential knowledge, such as perceptual knowledge, to the temperature-readings given by a thermometer:

> In some cases, the thermometer-readings will fail to correspond to the temperature of the environment. Such a reading may be compared to non-inferential false belief. In other cases, the reading will correspond to the temperature of the environment. Such a reading is like non-inferential true belief. The second case, where reading and actual environmental temperature coincide, is then sub-divided

---

[49] Notice that this is a bit different from Lewis's semantics (Lewis 1973a; also see Chapter Six). Williamson (2000) and Sosa (2000) have adopted Nozick's interpretation.

into two sorts of case. First, suppose that the temperature is a bad one, but that, on a certain occasion, the thermometer-reading coincides with the actual temperature. (Cf. the stopped clock that shows the right time twice a day.) Such a reading is to be compared with non-inferential true belief which falls short of knowledge. Suppose finally that the thermometer is a good one, so that a reading of 'T°' on the thermometer ensures that the environmental temperature is T°. Such a reading is to be compared with non-inferential *knowledge*. (Armstrong 1973, 166, original italics)

Focusing on the analogy between the temperature-readings given by a thermometer and one's believing that p in terms of p being the case, Armstrong manifests that his account is based on the idea of tracking—after all, it is very natural to interpret that thermometer tracks the temperature of the environment. Also, like Nozick, Armstrong endorses the belief formulation—the tracking involved in knowledge is a property of one's belief.

However, unlike Nozick, Armstrong does not mistakenly regard adherence as necessary for knowledge. Instead, he has explicitly denied that adherence is a necessary condition for knowledge:

[One feature of a reliable thermometer is:] given that the environmental temperature is T°, then the good thermometer will register 'T°'. Now this … feature is not necessary for non-inferential knowledge. Given that a person knows that c is J non-inferentially, then his belief is a reliable one, in the way which we have spelt out. He is in some condition and/or circumstances H …. such that if a person is in H, and believes something of that sort, then, as a matter of law-like connection, that thing believed is the case. But this does not follow that if a person is in circumstance/condition H, and c is J, then that person believes that c is J. Such a person may have no opinion at all on the matter. He may fail to register 'Jc'. (Armstrong 1973, 174–5)

As Armstrong points out, adherence is a good epistemic feature: One of the features of a reliable thermometer is that it is adherent—if the temperature is T°, the thermometer will register 'T°'. But this is not a feature that is necessary for knowledge. Admittedly, it suffices for one to know that p if one believes p *because* of the truth of p. However, it

does not follow that if p is (were) true, one has (would have) to believe (and thus know) p.[50]

I think the point can be generalized; adherence is at any rate irrelevant to knowledge. The following cases illustrate this point nicely:

> *Fan A*. I have a friend, John, who is a big fan of the Lakers and will watch every game of the team. If the Lakers win a game, John will call me to tell me the news. Hence, I always know whether or not the Lakers win a game. Now, suppose that John called this morning telling me that the Lakers won yesterday's game.

> *Fan B*. Like case *A*, John is a big fan of the Lakers and will watch every game of the team. Unlike the former case, John does not always call me to tell me the result of the game. He only does that occasionally. Hence, I do not always know whether the Lakers have won a game or not. Now, suppose that John called this morning telling me that the Lakers won yesterday's game.

Intuitively, I know that the Lakers won yesterday's game in both *Fan A* and *Fan B*. But my belief that the Lakers won yesterday's game is adherent in the first but not the second case. This should be enough to show that adherence has nothing to do with knowledge. The most important point of *Fan* is that it clearly shows that John's testimony in *Fan A* and *B* is equally good *as far as knowledge is concerned*. If I am to know that the Lakers won yesterday's game via John's testimony, my knowledge in *Fan A* is on epistemic par with my knowledge in *Fan B*, regardless of my belief that the Lakers won yesterday's game is adherent in the first case but not in the latter. The adherence condition, thus, plays no significant role in characterizing whether one is in a position to know.

---

[50] Kripke has argued the same point by the following counterexample:
> *Photon*. Suppose that Mary is a physicist who places a detector plate so that it detects any photon that happens to go to the right. If the photon goes to the left, she will have no idea whether a photon has been emitted or not. Suppose a photon is emitted, that it does hit the detector plate (which is at the right), and that Mary concludes that a photon has been emitted. Intuitively, it seems clear that her conclusion indeed does constitute knowledge. But is Nozick's [condition of adherence] satisfied? No, for it is not true, according to Nozick's conception of such counterfactuals, that if a photon had been emitted, Mary would have believed that a photon has been emitted. The photon might well have gone to the left, in which case Mary would have had no beliefs about the matter. (Kripke 2011, 178)

Admittedly, John's testimony in *Fan A* is somehow *better* than the one in *Fan B*, since I would have more knowledge in *Fan A* than I would have in *Fan B*. But this desirable feature of John's testimony is irrelevant to whether the quality of John's testimony is sufficient for giving rise to knowledge—in both cases, John's testimony is *equally* good in rendering my true belief knowledge.[51]

How about the examples Nozick uses to motivate the adherence condition (see Section 4.2)? Are those examples misguided? I agree that those examples are not essentially misguided. But admitting this does not force us to accept adherence. At any rate, the intuitive ideas prompted by such examples are not adherence.

4.4 The Belief Formulation

The belief formulation of tracking takes tracking as a property of belief. In contemporary literatures, this formulation has been widely adopted (Dretske (1971a) is a rare exception. I will discuss Dretske's account later). Many proponents of sensitivity and/or adherence have formulated sensitivity and adherence as a property of belief (e.g. DeRose (1995) and Black (2008)). In two recent works, Roush (2005) and Zalabardo (2012) both reject the counterfactual-conditional formulation. Nonetheless, they still accept the belief formulation of sensitivity and adherence. Some philosophers claim that neither sensitivity nor adherence is necessary for knowledge. They argue, instead, that knowledge requires safe belief:

---

[51] The same point holds even for Roush's (2005) formulation of adherence, which relies on conditional probability.

> *Safety*. S's belief that p is safe if and only if: If S were to believe that p, p would be true.[52]

Safety, as an alternative to sensitivity and adherence, is still a property of belief.

I think the belief formulation is either mistaken or misleading (getting things backward). Why has the belief formulation been so widely adopted by philosophers with such widely different views on epistemic position? One important reason, as I see it, is that the tracking account has at the very beginning been developed as a solution to the Gettier problem. Ever since Gettier points out that justified true belief is not sufficient for knowledge, it has been widely accepted that there is a logical gap between (justified) true belief and knowledge (in particular, a gap between true belief and knowledge). Philosophers have since then been zealous in searching for the condition that turns (justified) true belief into knowledge. The (various versions of the) tracking account has (have) been developed with an eye toward the Gettier problem. Nozick, for instance, takes sensitivity and adherence to be individually sufficient and jointly conditions for turning one's true belief into knowledge. DeRose takes sensitivity alone to be the condition that does the same job. Sosa, while denying sensitivity, takes safety to be a necessary condition for the condition that turns one's true belief into knowledge. In other words, since tracking has been widely regarded as the condition that 'bridges' the logical gap between true belief and knowledge, it is thus natural to also regard tracking as a property of belief.

A closer scrutiny shows that the belief formulation is not supported by the discussion of the Gettier problem, since tracking alone (sensitivity and/or adherence) is not enough to turn one's true belief into knowledge. First of all, as Fogelin (1994) points

---

[52] This is Sosa's (2000) formulation. For other formulations of the safety condition, see Rabinowitz (2011).

out, a closer scrutiny of the Gettier problem shows that there are two conditions determining one's standing on knowing dimensions with respect to p: (i) how strong one's epistemic position with respect to p is and (ii) whether one is epistemically responsible in forming one's belief that p. Let me call (i) and (ii) the *epistemic-position* and *doxastic-responsibility requirement (of knowledge)* respectively (also see Goldman 1988). These two requirements can come apart.

In Gettier's original paper, he considers mainly examples that satisfy the doxastic-responsibility requirement but not the epistemic-position requirement. That is, examples that the subject has been epistemically responsible in forming his belief and yet his belief fails to be grounded adequately (Fogelin 1994). Consider the example of Smith's true belief that someone in his office owns a Ford (see Section 4.2). Smith believes so since he is justified in believing (falsely) that John owns a Ford. Yet, unbeknownst to Smith, Mary, who is a co-worker of Smith, in fact owns a Ford. In this case, Smith has been epistemically responsible in forming the belief that someone in his office owns a Ford, since he is justified in believing that John owns a Ford. Yet, Smith lacks knowledge, since his ground for *Someone in his office owns a Ford* is not adequate—particularly, it consists of a false belief.

As the discussion of the nature of knowledge proceeds, philosophers have also constructed cases showing that forming one's belief in an epistemically responsible way is also necessary for knowledge. Consider Laurence BonJour's famous example:

> *Norman.* Norman, under certain conditions which usually obtain, is a complete reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, although he has no evidence either for or against this belief. In the fact the belief is true and results

from his clairvoyant power under circumstances in which it is completely reliable. (BonJour 1985, 41)

Intuitively, Norman does not know that the President is in New York City in this case. Yet Norman satisfies the epistemic-position requirement. Since Norman, arguably, fails to satisfy the doxastic-responsibility requirement, this case shows that the latter is also necessary for knowledge.

If the epistemic-position and doxastic-responsibility requirements are both necessary for knowledge, it seems that there are indeed two steps in bridging the gap between (true) belief and knowledge. The first step calls for the specification of whether one's epistemic position is good enough for knowledge, or whether one is in a position to know. The second step requires the specification of whether one forms the belief that p by an epistemically responsible way, e.g., whether one forms the belief on the basis of one's epistemic position with respect to p. Hence, contrary to Nozick's contention, tracking alone is not enough for bridging the logical gap between true belief and knowledge.

More importantly, there is a principled reason not to treat tracking as a property of (true) belief. In his *Evidence and Evolution*, Elliott Sober has distinguished three questions regarding the concept of evidence (Sober cites these questions from Richard Royall (1997)) (Sober 2008, 3ff.):

(1) What does the present evidence say?

(2) What should you believe?

(3) What should you do?

It is clearly that these questions are of different 'levels'. (1) is a question about the quality of one's evidence: whether one's evidence supports p, to what extent one's evidence

confirms (or disconfirms) p, etc. (2) is a question about whether one should form a belief (given one's evidence). For instance, whether one should believe that p given that one's evidence confirms that p to the degree d. (3) is then a question about what one should act given one possesses a certain belief: Is it rational to do e given that one believes that p? These three are different questions and should be treated differently.

Obviously, the same applies to the issue of knowledge. We may thus distinguish the following three questions:

(1)*   What is one's epistemic position with respect to p?

(2)*   Should one believe that p (when based, appropriately, on one's epistemic position with respect to p)?

(3)*   What should one do (given that one believes that p based, appropriately, on one's epistemic position with respect to p)?

Likewise, (1)* to (3)* are distinct. (1)* is the question about one's epistemic position: How strong one's epistemic position with respect to p is, whether one's epistemic position is good enough or not, etc. (2)* is a question of what one should believe given one's epistemic position with respect to certain propositions. (3)* is a question about what one should do given one's belief and/or knowledge.

(2)* and (3)* correspond to the epistemic-position and doxastic-responsibility requirements respectively. The reason that these two requirements may come apart is precisely that (2)* and (3)* are questions of different levels. As far as tracking is necessary for one being in a position to know, tracking is a property of one's epistemic position. That is, the tracking account is designed to answer (1)*, rather than (2)*, nor (3)*. Whether one tracks the truth of a certain proposition determines whether one is in a position to know that p. But one can be in a position to know that p without believing

(and hence knowing) that p. Let alone to act on the basis of believing and/or knowing p.

To sum up, the belief formulation conflates (1)* and (2)*, treating what should be an

account of (1)* to be an account of both. On pain of conflating questions of different

levels, we should not treat tracking as a property of (true) belief.[53]

---

[53] The fact that the belief formulation conflates two levels of questions is manifested by Nozick's further refinements of sensitivity and adherence. Recall that Nozick's characterization of sensitivity and adherence:

> *Sensitivity*. If p were false, S would not have believed that p.
> *Adherence*. If p were true, S would have believed that p.

Nozick has soon realized that these conditions, formulated as such, call for refinement. Consider the following case:

> *Grandmother*. A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. (Nozick 1981, 179)

Intuitively, the grandmother knows that the grandson is well. But her belief that the grandson is well is not sensitive: if he were sick, she would still believe that he is well.

Adherence also faces the same type of problem. Nozick talks about the following case:

> *Robber*. [T]he bank-robber's mask slips off as he is escaping and the bystander sees it is Jesse James, whose picture is on many wanted posters. (Nozick 1981, 193)

Intuitively, the bystander knows that Jesse James is robbing the bank. But the bystander's belief is not adherent: If Jesse James were robbing the bank, it might be that the bystander did not believe that Jesse James were robbing the bank—the reason is that in all similar circumstances where Jesse James were robbing the bank, his mask might not slip off as he was escaping.

To accommodate cases like these, Nozick suggests that we should relativize sensitivity and safety with *method*. Suppose that S believes truly that p via method M in the actual case. The tracking conditions can thus be refined as follows:

> *Sensitivity**. If p were false, S would not believe that p via M.
> *Adherence**. If p were true, S would believe that p via M.

The grandmother's belief that the grandson is well is insensitive, but sensitive*: Presumably, the grandmother believes that the grandson is well via seeing. However, if the grandson were not well, she would not believe that he is well via seeing (instead, she would believe that he is well via others' testimony).

By the same token, although the bystander's belief that Jesse James is robbing the bank is not adherent, it is adherent*: In the actual situation, the bystander believes that Jesse James is robbing the bank by the method of seeing his face. In cases where his mask did not slip off, the bystander would not fail to believe that Jesse James is robbing by using the same method, i.e., seeing his face.

Taking tracking as a property of belief immediately gives rise to the requirement of method. Anyone who has seriously pondered how to fresh out sensitivity and/or adherence and/or safety along this line soon realizes that the relativization to method is called for in order to deal with cases like *Grandmother* and *Robber* (cf. DeRose (1995), Black (2008), Sosa (2000)). If tracking requires that one's belief tracks the truth, it requires that one's belief tracks the truth via method M.

Why is method-relativization necessary? Because the belief formulation conflates questions of two levels—the question about one's epistemic position and the question about whether one should form a particular belief on the basis of one's epistemic position. Given that there are in fact two questions, as opposed to one, involved, it should be no surprise that some condition is required to connect them. This is exactly what we find. In *Grandmother*, for instance, the grandmother believes that the grandson is well in the actual situation as well as in the counterfactual. But her belief is formed on different basis in these two situations. Formed on the basis of her perception, her belief is sensitive, while formed on the basis of

Although the belief formulation has been widely adopted, Fred Dretske's account of knowledge has been a noticeable exception. Dretske's (1971) account of knowledge has explicitly specified sensitivity as a property of one's epistemic position. Like Nozick, Dretske also think that sensitivity is essential to knowledge. But, unlike Nozick, Dretske does not regard sensitivity as a property of one's belief. For Dretske, knowledge requires conclusive reasons:

(D)   If S knows that p solely on the basis of R, then R is a conclusive reason for p.

Conclusive reasons are characterized in terms of sensitivity:

(R)   If R is a conclusive reason for p, then R would not be the case unless p were the case.

For Dretske, R is a conclusive reason does not imply believing: "R may be a conclusive reason for believing P without anyone believing P, much less having R as his reason for believing" (1971a, 17). Sensitivity is a property of one's (conclusive) reason, not belief. Dretske's formulation of sensitivity, published before Nozick's, is in fact a more promising account, since it does not conflate two levels of questions. It is historically unfortunate that Dretske's formulation is not well adopted by the philosophical circle.

At this point, some readers may complain that I have made a fuss over the belief formulation. The readers might argue that the belief formulation, while strictly incorrect, does no real harm: Philosophers have been careless in assuming that sensitivity (or

---

others' testimony, insensitive. The notion of method plays the role to make sure that one's belief is formed on the same basis in both situations.

The introduction of method renders Nozick's account tediously complicated. It plays a very weird role as well: It seems that in most case, the relativization of method does not bother us. They are necessary only in cases like *Grandmother*. It should be noted that this move is not mandatory at all if we keep the distinction between questions (1)* and (2)* straight in the first place. That is, if we take tracking as the property of one's epistemic position (rather than one's belief), then cases like *Grandmother* ceases to become a problem at all (see Section 5.4).

tracking) is a property of belief, but the arguments for or against sensitivity (or tracking) can be carried over into the correct formulation such as Dretske's without losing any significant content; this might also partly explain why philosophers tend to be careless in this way—why be over cautious when no harm could be done?

Those who do not think that the denial of the belief formulation is a big deal is perhaps unaware of how prevalent and influential the formulation has been in the current literature. If the aforementioned discussion appears to be too intricate and tedious, it is because I want to give the fair hearing to the widely accepted belief formulation. Moreover, it is wrong to think that the belief formulation is a harmless assumption in the first place. The adaptation of the belief formulation has far reaching ramifications. Noticeably, many have constructed 'counterexamples' to the tracking (sensitivity) account of knowledge without realizing that the alleged counterexamples immediately lose their teeth once the belief formulation is discarded. Given that the general impression about the tracking (sensitivity) account is that it suffers from insurmountable counterexamples, the endorsement of the belief formulation in fact mistakenly and prematurely suggests (at least partly) the doom of the tracking (sensitivity) account. (I will discuss how the denial of the belief formulation frees us from some alleged, yet prominent, counterexamples to the tracking (sensitivity) account in the next chapter).

4.5 Tracking and Closure

Both Nozick and Dretske reject closure, as addressed above. Nozick in particular thinks that the denial of (Closure) is a plus to the tracking account, given that it allows him to explain the skeptical puzzle. But he does not seem to bother to explain why the

plausibility of (Closure) should not be respected in the first place. Dretske tries to argue that (Closure) is only plausible superficially. He proposes counterexamples to (Closure) intending to deprive the principle of its plausibility. Many have thought that those examples do not have a real bite on (Closure) (cf. Vogel 1990; Stine 1976).

At any rate, what (Closure) tries to capture is the intuitively plausible idea that "deduction is a way of extending one's knowledge" (Williamson 2000, 17). Neither Nozick nor Dretske has offered convincing reasons against this idea.

In fact, both Nozick and Dretske fail to recognize that the denial of (Closure) is at any rate not necessary for the tracking account to resolve the skeptical puzzle (e.g., DeRose 1995; Cohen 1988; 1999).

Recently, Roush has proposed to incorporate (Closure) into the tracking account:

> Closure is good, and tracking is good, and though no actual theory has yet put the two features together, that does not imply that it is logically impossible to do so. (Roush 2005, 42)

The resulting account, roughly, is as follows:

(K_R)   S knows that p iff (i) S believes truly that p, and (ii) S tracks the truth of p or (iii) for some q such that q implies p, S knows that q, and that q implies p.

Roush calls her view "the recursive tracking view of knowledge"[54]. The term 'recursive' does not modify the word 'tracking', but the word 'view' (Roush 2005, 47). The idea is to define knowledge recursively, with known implication as the mechanism of recursion (Roush 2005, 28).

---

[54] Roush's account of tracing is inspired by Nozick's but different from the latter in significant ways. More precisely, she uses the notion of conditional probability, rather than counterfactual conditionals, to account for tracking. For simplicity's sake, I will not deal with Roush's account here.

(K_R) implies that if S knows that p, then either S knows that p by tracking the truth of p or by knowing some propositions that imply p, assuming that the implication is also known. The recursive tracking view has the advantage of preserving the idea that tracking is essential to knowledge while at the same time respecting the intuitive idea that "deduction is a way of extending one's knowledge". The only cost of having this nice theory, according to Roush, is that it does not offer a solution to the skeptical puzzle:

> The recursive tracking view does not settle the skeptical question… According to this view of knowledge I may know that there is a table in front of me, in which case I also know that I am not a brain in a vat (by known implication), or I may not know that there is a table in front of me, because I do not know that I am not a brain in a vat. Recursive tracking does not determine which of these positions one must adopt, but does tell us what my knowing that I am not a brain in a vat depends on, and thus why my situation is frustrating [i.e., that I am in no position to refute the skeptic]. (Roush 2005, 55)

The reason why Roush's view does not settle the skeptical puzzle is that she takes knowledge to be invariant, namely, knowledge does not manifest any epistemically interesting feature of context-sensitivity. Hence, Roush's view predicts that either skepticism or common sense is right. Either I know that there is a table in front of me or I do not. In principle, the tracking view does not tell us which one we should pick.

Roush is satisfied with the idea that the recursive tracking view does not seem to settle the skeptical puzzle. "[T]he question what to say about skeptical hypotheses should not," Roush said, "drive one's view of what knowledge is" (2005, 57). I do not agree with Roush. To me, it is a legitimate move to shape our view of knowledge on the basis of skeptical considerations.

More importantly, while Roush wants to remain neutral to the skeptical puzzle, the recursive tracking view is by no means neutral in this sense. Notice that the recursive view, if correct, would mandate the appropriate solution to the skeptical puzzle to take

the form of what I call 'the repudiating account' (see Section 1.3). The repudiating account claims that (at least) one of our initial intuitions regarding 1.1-1.3 is false, since the recursive tracking implies that either 1.1, i.e. S knows that op, or 1.2, i.e., S does not know that not-sh, is false, the theory also implies that our initial intuitions regarding the false statement are fallacious.

I have pointed out that the repudiating account faces two difficult problems (see Section 1.5). On the one hand, it is very hard to explain what is wrong with our initial intuitions regarding 1.1 and 1.2. On the other hand, it is unpromising to explain away our intuitions as commitment of cognitive error of some sorts. These two general problems are not insurmountable. But they do indicate that not only the recursive view is not neutral to the skeptical puzzle, the view actually faces some serious problems when accounting for the puzzle. Unless Roush also wants to argue that a view of knowledge does not need to be able to resolve the skeptical puzzle altogether, the recursive tracking should be evaluated (at least partly) against its prospect of solving the skeptical puzzle.

I think the recursive tracking view can be improved by incorporating some variantist factors into the view. But we need to have some principled way to justify this move (I will take up this task in Chapter Seven).

4.6 The Tracking-Based Account

Let me take stock. I think that the notion of tracking has captured some significant insight regarding the notion of knowledge. In the previous sections, I have considered and objected to Nozick's influential tracking account. Specifically, I argue that (a) tracking does not consist in adherence and that (b) it is not a property of belief. I also think that (c)

Roush's innovative proposal of incorporating (Closure) into the tracking account is promising.

Learning from the morals from the previous discussion, I propose, tentatively, a new recursively defined tracking account of knowledge:

(K$_T$)   S knows that p if and only if

    (i)   S believes truly that p on the basis of one's (sole) reason for p. And

    (ii)   if R is S's sole reason for p, R tracks the truth of p, or

    (iii)   for some q such that q implies p, S knows that q and that q implies p.

Some clarifications are called for. The satisfaction of (K$_T$)(i) guarantees the satisfaction of the doxastic-responsibility requirement. For the purposes of this dissertation, I will put aside the question of the nature of this requirement. (K$_T$)(ii) discards the belief formulation; tracking is regarded as a property of one's reason R, rather than one's belief. Like Dretske, I take it that R can be one's reason for p even though one does not believe p on the basis of R (or does not believe p at all). Hence, (K$_T$)(ii) does not conflate questions of different levels as the belief formulation (at least when taken at face value) does. As we will see in next chapter, this allows (K$_T$) to avoid many alleged counterexamples to the tracking (sensitivity) account. (K$_T$)(iii) takes up Roush's proposal of a recursive definition of the notion of knowledge. It consists of a base case, specified by (K$_T$)(i) and (K$_T$)(ii), and an inductive clause, specified by (K$_T$)(iii). Like Roush's, the present account incorporates (Closure) into the inductive clause, respecting the intuitive idea that "deduction is a way of extending one's knowledge". However, one disadvantage of Roush's account is that it does not take induction into account. Given that induction has played a much larger role in our ordinary acquisition of knowledge than deduction

does, an account of knowledge that dismisses induction is intolerable. To right this wrong, we will have to modify (K_T)(iii) to incorporate not only deduction, but induction as well in the inductive clause.

However, before we get to that, it is worth noting that, although widely adopted, (K_T)(iii), as a formulation of (Closure), is less than satisfactory on two scores. First, as John Hawthorne correctly points out, "I know that p and know that p entails q, I may still have to *do* something—namely perform a deductive inference—in order to come to know that q. Until I perform that inference, I do not know that q" (Hawthorne 2004, 32). Secondly, even if one performs the required inference, one may still fail to know that p if one does not thereby form a belief in p, or if one fails to maintain one's knowledge of p throughout the entire process (Hawthorne 2004, 32-4).

To overcome these difficulties, I will thus reformulate (K_T)(iii), tentatively, in the way suggested by Hawthorne (cf. Hawthorne 2004, 34):

(K_T)(iii)*   for some q such that q implies p, S knows that q and competently deduces p, and thereby comes to believe p, while retaining knowledge of q throughout.

In general, (K_T)(iii)* is sufficient for the purpose of capturing the idea that knowledge is closed under known entailment. However, it may not be enough for an account of knowledge, as we not only think that "deduction is a way of extending one's knowledge", but also that induction is (in many cases) a way of extending one's knowledge. If the principled reason for incorporating (Closure) into the tracking account is to square well with the way we gain knowledge, we should also take induction into account as well.

Let us use the term 'inference' to denote both deductive and inductive reasoning. An inductive inference may be strong or weak, but even a strong inductive inference is not truth-preserving; the premises of a strong inductive inference do not entail its conclusion. I will call a strong inductive inference *correct* if its premises in fact lead to the truth of its conclusion—a correct induction is an inductive inference that works without a glitch. Strong inductive inferences have a higher chance of being correct, but there is no guarantee that a strong inductive inference is always correct. An inference is called *adequate* if it is either a valid deduction or a correct (strong) induction. With these terminologies at hand, I propose that we formulate a more general idea of (Closure) as follows:

$(K_T)(iii)'$ for some q, S knows that q and competently and adequately infers p from q, and thereby comes to believe that p, while retaining knowledge of q throughout.

$(K_T)(iii)'$ implies that knowledge is closed under both competent valid deduction and correct (strong) induction. It respects the idea that both deduction and induction are ways of extending one's knowledge. $(K_T)(iii)'$ does not imply that strong inductive inferences can always extend one's knowledge, since strong induction may not be correct.

What about tracking? Nozick characterizes tracking in terms of both adherence and sensitivity. However, for reasons mentioned above, adherence is not a necessary condition for knowledge. I propose to take sensitivity as the sole condition for tracking— this is not an implausible move given that sensitivity is well supported by ordinary knowledge attributions (see Section 4.2).

I also accept the counterfactual formulation, since as I see it counterfactuals has the potential to shed new important light on our understanding of epistemic position (I will come back to this issue the next chapter). Hence, I formulate tracking as follows:

(T)   If R is S's sole reason for p, then R tracks the truth of p iff: If p had been false, S would have lacked R.

Hence, we can reformulate $(K_T)(ii)$ as follows:

$(K_T)(ii)^*$ if R is S sole reason for p, then if p had been false, S would have lacked R.

Two clarifications are called for. Firstly, notice that $(K_T)(ii)^*$ and (T) have made use of counterfactual conditionals. Many questions thus naturally arise. How should we understand such conditionals? What are their truth conditions? Are counterfactual conditionals useful in illustrating the notion of knowledge? What is the epistemic implication of characterizing epistemic position in terms of counterfactual conditionals?

Secondly, what is the nature of R? The term 'reason' is adopted from Dretske, who uses it roughly as a catchall term referring to "evidence, grounds, or reasons" for knowing that p. For Dretske, when S knows that p on the basis of R, the question "How does S know?" can "sensibly be asked and answered" (1971a, 4). But what are reasons of this sort? Dretske does not tell us much. At some point, he defines (stipulates) reason as "some experiential state of S" (1971a, 17). My usage of the term 'reason' is basically in line with Dretske's with certain additional specifications. In my usage, I take S's reason R (for p) as a *type* of experiential state, since I allow two agents to have the same reason R for p. Moreover, having R for p does not necessarily imply p; it is possible that S

possesses R for p and yet p is false.[55] However, in order to maintain the fluency of the discussion, I will often speak informally, and take one's reason to be one's memory or one's action. The readers should always keep in mind that the reason involved in knowledge, strictly, pertains to one's experiential states.

To sum up, I propose the following neo-tracking account of knowledge:

(K$_T$) S knows that p if and only if

(i) S believes truly that p on the basis of one's (sole) reason R for p. And

(ii) if R is S's sole reason for p, R tracks the truth of p in the sense that S would have lacked R if p were false, or

(iii) for some q, S knows that q and competently and adequately infers p from q, and thereby comes to believe that p, while retaining knowledge of q throughout.

Call it 'the tracking-based account (for knowledge)'. This account features the twofold requirement for S's standing on knowing dimensions with respect to p given that S believes truly that p, namely, the epistemic-position requirement—S has sufficient ground for p—and the doxastic-responsibility requirement—S believes that p on the basis of her ground for p.

S's epistemic position with respect to p is characterized by a recursive definition taking the property of tracking, i.e., the sensitivity of one's reason, as the base and the idea that new knowledge can be generated by old knowledge plus competent and adequate inference has figured into the inductive step. The tracking-based account implies that every pieces of knowledge will have to rely directly or indirectly on one's reason being sensitive. No knowledge without tracking (sensitivity).

---

[55] This idea captures fallibilism, which has been widely endorsed by philosophers. For fallibilism, see, e.g., Cohen (1988).

This can partly explain why skepticism is so powerful in challenging ordinary knowledge. The skeptic has made a strong case that we fails to track the truth of ordinary, empirical propositions, those that we normally thought we know. Since knowledge requires tracking as a base, if the skeptic has successfully reject the base case, he is in a position to deprive us of many other (if not all) empirical knowledge (see Chapter 7).

Before we get to the issue of skepticism, it is worth emphasizing that the tracking-based account, as presented here, does not imply that one's epistemic position may vary across different contexts while holding other things fixed. To see this, we will need to look closer into the idea of sensitivity. As I will argue in Chapter Seven, the analysis of the truth condition of counterfactual shows that whether one's reason for p is sensitive or not depends, among other things, on what epistemic inquiry one is engaging in. The latter is then affected by pragmatic considerations.

4.8 Taking Stock

In this chapter, I propose an account of one's standing on knowing dimensions. On this account, one's standing on knowing dimensions with respect to p should be characterized in terms of two factors, namely, the epistemic-position requirement—whether one's epistemic position with respect to p satisfies the standards for knowledge—and the doxastic-responsibility requirement—whether one's is epistemically responsible for believing that p.

My focus in the chapter has been the epistemic-position requirement. Opting for the tracking reliabilist account, I characterize one's epistemic position with respect to p in terms of how good one tracks the truth of p. Traditionally, the tracking account (of

knowledge) has been deeply influenced by Nozick's tracking account in three aspects: First, it is widely accepted that the relation of tracking pertaining to knowledge should be defined by sensitivity and adherence. Second, sensitivity and adherence, in turns, are normally characterized by the counterfactual-conditional and the belief formulation. Third, the tracking account has been regarded as incompatible with (Closure).

A closer scrutiny, however, shows that Nozick's account is problematic in the following respects: First, adherence is not necessary for knowledge. As shown in Section 4.3 (e.g., *Fan*), one's status of knowing can remain constant even if we shift one's status of adherence. Second, the belief formulation is a (at best) misleading formulation of sensitivity. Sensitivity, as far as this notion is related to knowledge, is not a property of belief, but rather a property of one's reason. In Section 4.4, I emphasize that the belief formulation ignores the distinction between the epistemic-position and the doxastic-responsibility requirement in such a way that it (implicitly) takes sensitivity to be a property of one's standing on knowing dimensions without realizing that sensitivity is only pertains only to the epistemic-position requirement. Third, I argue that proponents of the tracking account should not deny (Closure); we should respect the intuitive idea that deduction is a way of extending one's knowledge. More precisely, I argue that we should respect the more general intuitions that inference (deductive or inductive) is a way of extending one's knowledge.

Learning from Nozick's mistakes has been extremely fruitful since it indicates the correct path to construct the tracking account, or so I have argued. I propose a the neo-tracking account, i.e., the tracking-based account, which is constructed (partly) on the basis of correcting what goes wrong in Nozick's account. The tracking-based account

departs from Nozick's account in that, firstly, it takes sensitivity but not adherence as crucial to the property of tracking pertaining to knowledge. Secondly, it rejects the belief formulation and takes sensitivity as the property of one's reason. Thirdly, it respects the general intuition that deductive or inductive inferences are generally a good method to extend one's knowledge.

Historically, the tracking account has been generally considered as one of the failure attempts of characterizing knowledge. As I see it, this is an unfortunate development since the main reason people give up the tracking account is that they have seen serious problems for Nozick's account. In the next chapter, however, I will consider all prominent objections to the tracking account (they mainly focus on Nozick's account) and see whether they cause serious problems for the present account as well. I will argue that they do not. Going through them has shed a new light on the present account, showing that it is more promising than Nozick's account and that the notion of tracking is still promising in characterizing the notion of knowledge.

Chapter Five

Counterexamples Reconsidered

5.1 Introduction

In the last chapter, I propose that the tracking-based account (of knowledge) is the most promising tracking account of knowledge:[56]

(K_T) S knows that p if and only if

(i)   S believes truly that p on the basis of one's (sole) reason R for p in an appropriate way. And

(ii)   if R is S's sole reason for p, R tracks the truth of p in the sense that S would have lacked R if p were false, or

(iii) for some q, S knows that q and competently and adequately infers p from q, and thereby comes to believe that p, while retaining knowledge of q throughout.

One of the merits of this account is that it can handle *most* alleged counterexamples to the tracking account. Admittedly, we have not said much about sensitivity, except that it should be characterized in terms of counterfactuals. More will be said in the next chapter on this issue. In the meantime, with (K_T) at hand, we are already in a position to see how this neo-tracking account is immune from famous objections to the traditional (Nozickean) tracking account. As we will see, some of the objections can be met by incorporating (Closure) into our account, and others can be resolved by discarding the

---

[56] I personally think the tracking-based account is the most promising account of knowledge. But I will settle with this weaker claim for now.

belief formation—by distinguishing the epistemic-position requirement from the doxastic-responsibility requirement—and by taking sensitivity as a property of one's reason (as opposed to one's belief).

Going through all prominent counterexamples will prove to be fruitful. Not only does this process allow us to see that the tracking-based account has an edge over the Nozickean account, it also sheds a light on the tracking account in general.

5.1 Counterexample I: Cases of Induction

We have seen that sensitivity alone is incompatible with the idea that knowledge is closed under known entailment (see Section 4.5). Early proponents of the sensitivity condition such as Dretske and Nozick have thus argued that we should reject deductive closure in order to preserve sensitivity. Little has these authors, however, fully realized the serious tension between sensitivity and inductively learned knowledge.

Early commentators of Nozick's *Philosophical Explanations* have quick realized that some pieces of seemingly unproblematic inductively acquired knowledge are nonetheless insensitive in Nozick's sense. The line of objections has been repeatedly raised by Jonathan Vogel, for example (cf. Vogel 1987; 2000). Consider the following case:

> *Ice Cubes*. Imagine it's a hot day in August, say 95 degrees in the shade. Several hours ago, you left some ice cubes in a glass out in the direct sun, and since that time you've gone inside to get out of the heat. You think about the ice cubes, and it occurs to you that the ice you left outside must have melted by now. Despite the fact that you are not, at that moment, perceiving the shallow layer of water at the bottom of the glass, you know that the ice has melted. (Vogel 1987, 206)

Despite the fact that you know that the ice has melted, your belief that the ice has melted is not sensitive in Nozick's sense, since if the ice had not melted—say, your wife saw the

glass and put it in the frig—you would still have believed that the ice has melted. Nor is it sensitive in the sense of (K_T)(ii). If the ice had not melted, you would still have the same reason (i.e., your memory that you left the ice cubes outside and that the temperature was high out there, etc.) for *The ice has melted*. Hence, Nozick's account fails to predict what appears to be a piece of unproblematic inductively learned knowledge (the same, however, cannot be said of the tracking-based account since, unlike Nozick's account, the tracking-based account accepts a general closure principle).

This kind of counterexamples focuses on the inherent tension between sensitivity and inductive knowledge, especially when one knows something solely on the basis of inductive reasons. More precisely, such cases highlight the fact that if we are willing to admit that S knows that p and a certain inductive reason R is S's sole reason for p, we have to admit that S might still believe (or have the same inductive reason for) p even if p were false, since inductive arguments may be invalid. For instance, suppose that S's sole reason for believing that the sun will rise tomorrow is his observations in the past that the sun has risen. It seems plausible to claim that S knows that the sun will rise tomorrow. However, if the sun were to cease to rise tomorrow, S would still believe that it does (or have the same reason for *The sun will rise tomorrow*).[57]

In order to preserve inductive knowledge, Nozick has tried to argue that some inductively learned knowledge is indeed sensitive:

---

[57] Ernst Sosa has offered a similar counterexample:

> *Chute*. On my way to the elevator I release a trash bag down the chute from my high rise condo. Presumably, I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) were not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance. But none such could affect my predictive belief as I release it, so I would still predict that the bag would soon arrive in the basement. My belief seems not to be sensitive, therefore, but constitutes knowledge anyhow, and can correctly be said to do so. (Sosa 1999, 145–6)

> Do we know that the sun will rise tomorrow, that the earth will continue to rotate its axis during the next 24-hour? If the sun were not going to rise tomorrow, would we have seen that coming, would that alternation in the earth's rotation have been presaged in the facts available to us today and before? If so, then we do know that the sun will rise tomorrow; our belief that it will [track] the fact that it will, by being based on fact that would have been different otherwise. (Nozick 1981, 222)

The idea is that for some p such that S believes that p solely on some inductive basis R, if p were false, then the world would have been different in a certain way such that S would not believe that p (or have R) in the first place. Nozick's reply, however, still fails to satisfactorily account for all counterexamples from induction. Let me elaborate.

As Nozick also admits, some sort of inductive beliefs is characteristically *in*sensitive:

> But isn't it logically possible that everything was as it was until now, yet the earth will not continue to rotate tomorrow? Yes, there are such skeptical logical possibilities SK: the bread no longer nourishes us, the sun stops in the sky, an event of a certain sort no longer continues its usual effects. If they are elaborated suitably, so that everything we can detect up until now would have remained the same, then we don't know they do not hold. (Nozick 1981, 222)

Nozick wants to distinguish between inductively learned beliefs "whose falsity would have been reflected back and presaged" in the facts on the basis of which we form our belief and inductively learned beliefs whose falsity would have been not reflected back and presaged in such a way. Two problems immediately arise.

First, this distinction is, as I see it, problematic. Note that the falsity of *The sun will not rise tomorrow* will not necessarily have been reflected back and presaged in the facts on the basis of which we form our belief. If the sun were to be destroyed tonight, it would not rise tomorrow. Clearly, the sun being destroyed tomorrow does not need to be reflected back and presaged in the facts on the basis of which we form our belief. In general, I think *most* (if not all) counterfactual statements, under appropriate

specification, allow for both interpretations: the interpretation that its falsity would not have been presaged in the facts on the basis of which we form our belief and the interpretation that its falsity would have been not presaged in such a way. The reason is that this distinction in effect consists in the distinction between the *backward* and *forward* interpretations of counterfactuals, and that, arguably, the same counterfactual statement can be interpreted as backward or forward, depending on the contexts (see Chapter Six).[58]

Nozick's incorporation of inductive knowledge has committed to a very serious problem in conflating the forward and backward interpretations of counterfactuals—at least, he seems to allow both interpretations for the sensitivity condition. This causes a serious problem since the backward and forward interpretations of counterfactuals specify different relationships between the subject (or her reason) and the believed proposition p. And as far as I can tell, the sensitivity involved in knowledge only deal with the forward interpretation, but not the backward one (see Chapter Six).

Even if we put the problem of counterfactuals aside, Nozick's reply still faces another serious problem. For the sake of discussion, let us suppose that there is a standard (or preferable) way of interpreting counterfactuals. More precisely, let us assume that Nozick is right that, on the standard interpretation, S's belief that the sun will not rise tomorrow is *sensitive*. But this does not save Nozick's account from counterexamples from induction, since his account still renders many intuitively knowable inductive beliefs unknowable. As Vogel points out, according to Nozick's distinction, we would

---

[58] The terms 'backward' and 'forward' counterfactual conditionals are borrowed from Jonathan Bennett:
> A forward counterfactual conditional is one whose consequent is about a later time than any the antecedent is about… The consequent of a backward conditional is about a time earlier than any that its antecedent is about. (Bennett 1984, 57)

not know (via induction) that the earth will not *suddenly* stop spinning (1987, 205), since according to the standard interpretation, the counterfactual *If the earth were to suddenly stop spinning, we would still believe that it* seems true: On the standard interpretation, if the earth were to suddenly stop spinning, it might be that some drastic damages have occurred suddenly to the earth, and since our belief that the earth is spinning is based on past observation (or testimony, etc.), we would believe that the earth does not suddenly stop spinning (or we would still have the same reason for *The earth does not suddenly stop spinning*). But this conclusion seems incorrect. Intuitively, it seems that we not only know (via induction) that the earth will not stop spinning, we also know (via induction) that it will not suddenly stop spinning (or that if we know (via induction) that the earth will not stop spinning, we also know (via induction) that it will not suddenly stop spinning.

While inductive knowledge is a serious problem for Nozick's account of tracking, it is not a problem for the tracking-based account, since this account takes knowledge to be closed under adequate inferences. Suppose that S believes that the sun will rise tomorrow solely on the basis of induction. Does S know that the sun will rise tomorrow according to the current account? Well, normally, the answer is "Yes". According to the present account, S knows that the sun will rise tomorrow, say, if (i) S's true belief that scientists around the world say that the sun will rise tomorrow is S's sole reason for *The sun will rise tomorrow*, and if (ii) S competently and adequately infers *The sun will rise tomorrow* from *Scientists around the world say that the sun will rise tomorrow*. Under normal circumstances, (i) and (ii) can be easily satisfied. As a result, S knows that the sun will rise tomorrow in normal circumstances.

5.2 Counterexample II: Not Falsely Believing

There is also a very influential line of objection to sensitivity which claims that the sensitivity requirement cannot account for one's knowledge that one does not falsely believe that p. Here is Sosa's example:

> *False Belief.* Suppose first we have two propositions as follows: (a) that p, and (b) that I do not believe incorrectly (falsely) that p. Surely no-one minimally rational and attentive who believes both of these will normally know either without knowing the other. Yet even in cases where one's belief of (a) is sensitive, one's belief of (b) could never be sensitive. After all, even if (b) were false, one would still believe it anyhow. Still it is quite implausible that the assertion that I know (b) could never be true, not even in many situations where the assertion that I know (a) *would* be true. (Sosa 1999, 145)[59]

Sosa's objection, at its core, relies on two intuitively plausible claims:

(A)   Generally, the second-order belief that one does not falsely believe that p is insensitive.

(B)   Generally, one can know that one does not falsely believe that p on the basis of one's knowledge that p.

(A) relies on an unproblematic assumption that in general if one falsely believes that p, one would still think that one does not falsely believe that p—it is paradoxical to assert Moore's paradox, e.g., "I believe that p but p is false". (B) is unproblematic since it requires a moderate claim that knowing p put one in a position to believe that one does

---

[59] Vogel also gives a similar objection:

> *False Belief*.* [Nozick's account] requires that, if you know X, then you would not believe X, if X were false. As things actually are, you believe that your belief that Omar has new shoes is not false. What if it were? If somehow your belief that Omar has a new pair of shoes were false, you would still believe that your belief was true, not false. The alternative is hard to fathom. It is difficult to conceive of your not believing that something you believe is true, whenever the matter happens to cross your mind. So, if your belief that Omar has new shoes were false, you would still believe that your belief was true, not false. You thereby fail to satisfy [sensitivity]. According to [Nozick's account], then, you do not know that your belief that Omar has new shoes is true, not false. (Vogel 2000, 610)

not falsely believe that p—although not all pieces of knowledge guarantee such transition, but many do.

(A) and (B) imply the denial of Nozick's account of knowledge; if (A) and (B) are correct, then there are some second-order knowledge which is insensitive: it is possible that S knows that S does not falsely believe that p and S's belief that S does not falsely believe that p is insensitive.

Cases like *False Belief* also pose a special problem for the pluralistic solution of the skeptical paradox. Recall that the pluralistic solution claims that the lack of knowledge of not-sh implies that the lack of knowledge of op. Now, let 'sh$_1$' and 'op$_1$' stand for *One falsely believe that one has hands* and *One has hands* respectively. On the pluralistic solution, if one does not know that not-sh$_1$, one does not know that op$_1$. But this is absurd, since sh$_1$ is not a well-constructed skeptical hypothesis in the first place— while one may try to persuade (even successfully) others in admitting that they do not know that they have hands by telling them the story that they may be a brain in a vat being fed by experiencing that they have hands, it is unlikely that one will achieve the same success by mentioning that they may believe falsely that they have hands. This shows that insensitivity is not enough to explain why we lack knowledge of not-sh; something more is called for.

DeRose tries to reply to the second objection by claiming that the genuine skeptical hypotheses—hypotheses that generate the skeptical puzzle—are different from the specious ones in that the former, but not the latter, also explain how one comes to falsely believe that op (DeRose 1995, 23). For instance, the skeptical hypothesis that one is a brain in a vat stimulated to experience that one has hands explains how one comes to

falsely believe that one has hands in the way that the skeptical hypothesis that one falsely believe that one has hands does not, and, as a result, only the former generates the skeptical puzzle.

DeRose's reply relies on a vague notion of explanation, which itself may bring more problem than solution. But even if we set the issue of explanation aside, DeRose's suggestion that explanatory power, accompanied with the sensitivity requirement, somehow distinguishes genuine skeptical hypotheses from specious ones is still vulnerable to counterexamples. Timothy Williamson offers the following counterexample:

> *BIV*.* For example, let α be a case in which I am climbing a mountain and appear to myself to be climbing a mountain; things appear to me as they are. In case β, I am a brain in a vat, but things appear to me generally just as they do in α. If I were in β, I should believe that I was not in β. If p is the proposition that I am not in β, then I do not sensitively believe p. Nevertheless, in my present case, I know that I am not in β, because I am in neither α nor β; things do not appear to me at all as they would in β. I appear to myself to be sitting in front of a computer screen in my office. I do not appear to myself to be climbing a mountain as I would in β. No matter what my situation, I cannot sensitively believe p. … Moreover, my false beliefs when I am in β is just explicable as in other skeptical scenarios, while the example does not involve false beliefs when I am in either α or my actual case. (Williamson 2000, 158)

The crux of this kind of objection is that there are some skeptical hypotheses which are insensitive and have the explanatory power of explaining how we come to falsely believe them, and yet we normally regard ourselves as knowing their falsity. Therefore, the Nozick's sensitivity account, even if complemented by DeRose's suggestion of explanatory power, still falls short of being a correct account of knowledge.

Surely, there are replies, but it is not clear how successful they are. To its merits, the tracking-based account can offer a satisfactory reply to this line of objection. Consider, firstly, *False Belief.* On my account, one can generally know that one does not

falsely believe that p on the basis of one's knowledge that p only if either one's sole reason R for *One knows that p* is sensitive or for some q, one knows that q and competently and adequately infer *One knows that p* from q. This makes sense once we consider actual cases. Suppose that I know that New York is not the capital of USA. Also, suppose that my (sole) reason for my belief that I knowledge that New York is not the capital of USA is that I remember I have read this piece of information in a textbook. Now, can I know that I do not falsely believe that New York is not the capital of USA based on my knowledge that New York is not the capital of USA? In many cases, I can. There are many ways I might have had falsely believed that New York is not the capital of New York. Suppose that in the counterfactual situation I falsely believed that New York is not the capital of USA because New York were indeed the capital of New York. In this case, it seems that my reason for *I do not falsely believe that New York is the capital of USA*—my knowledge that New York is not the capital of USA—is sensitive: In this case, I would have lacked the same reason due to the fact that I would have read something different from the textbook.

Another way for one to know that one does not falsely believe that p is for one to competently and adequately infer p from some known q. Suppose that I know that the capital of USA is Washington, DC., and I know that I know. Suppose also that I am capable of inferring *I do not falsely believe that New York is not the capital of USA* from *I know that the capital of USA is Washington, DC.* According to the tracking-based account, I thus know that I do not falsely believe that New York is not the capital of USA. Hence, unlike the traditional sensitivity account, the present account allows one to have second-order knowledge such as one knows that one does not falsely believe that p.

142

Secondly, consider *BIV**. *BIV** and the like does not constitute counterexamples to the present account since *BIV** is totally compatible with S's knowledge that S is not a brain in a vat that experiences climbing a mountain. More precisely, on the tracking-based account, as long as there is some q such that S knows that q, and competently and adequately infers *S is not a brain in a vat that experiences climbing a mountain* from q, S is in a position to know that S is not a brain in a vat that experiences climbing a mountain. This should be quite intuitively plausible. Suppose that S knows that it appears to S that S is sitting in front of a computer screen in S's office, and S competently and adequately infers *S is not a brain in a vat that experiences climbing a mountain* from *It appears to S that S is sitting in front of a computer screen in S's office*, don't we also want to say that S knows that S is not a brain in a vat that experiences climbing a mountain?

5.3 Counterexample III: Martin's Case

The denial of the belief formulation consists in the distinction between doxastic-responsibility and truth-relevant dimensions of knowledge. As it turns out, keeping this distinction helps us to avoid some prominent (obvious?) counterexamples to the tracking account.

Raymond Martin has offered a famous counterexample to Dretske's account of knowledge, which, like the tracking-based account, takes sensitivity to be a property of one's reason for p:

> *Horse Racing.* Suppose that S buys a ticket at a local horse track. The ticket is such that if S picks the winner in the first race or the winner in the second race, or both, S wins; otherwise, S losses. S picks Gumshoe to win the first race and Tagalong to win the second. Before the first race begins S is called away from the

143

track and does not return until the end of the second race. While S is away Gumshoe wins the first race … Tagalong finishes last in the second … On returning to the track and still not knowing the results of either race S goes to the cashier's window and presents his ticket. Without speaking a word, the cashier gives twenty dollars in exchange for S's ticket. (Martin 1975, 215–6)

Suppose that S believes that Gumshoe wins the race solely on the basis of receiving twenty dollars from the cashier. It seems obvious that S does *not* know that Gumshoe wins the race, but S's reason, i.e. that S receives twenty dollars from the cashier, is *sensitive*: If Gumshoe were to lose, S would not receive twenty dollars from the cashier.

Admittedly, cases like *Horse Racing* are intuitively persuasive. They have forced Dretske to change his mind later (cf. Adams and Clarke 2005, 208), and struck some (many?) as "completely decisive" (Fogelin 1994, 69). Fred Adams and Murray Clarke have tried to defend Dretske's account of sensitivity against Martin's counterexample. According to them, the sensitivity of one's reason holds only "relative to a specific set of circumstances C" (Adams and Clarke 2005, 210). In *Horse Racing*, S's reason R (i.e., a payout) for p (i.e., *Gumshoe wins*) is sensitive "when, given C, it is not possible to have R and not p" (2005, 210). They argue that Martin's counterexample does not constitute a genuine counterexample to Dretske's account since the circumstance relative to the sensitivity of R has been implicitly shifted:

> In Martin example, [S] buys one ticket prior to the race giving him Gumshoe in the first and Tagalong in the second. The circumstances C become fixed at the time of purchase. So a payout by a teller (R, with no more information) says only that either Gumshoe won or Tagalong won. We claim that Martin changes the circumstance from C to C′ after the first race. Then Martin claims, given that Tagalong lost (C′), a payout (R) says that Gumshoe won (that p) because given C′ it is not possible to have R and not p (not possible to receive payment unless Gumshoe won). But R's status as a [sensitive] reason is tied to circumstance C (not C′) because [S] purchased the ticket in C (not C′). R is not a [sensitive] reason relative to C, and that is why [S] does not know that Gumshoe won on the basis of R. The most R says, relative to C, is that either Gumshoe won or Tagalong won (but not which). (Adams and Clarke 2005, 210)

Adams and Murray's reply, however, is unsuccessful. First, there is a sense of ad hocery in this reply: In *Horse Racing*, S has received the payout (R) at C′ (after the first race), not at C (before the race). There is no principled reason for taking R to be related to C rather than to C′ (the main motivation to this move, it seems, is to get Dretske's account off the hook). Second, and more importantly, such an interpretation of the sensitivity of reason simply does not hold up to close scrutiny. Suppose that we modify *Horse Racing* such that on returning to the track to receive his payout, S has already known that Tagalong lost. By receiving the payout, it seems that S thus *knows* that Gumshoe wins. But if Adam and Murray were right—if S's R (i.e., the payout) is relative to C (before the race)—then it should be the case that S's reason for R is still *not* sensitive in C′, and as a result, S should not know that Gumshoe wins. But this is absurd.

Cases like *Horse Racing* do not pose a serious problem to the tracking-based account, or so I will argue. The key, as I see it, is to notice that cases like *Horse Racing* will pose a serious problem only if we take one's standing on knowing dimensions with respect to p to consist solely in the sensitivity of one's reasons for p. But as I argue above, the sensitivity of one's reason for p itself is *insufficient* for rendering one's true belief knowledge. Let me elaborate.

Recall that one's standing on knowing dimensions with respect to p is a placeholder for whatever condition that turns one's true belief into knowledge. I have argued that one's standing on knowing dimensions with respect to p consists in two general conditions: the epistemic-position and doxastic-responsibility requirement (see Section 4.4). Satisfying the epistemic-position requirement with respect to p puts one into a position to know p, but does *not* guarantee that one knows. One can be in a position to

know that p without knowing that p. Now, on the tracking-based account, the epistemic-position requirement consists in whether one's sole reason for p is sensitive. The former is satisfied if one's reason is indeed sensitive. That is, on the tracking-based account, if one's reason for p is sensitive, one is only in a position to know that p; one may not know that p even if one's sole reason for p is sensitive.

How will this view account for cases like *Horse Racing*? On the tracking-based account, S actually satisfies the epistemic-position requirement. In other words, S is in a position to know that Gumshoe wins the race. This is not implausible, since in this case, S's reason (receiving the payout) is in fact a good indicator of the truth of *Gumshoe wins*—if Gumshoe were to lose, S would not have received the payout. But the tracking-based account does not imply the implausible result that S does know that Gumshoe wins. Note that the main reason that prompts us to think that S does not know that Gumshoe wins in *Horse Racing* is the specification that S still does not know the result of either race when he presents his ticket to the cashier's window. It is obvious that S does not know since he has no information (presumably) to believe that receiving the payout indicates that Gumshoe wins. If S forms a belief that Gumshoe wins solely on the basis of him receiving the payout, S's belief will not satisfy the doxastic-requirement, since S's belief is not formed *in an appropriate way*.

Viewed in this light, we can see the parallel between *Horse Racing* and BonJour's *Norman* mentioned above. In *Norman*, Norman is in a position to know that the president is in New York City. But he does not actually know. The moral we learn from cases like this is that satisfying the epistemic-position requirement is not enough for knowledge; the satisfaction of the doxastic-responsibility requirement is also called for. In *Norman*,

Norman has no reason to believe that he is clairvoyant (and perhaps has plenty of reason to deny his clairvoyant capacity). Hence, Norman is epistemically irresponsible to form the belief that the President is in New York City solely on the basis of his clairvoyance. Similarly, in *Horse Racing*, S has no reason to believe that Tagalong does not win, and (presumably) S has reason to believe that he will receive twenty dollars from the cashier if either Gumshoe wins or Tagalong does (or both). Hence, it would be epistemically irresponsible for S to form the belief that Gumshoe wins solely on the basis of receiving the payout from the cashier. In both *Norman* and *Horse Racing*, the subject is in a position to know without actually knowing, since the doxastic-responsibility requirement is not satisfied.

The distinction between the epistemic-position and doxastic-responsibility requirement also helps us see clearly why the following case is not a genuine counterexample to the tracking-based account:

> *Placebo.* Consider a medical experimenter, testing the efficacy of a new drug on a certain disease. Suppose he neglects to apply proper experimental controls. To take an extreme case, suppose he neglects to give a placebo to a control group… Suppose, however, that in fact placebos are completely (or nearly completely) ineffective against this particular disease, although neither the experimenter nor anyone else in the medical community has any reason to suppose that this is so. Then, if the experimenter concludes that he has tested a drug that is chemically effective against the disease, his belief satisfies Nozick's conditions for knowledge. In particular, despite the fact that any medical scientist would say that the experimenter has not ruled out a relevant alternative, Nozick's [condition of sensitivity] is satisfied. For if the experimenter had tested a drug that was not chemically effective against the disease, little improvement in the patient would have resulted, since the disease is impervious to placebo effects. (Kripke 2011, 170–1)

Kripke's own judgment is that the experimenter does not know that he has tested a drug that is chemically effective against the disease since his belief is "the fruits of much sloppy experiment and research" (2011, 170). But since the experimenter's belief

satisfies Nozick's conditions for knowledge, he would be considered as "knows" by Nozick's account. Kripke then accuses Nozick's account as being "inimical to the very idea of proper experiment controls" (2011, 170).

As I see it, *Placebo* is structurally similar to *Horse Racing* and *Norman*: the subject is in a position to know that p since the subject's sole reason for p is sensitive, and yet the subject does not know due to the failure of satisfying the doxastic-responsibility requirement. In *Placebo*, the test is in fact a good indicator of the truth of *The drug is chemically effective against the disease*—it is the assumption of this case that the placebo effect does not obtain. However, the experimenter has no reason to suppose that the placebo effect does not take place (if the experimenter had good reason to believe that if he used the controls, the placebos would have been negative, we will be happy to grant that the experimenter in fact *knows*). As a result, if the experimenter is to form the belief that the drug is chemically effective against the disease solely on the basis of the result of the test, his belief will be epistemically irresponsible. No wonder we regard him as not knowing.

5.4 Counterexamples IV: Nozick's *Grandmother* and Goldman's *Oscar*

Nozick has considered a counterexample to the tracking account:

> *Grandmother.* A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. (Nozick 1981, 179)

Intuitively, the grandmother knows that the grandson is well. But her belief that the grandson is well is not sensitive in Nozick's sense: If he were sick, she would still believe that he is well. In order to overcome cases like *Grandmother*, Nozick suggests that we

should relativize the notion of sensitivity to method such that we should take one's belief that p as sensitive just in case if p were false, one would still have believed that p via method M (provided that one believes that p via M in the actual case).

Nozick's relativity of method gives rise to its own problem.[60] Fortunately, such a move is not even required for the tracking-based account. In *Grandmother*, presumably the grandmother's sole reason for *The grandson is well* is her perception (or visual experience) that the grandson is alive and kicking. However, if the grandson was not well, she would not have such perception (or visual experience). Hence, the grandmother's reason for *The grandson is well* is in fact sensitive; this explains why we intuitively think that she knows that the grandson is well.

Cases like *Grandmother*, in general, show that there is a logical gap between believing that p and having R as the sole reason for p. Under appropriate stipulation, one can believe p in different cases while one's sole reason for p is different. For instance, the reason for the grandmother's belief that the grandson is well in the actual case is her observation, while her reason for the same belief in the counterfactual case is others' testimony. The gap between believing p and having reason for p has prompted Nozick to relativize tracking to the method of acquiring a proposition. But this move is totally *ad hoc* and unnecessary from the perspective of the tracking-based account.

The way of the tracking-based account accounts for *Grandmother* is applicable to the famous counterexamples to sensitivity proposed by Alvin Goldman. In his (1976), Goldman offers a very famous counterexample to Nozick's account of sensitivity:

> *Oscar.* Suppose Oscar is standing in an open field containing Dack the
> dachshund, Oscar sees Dack and (noninferentially) forms a belief in (P):

---

[60] See, for example, Vogel (2012).

(P)   The object over there is a dog.

Now suppose that (Q):

(Q)   The object over there is a wolf.

> … wolves are frequenters of this field... Further suppose that Oscar has a tendency to mistake wolves for dogs (he confuses them with malamutes, or German shepherds). Then if the object Oscar saw were Wiley the wolf, rather than Dack the dachshund, Oscar would (still) believe (P). This means that [Nozick's sensitivity requirement] is violated. But surely it is wrong to deny… that Oscar *knows* (P) to be true. The mere fact that he would erroneously take a wolf to be a dog hardly shows that he doesn't know a *dachshund* to be a dog! (Goldman 1976, 779)

In *Oscar*, Oscar's belief that the object over there is a dog is insensitive: If Oscar did not see Dack the dachshund, he would have seen Wiley the wolf, and in that case, he would still believe that the object over there is a dog.[61]

The tracking-based account handles cases like *Oscar* nicely. In *Oscar*, Oscar's sole reason for the object over there is a dog is his perception of Dack. If he were to perceive Wiley, he would lack the perception of Dack. As a result, the tracking-based account is consistent with the judgment that Oscar knows that the object over there is a dog.

---

[61] Cases like *Oscar* also put pressure on Nozick's refined version of tracking which relativizes tracking with method. Regardless of how one individuates the method in play, it seems clear that one has to grant that grandmother uses the same method to form her belief about the grandson's health status if she observe the grandson behavior in the actual and counterfactual cases. More precisely, suppose that in *Grandmother*, if the grandson were sick, the grandmother would have seen him sick, and thus come to believe that he is sick. Nozick's account should conclude that the grandmother knows that the grandson is well since her belief that the grandson is well is sensitive: In the actual world, she believes that the grandson is well via M1 and if the grandson was not well, she would not believe that he is well via M1. If this was the right interpretation of *Grandmother*, the same should be said about *Oscar*. That is, if Oscar is to believe that the object over there is a dog via M2 in the actual case, he would have to believe that the object over there is a dog via M2 in the counterfactual case. But this last claim implies that Oscar does not know that the object over there is a dog, since if the object there is not a dog, it would be a wolf, and Oscar would still have believed that the object over there is a dog via the same method (cf. Williamson 2000, 157).

Oscar is different from *Grandmother* in that *Oscar* seems to use to same method. Yet, we do not want to say that just because Oscar cannot distinguish malamutes or German shepherds from wolves it does not mean that he cannot distinguish dachshunds from wolves.

Before we move on to the next type of counterexamples, it is worth mentioning that there is a convenient (common) way to construct "counterexample" to the sensitivity condition. The idea is that we can just stipulate the scenario such that if p did not obtain, then q would obtain, and the subject was incapable of distinguishing p from q. For instance, one can modify *Grandmother* such that if the grandson were sick or dead, others would have another person disguised as the grandson and the grandmother could not tell between the real grandson and the fake. This maneuver is convenient and has been exploited by many. I will call such cases "counterexamples from pseudo skeptical hypotheses". For many of the following cases, one can make a twist and turn them into counterexamples from pseudo skeptical hypotheses using the guideline mentioned here. For simplicity's sake, I will not consider such a twist when talking about other cases. I discuss counterexamples from pseudo skeptical hypotheses in Section 5.6.

5.5 Counterexample V: Kripke's *Red Barn*

Viewed in this light, we can also account for Kripke's (another) famous counterexample to sensitivity. To appreciate Kripke's objection, we need to first recognize how Nozick's account accounts for Goldman's famous barn case:

> *Barn*. [Suppose that Henry is driving in the countryside which] is full of papier-mâché facsimiles of barns. These facsimiles look from the road exactly like barns, but are really just facades, without backs walls or interiors, quite incapable of being used as barns. They are so cleverly constructed that travelers invariably mistake them for barns. Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the object on that site were a facsimile, Henry would mistake it for a barn. (Goldman 1976, 773)

"Given this new information," Goldman concludes, "we would be strongly inclined to withdraw the claim that Henry *knows* the object is a barn."

For the sake of argument, let us suppose that Goldman is right. It has been widely agreed that Nozick's account predicts Henry's lack of knowledge. In *Barn*, Henry is in an area full of fake barns such that if he did not look at a genuine barn, he would have seen a fake barn. In that case, his belief that the object is a barn is not sensitive: If the object was not a barn, Henry would still have believed that it is a barn. Hence, Henry does not know that the object is a barn.

Kripke has powerfully pointed out that the aforementioned explanation of *Barn* is problematic. He modifies *Barn* as follows:

> *Red Barn.* [N]ow let us suppose that the barn is red. Suppose further that any counterfeit erected in its place would have been green. (We can suppose, if we wish, that for some chemical reason the cardboard in the counterfeit barns cannot be painted red. Alternatively, those who erected counterfeits barns definitely preferred green ones, or even definitely preferred a green one in this particular location.) Now consider Henry's true belief … that there is a (genuine) red barn in the field. Now the [sensitivity] condition is satisfied. If there had not been a red barn in the field, then there would have been a green counterfeit, and Henry would have not believed that there was a red barn in the field. (Kripke 2011, 186)

Now, it seems that Henry's belief that there is a red barn in the field satisfies all Nozick's conditions for knowledge, and thus, on Nozick's account, Henry knows that there is a red barn in the field. But, the problem is, how can Henry knows that there is a red barn in the field but not that there is a barn in the field?

Ever since it is proposed, *Red Barn* has been widely considered as a very serious problem for the sensitivity account; one of those insightful counterexamples that proclaims the death of the sensitivity account. What would the tracking-based account say about this objection?

On the tracking-based account, Henry is in a position to know that there is a barn in the field as well as there is a red barn in the field. In *Barn*, Henry's sole reason for

There is a barn in the field, presumably, is his perception of the barn in front of him. I am inclined to think that Henry's reason is sensitive: If there were no barn in front of him, he would have not the same reason for *There is a barn in the field*—it seems very plausible that if the barn Henry is facing did not exist, he would be facing nothing, and as a result, lack the perception of a barn.[62] Now consider *Red Barn*. On the tracking-based account, Henry is also in a position to know that there is a red barn in the field since it is obvious that Henry's sole reason for *There is a red barn in the field*, i.e., Henry's perception of the red barn, is sensitive: If the red barn did not exist, Henry would not have had the same reason.

The tracking-based account has an edge over Nozick's account when facing Kripke's challenge. First, Nozick's account requires a different accounts and verdicts with respect to *Barn* and *Red Barn* of which many have felt (rightly, I think) uncomfortable. The tracking-based account, by contrast, offers a uniform account and verdict to both cases. Henry's sole reason for *There is a barn in the field* (*There is a red barn in the field*) is sensitive; Henry is in a position to know that there is a barn in the field (that there is a red barn in the field). More importantly, such a uniform account and verdict to *Barn* and *Red Barn* does not necessary requires a deviation from the widely adopted verdicts of such cases. For instance, the tracking-based account can agree with Goldman that Henry does not know that there is a barn in the field. The idea is that while Henry satisfies the epistemic-position requirement of knowledge—Henry's sole reason for *There is a barn in the field* is sensitive—Henry may not satisfy the doxastic-

---

[62] Of course, some might try to reply that we can stipulate *Barn* such that if the barn Henry is facing did not exist, a fake barn would have been built on the same location, and Henry would have thought that he was seeing a barn. This is precisely the kind of move that generates counterexamples from pseudo skeptical hypotheses. I will deal with such kind of cases in Section 5.6.

responsibility requirement, since Henry has failed to realize that he is in a field full of fake barns.

Second, while Nozick's account is bounded to give a *definite* answer as to whether Henry knows in *Barn* and *Red Barn*, the tracking-based account leave more room for different verdicts in such cases. In *Barn* (*Red Barn*), Henry is in a position to know that there is a (red) barn in the field. But being in a position to know does not guarantee knowledge. Hence, strictly, the tracking-based account is consistent with both the verdicts that Henry knows that there is a (red) barn in the field and that Henry does not know that there is a (red) barn in the field.

In my opinion, such flexibility is more than welcome in cases like *Barn* and *Red Barn*, since I do not have robust intuition regarding them in the first place. More precisely, contrary to what philosophers tend to take for granted, I do not have a clear intuition that Henry does not know that there is a barn in the field. The underlying thought for the common verdict of *Barn* seems to be that knowledge is incompatible with epistemic luck: Although Henry believes truly that there is a barn, his belief is true accidentally—if he were to look at a fake barn (which is a more likely scenario according to the story), he would still believe that there is a barn in the field.

However, knowledge is not incompatible with all sort of epistemic luck. Some epistemic luck seems to be compatible with knowledge. For instance, imagine a world that everybody has 20 percent chance of being a brain in a vat. I am lucky to live in the real world. But my luck does not seem to debar me from knowing that I have hands. Some luck is epistemically vicious, and undermines knowledge. If I make a wild guess that the President is in China at the moment and it happens to be right, my true belief fails

to constitute knowledge since its truth is based on (epistemically vicious) luck. So, the slogan that knowledge is anti-luck true belief should be changed as that knowledge is anti-vicious-luck true belief.

While epistemic luck does play a crucial role in Henry's acquisition of true belief in *Barn*, it is not obvious to me that the luck involved is necessarily epistemically vicious. Consider a similar case discussed by Fogelin:

> *Thermometer*. Suppose that S took the thermometer from a box of thermometers that were being discarded because the manufacturer had written saying they were defective. In fact, all the thermometers were defective except for the one chosen by S; it was in perfect working order. [Suppose S has placed the thermometer in his child's mouth and the thermometer reads, correctly, 98.6°.] In this case it is true that the thermometer would not have registered 98.6° unless the child's temperature were normal... (Fogelin 1994, 67)

Suppose that S believes that the child's temperature were normal based on his reading of the thermometer. S's belief is true because he is lucky—it happens that S has picked an effective thermometer. "As the story is told," Fogelin concludes, "it does seem clear that S does not know that his child's temperature is normal" (1994, 68). *Thermometer* is a lot like *Barn* in that the subject is supposed to be lucky in acquiring a true belief—in *Barn*, Henry is lucky to look at a real barn among all fakes, while in *Thermometer*, S is lucky to pick an effective thermometer out of all defective ones. Following the common view, Fogelin takes it for granted that S's true belief does not constitute knowledge in *Thermometer*.

It strikes me as dogmatic to conclude that S does not know that his child's temperature is normal. To me, it makes sense to say that S knows that his child's temperature is normal in *Thermometer* inasmuch the same way as it does to say that Henry knows that there is a barn in the field. However, it seems (to me at least) Fogelin's

verdict of *Thermometer* is somehow more plausible than Goldman's verdict of *Barn* since, in *Thermometer*, it has been explicitly mentioned that the box of thermometers has been "discarded" and the manufacturer had written saying that they are "defective". To appreciate how much weight this specification accounts for our intuitions. Consider a case where this stipulation is removed:

> *Thermometer**. Everything is like the previous one except that the box of thermometers is not discarded, and that the manufacturer has very low rate of producing defective thermometers. Out of sheer luck, the box of thermometers happens to contain all but one defective thermometer. Suppose that S has placed the thermometer in his child's mouth and the thermometer reads, correctly, 98.6°.

To me, although it still make sense to say that S does not know that his child's temperature is normal in *Thermometer**, it seems that our intuitions regarding S's lack of knowledge in *Thermometer** are significantly weaker than our intuitions regarding S's lack of knowledge in *Thermometer*.

It should also be noted that *Thermometer** is more similar to *Barn* than *Thermometer* is, since *Barn* also lacks the specification that the countryside is full of fake barns. By the same token, we can modify *Barn* so that it is more similar to *Thermometer*:

> *Barn'*. Just like *Barn*, Henry is driving in the countryside full of fake barns. Henry has been lucky to look at a rare genuine barn in the area and come to believe that there is a barn in the field. Unlike *Barn*, at the entrance of the countryside there is an obvious sign that says that the field consists of 99 percent of fake barns and 1 percent of genuine barns. When driving in the countryside, Henry is absent-minded and does not see the sign.

Barn' is more similar to *Thermometer* than *Barn* is since, according to the former (but not the latter), Henry has failed to notice an obvious "sign" indicating the existence of a large amount of fake barns. I suspect that our intuitions regarding Henry's lack of knowledge in *Barn'* is stronger than our intuitions regarding Henry's lack of knowledge in *Barn*,

parallel with our intuitions regarding S's knowing status in *Thermometer* and in *Thermometer**.

In sum, *Barn* is analogous to *Thermometer**, while *Barn'* is analogous to *Thermometer*. Although our intuitively think that Henry (S) does not possess knowledge in *Barn* (*Thermometer**) and *Barn'* (*Thermometer*), the strength of our intuitions vary with such cases. In particular, we are much willing to deprive Henry (S) of knowledge in *Barn'* (*Thermometer*) than we are with respect to *Barn* (*Thermometer**). This phenomenon, if real, needs clarification. The traditional Nozickean account of knowledge cannot account for this phenomenon since the sensitivity of Henry's (S's) belief does not vary with such cases.

The current account is in a position to explain the aforementioned phenomenon. On this view, in both *Barn'* and *Thermometer*, the subject has failed to recognize some important information related to their true belief—in *Barn'*, Henry fails to recognize the sign indicating the existence of fake barns, while, in *Thermometer*, S fails to recognize the sign indicating that the box is full of defective thermometers. More importantly, the subject's failure of recognition seems to be in tension with the satisfaction of the doxastic-responsibility requirement—it seems that the subject is somehow epistemic irresponsible in forming their respective beliefs. I propose that the seemingly failure to the satisfaction of the doxastic-responsibility requirement is responsible for our stronger intuitions that the subject does not possess knowledge in *Barn'* and *Thermometer*—we tend to think that Henry (S) does not know that there is a barn in the field (that the child's temperature is normal) in *Barn'* (*Thermometer*) since the story strongly suggests that Henry (S) has been epistemically irresponsible in forming his belief.

Now, I suggest that the reason that people are inclined to think that Henry (S) lacks knowledge in *Barn* (*Thermometer**) (if at all) is that they judge that Henry does not satisfy the doxastic-responsibility requirement. When the story is modified to emphasize (or indicate) Henry's (S's) epistemic irresponsibility such as *Barn'* (*Thermometer*), people are more willing to deprive Henry (S) of knowledge.

This explanation of the underlying mechanism responsible for our intuitions regarding *Barn* (*Barn'*) and *Thermometer* (*Thermometer**) is vindicated by considering cases similar except that the subject has been emphasized as not being epistemically irresponsible. Tamar Gendler and John Hawthorne have an interesting case showing exactly this point:

> *Diamond*. Recently introduced in our widely-publicized "Your friends Will Never Know You're Wearing a Diamond Ring" campaign, our costume jewelry collection offers you a way of preventing others from knowing that you are sporting some sort of valuable doo-dah. Just send us a photograph of your genuine gem, and we'll do the rest!
>
> Our diamond Ring Kit provides you with six phony diamond rings that look identical to your genuine rock. Slip them surreptitiously into your pocket, and whenever someone sees your rings, there will be lots of fakes in the area. Result? Even when their eyes chance upon it, your friends will not know that you're wearing a diamond ring! (Gendler and Hawthorne 2005, 334)

Well, if John sees Sarah wearing a diamond ring, while in fact she has 6 fake diamonds in her purse, does he know that Sarah is wearing a diamond ring? It seems that John does know. This result cannot be explained by Nozick's account since, according to this account, the subject's status of knowledge in *Diamond* is essentially similar to the one in *Barn*—if you deprive Henry of knowledge in *Barn*, you should do the same to John in *Diamond*.

The crucial difference between *Diamond* and *Barn* is that the former (but not the latter) explicitly emphasizes that John is not epistemically irresponsible in forming the belief that Sarah is wearing a diamond ring—it is absurd to require everyone to first make sure that the non-existence of fake diamonds in the area whenever they see someone wearing a diamond ring in the first place.

To sum up, I think that our intuitions regarding *Barn* are not robust. It seems to make sense to say that Henry knows that there is a barn in the field inasmuch the same way as it makes sense to say that Henry does not know. My suggestion is that it all boils down to whether one regards Henry as being epistemic irresponsible. If one interprets *Barn* as closer to *Barn'*, then one is inclined to accuse Henry of being epistemically irresponsible, and thus deprives him of knowledge. If one interprets *Barn* as closer to *Diamond*, one will be inclined to exempt Henry from being epistemically irresponsible, and thus grants him the knowledge. Nozick's account has made a wrong prediction that our intuitions regarding *Barn* are definite; the aforementioned cases (*Barn*, *Barn'*, *Thermometer*, *Thermometer**, *Diamond*, etc.) show that our intuitions in effect vary with the details of the story that are related to whether the subject appears to be epistemically responsible. The tracking-based account, by contrast, offers a neat and plausible explanation for *Barn* and its variants.

5.6 Counterexample VI: Pseudo Skeptical Hypotheses

So far, the tracking-based account has done a pretty good job in dodging the counterexamples targeting the traditional tracking account of knowledge. However, one

might still think that we can modify some of the counterexamples mentioned above to make them applicable to the tracking-based account.

For instance, one might modify *Oscar* as follows:

Oscar*. Like the original case, Oscar is standing in an open filed facing Dack the dachshund. He thus forms a belief that the object over there is a dog. Now suppose that someone has decided to make Oscar believes that the object over there is a dog. If Dack the dachshund is not in front of Oscar, this person will send a wolf painted to look exactly like Dack to face Oscar.

Let us assume that Oscar cannot distinguish between Dack and the painted wolf. It seems that in this case, if the object over there were not a dog, Oscar would have seen a painted wolf, and thus come to believe that the object over there is a dog. More importantly, some might argue that Oscar's sole reason for *The object over there is a dog* will remains the same—after all, Oscar cannot distinguish Dack the dachshund from the painted wolf. This seems to show that according to the tracking-based account, Oscar does not know that the object over there is a dog in *Oscar**. But is this a correct verdict?

Some might try to reply on behalf of the current account in the following way: It is in effect *plausible* to claim that Oscar does *not* know that the object over there is a dog, since Oscar's perceptual experience of seeing is a dog is not a good indicator of the existence of dog in this case—no matter whether the object Oscar is facing is a dog or not, he would still have the vision experience that he is facing a dog.

Unfortunately, this reply does not work. Even if we grand that Oscar does not know that the object over there is a dog, the same verdict cannot be made in other similar cases. Matthew McGrath once raised the following case to me:

Grandfather. Suppose that Nancy's grandfather was born in 1945. Nancy believes that he was born in 1945 via his testimony. Suppose, for some reasons, her grandfather are obsessed to make others to believe that he was born in 1945 such that if he were not born in that year, he would still try to deceive everyone to

believe that he was born in 1945 (suppose he is very good at this, and there is a very high chance for him to successfully deceive everyone around him).

It seems plausible to think that Nancy *knows* that her grandfather was born in 1945. But it seems that her sole reason (i.e., her grandfather's testimony) is not sensitive—if he were born in 1946, he would have deceived others that he was born in 1945, and Nancy would thus have believed that he was born in 1945 (or she would have the same testimony that he was born in 1945).

Colin McGinn offers a similar case in his (1984):

Benevolent Deity. Suppose we are living in a universe in which there also exists a benevolent deity who watches over our sensory input: he has the intention to preserve this input by artificial means in the event of a cataclysm in which the material objects that actually produce it should suddenly go out of existence. Let us suppose that this cataclysm is, in fact, physically possible and that the deity has the power to carry out his intention. Then it seems that we have the truth of this counterfactual: 'If the objects around me were to go out of existence, I would still believe that I was surrounded by those objects'—since the deity would see to it that my experience sustained this belief were the cataclysm to occur. (McGinn 1984, 531)

I agree with McGinn that it seems that the agent in such a universe knows that one is surrounded by material objects. However, it seems that the following counterfactual is also true in this universe: If there were not material objects, one would still have believed that one is surrounded by material objects.

The main point of the aforementioned cases is that they are so stipulated such that (i) the subject believes that p via the reason R, and that (ii) in some *pseudo* skeptical scenarios, if p were false, the subject would still have R. A skeptical scenario is pseudo if it is not sufficient for depriving one of knowledge.[63] Nonetheless, a pseudo skeptical

---

[63] Intuitively, we can tell whether a skeptical hypothesis is pseudo or not. But it is a pain in the neck to specify exactly what condition determines a pseudo hypothesis. For simplicity's sake, I will not address this issue any further here.

scenario, because of its *skeptical* feature, may still render the subject's sole R for p insensitive. Hence, on the tracking-based account, S should lack knowledge of p in pseudo skeptical scenarios. But our intuitions tell us the opposite: It seems obvious that the subject knows in such cases.

This objection is essentially misguided. However, showing how it is misguided has proven to be extremely useful. The core idea is that there are essentially two kinds of counterfactualization. More precisely, a statement with the form 'If p were not the case, then q would be the case' can be interpreted as what I call intervention- or inference-counterfactualization. In the next chapter, I will argue that the cases considered in this section have prompted us to use inference-counterfactual interpretation on "If p were not the case, S would still have R". I will show that those counterfactual conditional is indeed true under inference-counterfactual interpretation. However, the objection is still misguided since the interpretation that is relevant to knowledge is intervention-counterfactual, not inference-counterfactual. On the intervention-counterfactualization, "If p were not the case, S would still have R" is indeed false. Counterexamples from pseudo skeptical hypotheses are nonetheless very illuminating since they help us to focus on the nature of counterfactualization.

The distinction between intervention- and inference-counterfactualization has generally escaped epistemologists. But a comprehensive sensitivity account should not conflate them. I will come back to this issue the next chapter.

5.7 Taking Stock

In this chapter, I have considered the most prominent counterexamples to the tracking account of knowledge (specifically, the sensitivity condition). While many of these cases post serious problems to Nozick's account, they are not as effective as against the tracking-based account. The present account, as noted in the last chapter, has corrected some important flaws in Nozick's original account of tracking. It is a surprising and happy result that refining in such a way also exempts us from many prominent counterexamples on the market.

The moral, as I see it, is twofold. On the one hand, it shows that the main problem for the traditional (i.e., Nozickean) tracking account of knowledge consists in the fallacious assumptions people tend to make in characterizing the tracking property pertaining to knowledge—a tendency that is greatly influenced by Nozick's work, more precisely, the assumptions that (a) adherence is essential to tracking, that (b) sensitivity (and adherence) is a property of belief, and that (c) tracking is incompatible with (Closure). In the previous chapter, I offer principled reasons for rejecting all these assumptions. In this chapter, we see another merit for their denial—it frees us from the alleged counterexamples to the tracking account. On the other hand, we are now in a position to notice that the basic idea of the tracking account has so far remained intact. The tracking account is motivated by the consideration that knowledge typically requires tracking the truth; this motivation survives almost all prominent objections once the notion of tracking has been casted out properly. In sum, there is no decisive objection to the tracking account, at least not when it is characterized in terms of (K_T).

However, at the end of our discussion, I mention a set of counterexamples that involves pseudo skeptical scenarios. So far, I did not offer a reply to objections along this

line. The main reason is that an appropriate account requires us to have a deeper understanding of the truth conditions of counterfactuals, in particular, the distinction between what I call the interventional and inferential counterfactualization. In the next chapter, I will take up this task.

Chapter Six

Counterfactualization and Sensitivity

6.1 Introduction

Where are we now? This dissertation aims at arguing for a kind of standing-variantist solution, which claims that (i) the truth values of ordinary knowledge attributions are sensitive to pragmatic factors of some sort, and that (ii) such a pragmatic sensitivity of knowledge attributions gives rise to the fact that the skeptic's denials of knowledge and the ordinary person's attributions of knowledge are both true in their respective contexts. Starting from Chapter 4, I have been building up an account of knowledge—the tracking-based account—that justifies (i). The present chapter plays the role of connecting (i) and (ii). On the one hand, this chapter will offer an in-depth analysis of the notion of sensitivity that allows us to reply to a prominent line of objection, which has been left unresolved in the last chapter. On the other hand, the result of this chapter will shed light on how knowledge attributions can be sensitive to pragmatic considerations as suggested by standing variantism (this last claim will be further developed in the next chapter).

The tracking-based account relies on the *sensitivity condition*—whether one's epistemic position with respect to p satisfies the standards for knowledge in play or not depends on whether one's reason for p is sensitive or not (i.e., (K_T)(ii); see Section 4.6).

Following Nozick and Dretske, I formulate the sensitivity condition of knowledge in terms of counterfactual conditionals (i.e., the counterfactual-conditional formulation; see Section 4.2). That is, S's sole reason R for p is sensitive if and only if S would have lacked R if p were false. In Chapter 5, I argue that sensitivity, understood in this way, is immune to almost all prominent objections to the traditional sensitivity account (of knowledge). I have, however, left an important branch of counterexamples untouched, viz., cases involving pseudo skeptical hypotheses (Section 5.6).

In my opinion, the problem of cases involving pseudo skeptical hypotheses runs deep; we will not be able to see what goes wrong with such alleged counterexamples unless we are clear about the truth condition of counterfactual conditionals. Yet, so far, I have said little on this score. In this chapter, I will take up this task.

Traditionally, epistemologists have followed David Lewis (1973a) in analyzing the truth condition of counterfactual conditionals in terms of the *possible-worlds semantics*.[64] In this chapter, I want to suggest epistemologists to give up the possible-worlds talk and shift to the *causal-modeling* talk. That is, I suggest that epistemologists should opt for an alternative semantics—the *causal modeling semantics* proposed by Judea Pearl (2000), and others—when dealing with the sensitivity condition. Two reasons favor the causal modeling semantics over the possible-worlds one, at least from an epistemological point of view. First, a variety of cases indicate that the notion of causation is central in analyzing the notion of counterfactualization. The causal modeling

---

[64] Lewis (1973a) and Stalnaker (1968) have been two early developments of the possible-worlds semantics of counterfactuals. Nuances aside, the main difference between Lewis's and Stalnaker's semantics is that the former rejects while the latter accepts what Lewis calls 'the limit assumption', i.e., that there is a unique closest possible to the actual world. Also see (2003, Chapter 11). For simplicity's sake, I will focus on Lewis's semantics.

semantics, which basically characterizes counterfactualization in terms of causation, is capable of explaining such cases, while the possible-worlds semantics, at least in its orthodox form, fails to do so. Secondly, The casual modeling semantics, but not the possible-worlds semantics, is able (or has the resources) to account for the distinction between two kinds of counterfactualization, which I will call the *interventional* and *observational counterfactualization* respectively. This distinction, which has been unfortunately widely ignored, is of utmost importance to epistemology, since failing to recognize the distinction has been responsible for a common misunderstanding of the sensitivity condition, namely, the widespread idea that cases involving pseudo skepticism are genuine counterexamples to the sensitivity condition (see Section 6.11.1). Hence, for the purposes of epistemology, one should adopt the causal modeling semantics.

The following consists of eleven sections. In Section 6.2, I describe the standard form of the truth condition of counterfactual conditionals, which is shared by both the possible-world and the causal modeling semantics. In Section 6.3, I review the Lewis's possible-worlds semantics of counterfactuals, which has been widely adopted, in one way or another, by proponents of the sensitivity account. Epistemologists often assume that the comparative similarity between worlds should be characterized in terms of the notion of overall similarity. I argue, in Section 6.4, that the overall-similarity proposal simply does not work. In Section 6.5, I consider how philosophers may shift to adopt Lewis's famous system of weights of similarity. I argue that this move does not save epistemologists since there are some serious problems for Lewis's proposal. In Section 6.6, I look at Jonathan Schaffer's amendment of Lewis's system of weights of similarity, which appeals to the notion of causation. In Section 6.7, I discuss Robert Northcott's

objection to the Lewis-Schaffer semantics of counterfactuals. Northcott's objection leads us to the important distinction between the backtracking and forward-tracking counterfactualization, which poses an initial problem for the possible-worlds semantics. Lewis's original reply to this problem is to appeal to the asymmetry of temporal order in counterfactuals. In Section 6.8, I argue that temporal order is just a red herring. The distinction between backtracking and forward-tracking counterfactualization is a special case of a distinction between two kinds of counterfactualization that escape the possible-worlds semantics. In Section 6.9, I describe the general formulation of the causal modeling semantics. In Section 6.10, I use the causal modeling framework to account for the distinction between the interventional and inferential counterfactualization of which the distinction between the backtracking and forward-tracking counterfactualization is a special case. In Section 6.11, I show that distinguishing interventional counterfactualization from the inferential one enables us to see clearly that cases involving pseudo skeptical hypotheses are not genuine counterexamples to the sensitivity condition. The last section wraps up the main point of this chapter.

6.2 The Standard Form

For epistemologists, Lewis's possible-worlds semantics has been the orthodox view on counterfactual conditionals.[65] However, by arguing for the causal modeling semantics, I do not mean to sacrifice much of the (semantic) merits of the orthodox view. In fact,

---

[65] Nozick (1981) appeals to the possible worlds semantics in explicating the sensitivity and adherence conditions, although he explicitly mentions that the appeal to such a semantics serves only for illustrative purposes. Later on, it has become a prevalent trend for epistemologist to use the possible-worlds semantics to characterize the 'modal' or 'counterfactual' dimensions of the notion of knowledge (cf., e.g., DeRose 1995; Sosa 2000; Vogel 2000; Zalabardo 2012, to mention just a few).

these two semantics, while motivated by radically different considerations, are formally equivalent in the sense that they share the same general structure in characterizing the truth condition of counterfactual conditionals (call it the *standard form*) (Briggs 2012). Let 'A > C' stands for counterfactual conditional *If A had obtained, B would have obtained*. The standard form can be specified as follows:

> (S) "A > C" is true at a point *a* if and only if C is true at a selected set of points *b* in which A is true (more precisely, for all b ∈ *f*(A, *a*), C is true in *b*). (Briggs 2012, 140–1)

A *point* is a semantic counterpart of the worldly condition; the standard form leaves it open whether a point is to be interpreted as a possible world (Lewis) or a causal model (Pearl). The function *f* is a selection function which maps the antecedent and the point *a* to a set of points *b*; the standard form leaves it open whether *f* is to be characterized as a function of *comparative similarity* between worlds (Lewis) or a function of *intervention* on a casual model (Pearl).

(S) has manifested the general structure of how we normally evaluate whether a counterfactual conditional is true or not. For instance, people in general agree that the counterfactual conditional "If Oswald had missed, Kennedy would have been alive" is *actually* true (or true in the actual world). How do we arrive at this conclusion? Intuitively, we first envision the (relevant) counterfactual situation, which is significantly similar to the actual situation except that Oswald had missed. We then further decide whether Kennedy would have been alive in such a situation. If the answer is "Yes", we regard the conditional as true, otherwise false. The point in (S) is a semantic counterpart of an actual/counterfactual situation, while the selection function *f* is a semantic mechanism that specifies the relevant counterfactual situations based on the antecedent

and the actual situation.[66] We normally take a counterfactual conditional to be true if it is true in (all) the selected counterfactual situations. (S) is just a formal way to capture this normal process of evaluating counterfactual conditionals. Hence, by shifting from the possible-worlds semantics to the causal modeling semantics, we do not loss much of the intuitive features of the former.[67]

6.3 The Possible-Worlds Semantics

Let us first introduce Lewis's possible-worlds semantics. According to Lewis (1973a), the truth condition of a counterfactual conditional is specified as follows:

(PW)   "A > C" is (non-vacuously) true in w_1 if and only if some A-world w_2 in which C is true is closer to w_1 than any A-world in which C is false.

Some clarifications are called for. First, (PW) is a special case of (S) with *points* and the *selection function* being interpreted as *possible worlds* and the *function determining the distance between worlds* respectively. Second, 'p-world' denotes worlds in which p is true. Similarly, '~p-world' denotes worlds in which "not-p" is true, while 'p⊕q-world' denotes worlds in which "p⊕q" is true (where '⊕' stands for two-place truth functional connectives). Third, "A > C" is vacuously true if there is no A-world. Finally, the

---

[66] Note: A counterfactual situation is to be envisioned based on an actual situation. Here, whether a situation is considered 'counterfactual' is relative to a certain situation that is considered 'actual'. A situation that is considered 'actual' in this sense may not be an actual situation. For instance, people familiar with the story of *Superman* may agree that "If Superman had not been a journalist, he would have destroyed the earth" is false. Here, the falsity of this counterfactual condition depends on a counterfactual situation (i.e., that Superman had not been a journalist) and an 'actual' situation (i.e., that Superman is a journalist). In what follows, I will continue to use the term 'counterfactual' and 'actual situation'. The readers should keep in mind that what is labeled 'actual' may not be real.

[67] The possible-worlds semantics and causal modeling semantics are formally different in other aspects. As Briggs (2012) points out, the possible-worlds semantics and the causal modeling semantics, when elaborated, count difference inferences as valid. Specifically, *modus ponens* for counterfactual conditionals counts as valid inferences according to the possible-worlds semantics but not to the causal modeling semantics. In what follows, I will ignore the logical issues since my main concerns here is formal semantics, not logic.

selection function *f* is characterized as the *comparative distance* between worlds such that the function selects a set of A-worlds which is closer to w_1 than any other ~A-world is. The distance between worlds is further characterized as a function of the *comparative similarity* between worlds.

Now, epistemologists have appealed to (PW) to characterize the truth condition of the sensitivity condition (i.e., the condition that S would not have believed that p if p were false). They specifically take A-worlds w_2 in (PW) to be the set of *nearby* A-worlds—A-worlds in the close vicinity of the actual world. The similarity (or distance) between worlds is further determined by the overall similarity between worlds. Call it 'the overall-similarity proposal'.

6.4 Against The Overall-Similarity Proposal

While many epistemologists accept the overall-similarity proposal, little attention has been paid to the fact that the overall-similarly proposal simply does not square well with our intuitions.

Consider the following case:

Holocaust. Nixon was sitting in a room deciding whether to push a button in front of him, which would detonate the nuclear bombs. Lauren was outside the room and she knew that if the button were pressed, there would have been a nuclear holocaust. There had not been a nuclear holocaust.[68]

It seems that Lauren knew, in *Holocaust*, that Nixon did not press the button (since she did not see a nuclear holocaust happen). If knowledge requires sensitivity, Lauren's belief

---

[68] Readers familiar with literatures on counterfactual conditionals should notice that *Holocaust* is derived (with modification) from Kit Fine (1975). Fine's discussion indicates that the distance between worlds in (PW) cannot be evaluated in terms of overall similarity, otherwise (PW) is bounded to give incorrect verdicts to cases like *Holocaust*. This point has been widely accepted by philosophers of language. It is quite surprising that epistemologists are content to use a proposal that has been widely rejected in another sub-discipline.

that Nixon did not press the button has to be sensitive (if we take sensitivity to be a property of belief) or her reason for *Nixon did not press the button* has to be sensitive (if we take sensitivity to be a property of reason). Unfortunately, the overall-similarity proposal is unable to deliver this verdict.

Let w_0 be the actual world. Consider two possible worlds w_1 and w_2. In w_1, Nixon pushed the button, the nuclear bombs were detonated, and the holocaust occurred; in w_2, Nixon pushed the button, but due to a small miracle, the bombs were not detonated, and no holocaust occurred. On the overall-similarity proposal, w_2 is more similar to w_0 than w_1 is—the world without the nuclear holocaust is surely more similar to the world in which the holocaust occurs is. If we take w_2, but not w_1, to be the nearby world, Lauren's belief that Nixon had not pressed the button was not sensitive—in all nearby worlds in which Nixon pressed the button, Lauren would still have believed that Nixon did not push the button. Nor is Lauren's reason for *Nixon did not push the button* sensitive—in all nearby worlds in which Nixon pushed the button, Lauren would still have the same reason (e.g., not seeing the holocaust) for *Nixon did not push the button*. If we take both w_1 and w_2 to be nearby worlds, Lauren's belief would still be insensitive—in some nearby worlds (i.e., w_2), Lauren's belief is not sensitive, likewise for the sensitivity of Lauren's reason—in some nearby worlds (i.e., w_2), Lauren's reason for *Nixon did not press the button* is insensitive. This shows that the overall-similarity proposal fails to give the appropriate account of counterfactualization necessary for the sensitivity account.

6.5 Lewis' System of Weights of Similarity

Lewis (1979b) tries to avoid the problem posed by *Holocaust* (and others) by offering a system of weights of similarity. Perhaps epistemologists can adopt Lewis's proposal in explicating the standard of sensitivity. In this section, I want to argue that this does not work. More importantly, learning how Lewis's proposal fails paves the way for the causal-modeling semantics.

Lewis suggests that the comparative similarity between worlds should be weighed in the following ways:

(1)　　It is of the first importance to avoid big miracles or big quasi-miracles.

(2)　　It is of the second importance to maximize the region of perfect match.

(3)　　It is of the third importance to avoid small miracles or small quasi-miracles.

(4)　　It is of the fourth importance to maximize the region of imperfect match.[69]

Clarifications. *Miracles* here refer to the violation of the (actual) physical laws. Suppose that the actual physical laws are indeterministic. An indeterministic outcome is counted as a *quasi-miracle* if it seems to "conspire to produce a pattern" (Lewis 1986, 60). In other words, a quasi-miracle, though physically possible, is an event "which is both low probability and which has a pattern which is, by our lights, *remarkable*" (Hawthorne 2005, 398).

Let 'PRESS > HOLOCAUST' stand for *If Nixon pressed the button, there would have been a holocaust*, w_1 for the actual world, w_2 for the world in which Nixon pressed the button, the bomb was not detonated (due to a small miracle), and the holocaust happened (i.e., w_2 is a PRESS&~HOLOCAUST-world), and w_3 for the world in which

---

[69] Here, for the sake of discussion, I adopt Schaffer's (2004b, 301) formulation of Lewis's conditions.

Nixon pressed the button, the bomb was detonated, and the holocaust happened (i.e., w_3 is a PRESS&HOLOCAUST-world).

Accompanied with (1)-(4), (PW) can give a correct verdict to "PRESS > HOLOCAUST": while a PRESS&~HOLOCAUST-world like w_2 may maximize the region of imperfect match, it gains this feature at the cost of containing a small miracle, say, the signal of pressing the button miraculously fails to transmit to the bomb. By contrast, while a PREES&HOLOCAUST-world like w_3 may fail to maximize the imperfect match, it also does not contain small miracles. Given that (3) is weighed more than (4) is, w_3 is more similar to w_1 than any PRESS&~HOLOCAUST-world like w_2 is. As a result, "PRESS > HOLOCAUST" is true in w_1, according to (PW) and (1)-(4).

An important feature of Lewis's semantics, which is in sharp contrast to the causal modeling semantics, is that, when determining which worlds are relevant to the evaluation of counterfactuals, the notion of causation plays no significant role. This may seem surprising since psychological studies have shown that counterfactual reasoning is intertwined with causal reasoning (cf. Spellman, Kincannon, and Stose 2005).

Lewis's reason for not incorporating the notion of causation into (1)-(4) is that he wants to reduce the notion of causation to the notion of counterfactual dependency (cf. Lewis 1973b). Lewis's reductive approach to causation, however, is questionable. Firstly, notice that the notion of causation, arguably, is one of those primitive notions that constitute our basic ways of conceiving the world; causation is so basic that it could hardly be defined, at least not non-circularly, by others (Edgington 2011, 230–3). If one does not share Lewis's idea that causal dependency depends on counterfactual dependency, there seems to be no principled reason for accepting (1)-(4). Secondly, this

approach seems to get things backward. As we will see, (1)-(4) has suffered from counterexamples that seem to be insurmountable unless the notion of causation is somehow incorporated into them. In what follows, I will discuss three kinds of such counterexamples.

Ryan Wasserman (2006) has offered the following counterexample to (1)-(4):

Bomb. Imagine a deterministic world … that is much like our own in its distribution of objects and qualities, but which contains a black box in the middle of the Milky Way. In the black box there is a beetle and a button. If the button is pushed, a signal will run along a wire and out of the box. Beyond the wire, there are no causal avenues running out of the black box—whatever happens in the box stays in the box. The wire is connected to a 'mega-bomb' which is lightening fast and deadly powerful—if the mega-bomb explodes, everything in the future light cone of the bomb will be destroyed. But the universe is spared. The beetle does not strike, the bomb does not destroy. Let us suppose, finally, that the black box and all of its contents is destroyed in a lawful manner shortly after *t*. (Wasserman 2006, 59)

Let 'PUSH > DESTROY' stand for *If the beetle had pushed the button, the universe would have been destroyed*. Intuitively, "PUSH > DESTROY" seems true in *Bomb* but (1)-(4) give the opposite verdict. Let w_4 be the actual world in which the beetle does not push the button and the universe is not destroyed (i.e., w_4 is a ~PUSH&~DESTROY-world), w_5 be a world in which the beetle, due to a small miracle, pushes the button and the universe is destroyed (i.e., w_5 is a PUSH&DESTROY-world), and w_6 be a world in which the beetle, due to a small miracle, pushes the button but the universe is not destroyed (i.e., w_6 is a PUSH&~DISTROY-world)—after the beetle pushes the button, another small miracle occurs so that the signal of pushing the button is not transmitted to the mega-bomb. On (1)-(4), it seems that w_6 is closer to w_4 than w_5 is—w_6 contains more small miracles than w_5 does, but the former also has a larger region of perfect match (w_6 matches w_4 perfectly after the occurrence of the second miracle) than the latter does, and

given that the maximization of the region of perfect match (i.e., (2)) is weighed more than the avoidance of small miracles (i.e., (3)), w_6 should be weighed more similar to w_4 than w_5 is. In other words, (1)-(4) predict that "PUSH > DESTROY" is false at w_4, contradicting our intuitive judgments.

Bomb has presupposed determinism. Some might thus argue that the culprit was not (1)-(4), but determinism. In fact, given that the traditional mechanical (deterministic) idea of the physical world is incompatible with quantum physics, a plausible semantics for counterfactual conditionals, some might argue, needed not be compatible with determinism. This reply is unsuccessful on two scores. First, philosophers have generally agreed that the semantics for counterfactual conditionals should not depend on whether determinism is true or not. While people's opinions about determinism may be widely divided, their opinions regarding the truth values of counterfactual conditionals tend to converge. Second, and more importantly, there are other counterexamples to (1)-(4) that do not presuppose a deterministic world.

Michael Slote (1978) once reported Sidney Morgenbesser's counterexample as follows:

> *Bet*. Imagine a completely undetermined random coin. Your friend offers you good odds that it will not come up heads; you decline the bet, he flips, and the coin comes out heads. He then says: "you see; if you had bet (heads), you would have won." (Slote 1978, 27, Note 33)

Let "BET > WIN" stands for the utterance "If the hearer had bet (heads), the hearer would have won". Intuitively, "BET > WIN" seems true in *Bet*.[70] However, (1)-(4)

---

[70] Slote admits that he does not know how to handle this kind of cases. But he then quickly suggests that "perhaps it simply *isn't* correct and the correct retort to it is: "no, you're wrong; if I had bet (heads), the coin might have come up differently and (so) I might have lost—assuming the coin was random" (1978, 27). I am inclined to agree with Dorothy Edgington that this is just Slote's "wishful thinking (wishful

cannot account for our intuitions regarding *Bet*. Let w_7 be the actual world in which *Bet* occurs (i.e., w_7 is a ~BET&~WIN-world), w_8 be the world in which the speaker flips heads and the hearer bets heads and won (i.e., w_8 is a BET&WIN-world), w_9 be the world in which the speaker flips tails and the hearer bets heads and does not win (i.e., w_9 is a BET&~WIN-world). According to (1)-(4), w_8 and w_9 are equally similar to w_7: On the one hand, both w_8 and w_9 contain only one small miracle that changes the hearer's behavior from declining the bet to accept it. On the other hand, w_8 and w_9 have the same region of perfect and imperfect match (since the coin is completely random, either the coin lands heads or it lands tails is equally possible). However, if w_8 and w_9 are equally similar to w_7, (PW) and (1)-(4) predict that "BET > WIN" is not true in *Bet*, which is implausible.

Along a similar line, Jonathan Schaffer has proposed the following counterexamples:

> *Bet**. At indeterministic [w_{10}], Fred is a compulsive coin flipper, and Lucky a compulsive gambler. Lucky bets that Fred will not flip a million tails in a row. Fred flips a million tails in a row, so Lucky loses the bet. (Schaffer 2004b, 304)

Let "TAILS > WIN" stands for "If Lucky had bet all tails, he would have won". Intuitively, "TAIL > WIN" is true in w_{10}. However, (1)-(4) and (PW) seem to imply that "TAILS > ~WIN" is true. Let w_{11} be the actual world in which Fred bets all tails, and Fred flips a million tails in a row (i.e., w_{11} is a TAILS&WIN-world), and w_{12} be the world in which Lucky bets all tails, and Fred does not flip a million tails in a row (w_{12} is a TAILS&~WIN-world). Now, both w_{14} and w_{15} contain a small miracle, namely, Lucky bets that Fred will flip a million tails in a row (i.e., TAILS obtains). But unlike w_{12}, a

---

TAILS&WIN-world like w_{11} essentially contains a quasi-miracle that Fred flips a

millions tails in a row. Given (3), a world contains a quasi-miracle is less similar to the

actual situation than the one that lacks it is, other things being equal. It follows from (1)-

(4) and (PW) that w_{12}, which is a TAILS&~WIN-world, is more similar to w_{10} than any

TAILS&WIN-world (such as w_{11}) is. Therefore, "TAILS > ~WIN" turns out to be true in

w_{11} according to (1)-(4) and (PW), against our intuitions.

6.6 Schaffer's Refinement

The aforementioned counterexamples have shown clearly that Lewis's system of weights

of similarity fails to do justice to our intuitions regarding counterfactual conditionals.[71] A

moral usually drawn from them is that (1)-(4) need to be complemented by appropriate

causal constraints. For instance, Dorothy Edgington writes, "I don't see how Lewis can

handle these examples without appealing to the notion of causal independence" (2004,

20). Schaffer also draws the same conclusion, "These counter-examples all suggest a

need to invoke *causal independence*" (2004b, 305, original italics).

One way one might amend (1)-(4), hence, is to incorporate the appropriate casual

constraints. In what follows, I will focus on Schaffer's amendment:

(1c)   It is of the first importance to avoid big miracles.

(2c)   It is of the second importance to maximize the region of perfect match,
       from those regions causally independent of whether or not the antecedent
       obtains.

(3c)   It is of the third importance to avoid small miracles.

(4c)   It is of the fourth importance to maximize the spatiotemporal region of
       approximate match, *from those regions casually independent of whether*

---

[71] For more objections, see, e.g., Elga (2001) and Hawthorne (2005).

The main difference between (1)-(4) and (1c)-(4c) is that the latter, but not the former, has certain causal constraints on the maximization of the regions of perfect and imperfect match. On Schaffer's system, the maximization of the regions that are causally dependent of whether or not the antecedent of the counterfactual conditional obtains has no weight in the calculation of similarity between worlds; the perfect or imperfect match of the regions, which is dependent on the antecedent, does not matter as far as counterfactual conditionals are concerned.

(1c)-(4c) and (PW) are able to give intuitive verdicts to all the counterexamples mentioned above (Schaffer 2004b, 306–7). For instance, in *Bomb*, the PRESS&HOLOCAUST-world w_5 is more similar to w_4 (i.e., the world represents *Bomb*) than the PRESS&~HOLOCAUST-world w_6 is: w_5 contains fewer small miracles than w_6 does. Moreover, the region of perfect match of w_6 causally depends on whether PRESS obtains or not. Such a match should have no bearing on the weighing of similarity between worlds, according to (2c). Hence, other things being equal, (3c) predicts that w_5 is more similar to w_4 than w_6 is. By (PW), "PRESS > HOLOCAUST" is thus true in *Bomb*, as desired.

Consider *Bet**. That Fred flips a million tails in a row is causally independent of whether Lucky bets or not. Hence, the world in which Lucky bets Fred flips tails a million time in a row and Fred does (i.e., w_{11}) contains more causally independent imperfect match than does the world in which Lucky bets Fred flips tails a million time in a row and Fred does not (i.e., w_{12}). As a result, other things being equal, w_{11} is more similar to w_{10} than w_{12} is, according to (4c). Similar stories can be told about *Bet* as well.

(1c)-(4c) preserve the core of (1)-(4) and is able to dodge many prominent counterexamples of the latter. The price to pay here is that the possible-worlds semantics is no longer independent of the notion of causation; the notion of causal dependence is now built in (1c)-(4c). This will not be a problem if one does not also share Lewis's ambition of analyzing the notion of causal dependence in terms of counterfactual dependence.

Schaffer's amendment of Lewis's semantics is indeed a step closer to the casual modeling semantics, which appeals to the notion of causation in explicating the truth conditions of counterfactual conditionals. In the next section, I will argue that the possible-worlds semantics has failed to respect a very important distinction of counterfactuals, which I will call the distinction between *interventional* and *inferential counterfactualization*. As it turns out, this distinction is crucial to the specification of the sensitivity condition—the sensitivity of knowledge required interventional counterfactualization but not referential one. I take it that failure to distinguish these two kinds of counterfactualization has prompted philosophers to regard, mistakenly, cases involving pseudo-skeptical scenarios as genuine counterexamples to the sensitivity condition of knowledge (Section 6.11).

6.7 Northcott's Counterexample and Backtracking

In his (2009), Robert Northcott has tried to offered a counterexample to Schaffer's amendment to (1)-(4). Northcott's objection, to me, falls short of being conclusive against (1c)-(4c) since it relies on a kind of counterfactualization (i.e., backward counterfactualization), which has been explicitly shunned by traditional proponents of the

possible-worlds semantics. Nevertheless, Northcott's objection is still worth considering since it points to a direction in which the genuine problem of the possible-worlds semantics lies.

Consider the following cases:

Smoking. John is smoking in a poorly ventilated room. Two causal effects of the smoking are that the smell of tobacco smoke gets caught in his hair and that his eyes itch (cf. Northcott 2009, 337–8).

Now, let "SMOKE", "SMELL", and "ITCH" stands for *John is smoking in a poorly ventilated room, The smell of tobacco smoke gets caught in his hair*, and *John's eyes itch* respectively. Now, consider "(~SMELL & ~ITCH) > ~SMOKE", which Northcott interprets as "If we were to clean John's hair and relieve his eyes, we would thereby prevent his smoking". According to Northcott, "(~SMELL & ~ITCH) > ~SMOKE" is intuitively false in *Smoking*. Let w_{13} be the world of *Smoking* in which John smokes, the smell of tobacco smoke gets caught in his hairs, and his eyes itch (i.e., w_{13} is a SMOKE&SMELL&ITCH-world), w_{14} be a world in which John smokes, but due to a big miracle, the smell of tobacco smoke does not get caught in his hair, nor do his eyes itch (i.e., w_{14} is a SMOKE&~SMELL&~ITCH-world), w_{15} be a world in which John does not smoke, due to a small miracle. Since SMELL and ITCH are the effects of SMOKE, ~SMELL and ~ITCH hold in w_{15} as well (we may assume that there are no other causes for SMELL and ITCH in *Smoking*) (i.e., w_{15} is a ~SMOKE&~SMELL&~ITCH-world).

Northcott argues that "(~SMELL & ~ITCH) > ~SMOKE", while intuitively false, is predicted as true by (1c)-(4c) and (PW). The idea is that w_{15} is more similar to w_{13} than w_{14} is since the former contains a *small* miracle (e.g., John changes from smoking to not smoking) while the latter contains a *big* miracle (e.g., the smell of John's hair is totally

eradicated *and* his eyes totally relieved)[72], since it is more important to avoid big miracles (i.e., (1c)) than to avoid small miracles (i.e. (3c)). The. As a result, "(~SMELL & ~ITCH) > ~SMOKE" is true in *Smoking*, according to (PW) and (1c)-(4c), but this contradicts, Northcott argues, our intuitions. Of cause, as Northcott points out, w_{14} has a larger region of perfect match than w_{15} does—in the former (but not in the latter) John still smokes, but this consideration does not affect the previous conclusion, since avoiding big miracles (i.e., (1c)) is still more important than maximizing the region of perfect match (i.e., (2c)).[73]

Northcott's case is important, since the conditional involved in this case, i.e., "(~SMELL & ~ITCH) > SMOKE", is actually a *backward* conditional, whose consequent is about "a time earlier than any that its antecedent is about" (Bennett 1984, 57). Proponents of the possible-words semantics might reply to Northcott's objection by claiming that (1)-(4) (or even (1c)-(4c)) and (PW) do not deal with backward conditionals. Let me elaborate.

What is problematic about backward conditionals is that such conditionals, if true, indicate that an event at an earlier time might be counterfactually dependent of an event at a later time. This falls into a broader category of 'backtracking counterfactualizing'— counterfactualizing back in time (and then forward again) (Bennett 2003, 208). Lewis is well aware of the problems of backtracking counterfactualization: First, ordinary

---

[72] Lewis is unclear about what counts as 'big' or 'small' miracles. But for the purpose of this argument, we may bypass this issue. Notice that that (~SMELL & ~ITCH) is a conjunction. In principle, we can complicate this conjunction by adding more conjuncts, and doing so will not affect the core of Northcott's argument. Hence, one can add as many conjuncts as one likes, until one feels that the resulted complex conjunction counts as a big miracle. Northcott's original example has three conjuncts. I have simplified his example for the sake of discussion.

[73] The reader may notice that since Schaffer's amendment has no effect on determining the truth value of "(~SMELL & ~ITCH) > ~SMOKE", the same objection also applies to (1)-(4) as well.

attributions of counterfactual conditionals are, in general, not backtracking in character. Second, and more importantly, Lewis does not want his semantics to allow for the counterfactual dependence of earlier events on later events, since this view undermines his project of analyzing causal dependence in terms of counterfactual dependent—a cause, which usually happens before the effect, is not causally dependent on its effects. Hence, Lewis has explicitly banned backtracking counterfactualization in his semantics, claiming that it is *non-standard* usage of counterfactual conditionals (cf. Lewis 1979b). Following this line of thought, it is not hard to imagine that proponents of the possible-worlds semantics would dismiss Northcott's counterexample by also denying that the semantics has anything to do with backtracking conditionals.

In the next section, I will argue that Lewis's rationales against backtracking are untenable. More importantly, the orthodox possible-worlds semantics has failed to appreciate an important distinction between two kinds of counterfactualization, of which the distinction between backtracking and forward-tracking counterfactualization is a special case.

6.8 Toward the Causal Modeling Semantics

Backtracking counterfactual conditionals require counterfactualizing back in time and then forward again. Are they legitimate? Normally, when people talk about counterfactual conditional, they talk about forward-tracking conditional, which counterfactualizes forward in time. For instance, "If Oswald had not killed Kennedy, the President would not have died" is about how an earlier counterfactual event (*Oswald had not killed Kennedy*) affects a later event (*Kennedy would not have died*), or how the later

event is counterfactually dependent on the earlier event.

However, backtracking conditionals are not unheard of. The following, originally

discussed by Downing (1958), is a famous case of backtracking counterfactualization:

> *Ask*. Jim and Jack quarreled yesterday, and Jack is still hopping mad. We conclude that if Jim asked Jack for help today, Jack would not help him. But wait: Jim is a prideful fellow. He never would ask for help after such a quarrel; if Jim were to ask Jack for help today, there would have been no quarrel yesterday. In that case Jack would be his usual generous self. So if Jim asked Jack for help today, Jack would help him after all. (Lewis 1979b, 456)

Let "ASK > HELP" stand for "If Jim had asked Jack for help, Jack would help him".

Intuitively, "ASK > HELP" is true in *Ask* (some might disagree. I will come back to this

issue below). More importantly, the truth of "ASK > HELP" requires backtracking

counterfactualization: We reason, in a backward manner, that in the counterfactual

situation in which ASK obtains, some event earlier than ASK would have been different,

namely, there would have been no quarrel yesterday. We then reasons forward again,

concluding that given that there had been no quarrel yesterday, and Jim had asked Jack

for help, Jack would help him after all (i.e., HELP obtains).

Lewis claims that backtracking counterfactualization is by no means *standard*:

> We ordinarily resolve the vagueness of counterfactuals in such a way that counterfactual dependence is asymmetric (except perhaps in cases of time travel or the like). Under this standard resolution, back-tracking arguments are mistaken: if the present were different the past would be the same, but the same past causes would fail somehow to cause the same present effects. If Jim asked Jack for help today, somehow Jim would have overcome his pride and asked despite yesterday's quarrel. (Lewis 1979b, 458)

Still, Lewis is forced to admit that backtracking conditionals can be true in some specific

contexts:

> A counterfactual saying that the past would be different if the present were somehow different may come out true under the special resolution of its vagueness, but false under the standard resolution. If so, call it a back-tracking

counterfactual. Taken out of context, it will not be clearly true or clearly false. Although we tend to favor the standard resolution, we also charitably tend to favor a resolution which gives the sentence under consideration a chance of truth. (Lewis 1979b, 458)

Essentially, Lewis's approach to backtracking conditionals is to dismiss them as non-standard, hence sparing his semantics from dealing with them. But, to me, there is no principled reason for regarding the forward-tracking as standard, while backtracking as non-standard. Presumably, that forward-tracking conditionals are standard is not because they are used more frequently than backtracking ones are, since frequency is a contingent matter that varies with times and contexts. Lewis's second quotation above seems to suggest that backtracking counterfactuals do not have clear truth values "taken out of context". But this cannot be the reason why they are non-standard, if at all, since even forward-tracking conditionals may not have clear truth values taken out of context.[74] Things get worse when Lewis admits that backtracking conditionals are indeed true in some contexts. If some backtracking conditionals are indeed true in some contexts, a complete semantics for counterfactual conditionals should not dismiss them.[75]

Lewis takes forward-tracking conditionals as standard since he thinks that counterfactual dependence is temporally asymmetric in that events at an earlier time (usually) do not depend counterfactually on events at a later time. What is surprising, perhaps, is that temporal asymmetry may simply be a red herring. As Jonathan Bennett

---

[74] Later Lewis also complains that backtracking counterfactuals are marked by a syntactic peculiarity. For instance, it would be natural for us to say, in *Ask*, "If Jim asked Jack for help today, there would have to have been no quarrel yesterday" (1979b, 458). But, to me, syntactic peculiarity has nothing to do with the non-standard of backtracking counterfactuals either, since not all languages have different syntactic structures for backtracking and forward-tracking counterfactual conditionals. Mandarin, for one, uses the same syntactic structure for backtracking and forward-tracking counterfactuals.

[75] In the second quotation, Lewis suggests that backtracking conditionals are true due to the rules of accommodation (cf. Lewis 1979a). But I do not see a principled reason for the idea that being true due to accommodation renders backtracking conditionals non-standard.

points out, one can reconstruct *Ask* such that it does not contain backtracking counterfactualization at all:

> *Ask**. Jack is mad at Jim this morning; Jim had no idea why. We conclude that if Jim asked Jack for help today, Jack would not help him. But wait: Jim is a prideful fellow. He never would ask for help when Jack is mad at him; Jim would ask Jack for help only if Jack were not mad at him. In that case Jack would be his usual generous self. So if Jim asked Jack for help, Jack would help him after all (cf. Bennett 2003, 206).

Arguably, "ASK > HELP" seems true in *Ask**. The only difference between *Ask* and *Ask** is that the latter does not explicitly mention yesterday's quarrel. Hence, the truth of "ASK > HELP" in *Ask** does not requires backtracking counterfactualization (Bennett 2003, 206).

The important moral we should learn from *Ask* and *Ask**, *contra* Lewis, is not that backtracking counterfactuals are non-standard. Rather, it is that there are two distinct kinds of counterfactualization; a distinction of which the distinction between backtracking and forward-tracking conditionals is a special case.

To elaborate, notice that there are two ways one may counterfactualize *Ask**: On the one hand, since Jack is mad at Jim, Jim would have been rejected if he were to ask Jack to help. Hence, "ASK > ~HELP" is true. On the other hand, since Jim would not have asked for help unless he would not be rejected, if Jim were to ask, Jack would not have been mad at him, and he would not be rejected. Hence, "ASK > HELP" is true.[76] As we can see now, temporal asymmetry, which philosophers have spent so much time discussing about, is a red herring that does not reveal the core of the distinction between

---

[76] Bennett seems to think that cases allowing both kinds of counterfactualization are illegitimate. But as far as I can tell, there is nothing unnatural or absurd in cases like *Ask* and *Ask**. I am inclined to regard such cases as posting a theoretical challenge as opposed to manifesting a theoretical error. At any rate, this is the approach I will be pursuing here.

these two kinds of counterfactualization. To me, they are just two modes of counterfactual reasoning; both are legitimate (of cause, depending on one's current purposes).

These two kinds of counterfactualization are crucial to an account of knowledge. As I shall argue, the sensitivity condition should be interpreted in terms of one kind of counterfactualization, but not the other. Failure to realize this fact has prompted philosophers to discard prematurely the sensitivity account.

6.9 Causal Modeling Semantics

I have argued that Lewis's original semantics of counterfactual conditionals has suffered from serious objections indicating that the semantics needs to be amended by appropriate causal constraints. To amend Lewis's semantics along this line pushes it a step closer to the causal modeling semantics, which appeals to the notion of causation in characterizing the selection function of the standard form. We have also seen that even if we adopt Schaffer's refinement of Lewis's semantics, which arguably dodges all counterexamples mentioned in Section 6.4, something is still missing. That is, the possible-worlds semantics still fails to respect a distinction between two kinds of counterfactualization, manifested in *Ask* and *Ask**. I want to argue that this further gives us reason to adopt causal modeling semantics—unlike the possible-worlds semantics, causal modeling semantics has the resources to account for the distinction. In the present section, I will first introduce the causal modeling semantics. I will then point out how to use this semantics to account for the aforementioned two kinds of counterfactualization in Section 6.10.

Let us start with a variant the following story (this story is adopted, with modification, from Dretske (1970)):

> *Zoo*. John goes to a zoo today and sees a zebra in the pen (or has the visual experience of seeing a zebra in the pen). Suppose that whether the zebra will be exhibited in the zoo depends on whether there is a shortage of zebras in all the zoos around the country. Moreover, the zoo has been very determined in letting its visitors to experience seeing a zebra such that if the zebra is unavailable (due to a shortage of zebras), the zoo will put a cleverly painted mule, which looks exactly like a zebra, in the pen. (The zoo policy is unknown to the public.)

We can construct a *causal model* to represent *Zoo*. In general, a causal model is supposed to capture the causal structures or relations of the events in a situation. The idea is that we construct causal models to represent the situations we encounter.

Formally speaking, a causal model is a triple $<V, S, A>$.[77] V is a finite set of variables, $\{V_1, V_2, ... V_n\}$. Such variables represent certain events in the situation that the causal model is supposed to represent. Suppose that we want to construct a causal model Z to represent *Zoo*. In Z, it is natural to take V to contain the following variables:

SHORTAGE[78] represents whether there is a shortage of zebras around the country or not.

AVAILABLE represents whether the zebra is available or not.

ZEBRA represents whether the zebra is in the pen or not.

MULE represents whether the painted mule is in the pen or not.

EXPERIENCE represents whether John has the visual experience of seeing a zebra in the pen or not.

---

[77] The causal modeling semantics has been developed by Jude Pearl (cf. Pearl 2000; also see Galles and Pearl 1998). The following formulation has been influenced by Briggs (2012). Hiddleston (2005) has constructed a different type of causal modeling semantics.

[78] In what follows, the bold character of each variable is used to indicate the very variable when we try to present the causal model using a directed acyclic graph (DAG). For DAGs, see below.

In general, each $V_i \in V$ admits a range of values. For the present purposes, we may assume that the value assignment of each V_i is mutually exclusive and jointly exhaustive, and that V_i takes on a finite number of values. In the simplest case such as Z, the variables take on two possible values—either "Yes" or "No" for each variable (for simplicity's sake, I will consider binary variables in the following discussion).

It is a convention to use a sentence of the form '$V_i = v_i$' to represent the proposition *the variable V_i takes on the value v_i*. For binary variables such as SHORTAGE, AVAILABLE, ZEBRA, MULE, and EXPERIENCE, we may use the value 1 and value 0 to stand for the value of "Yes" and "No" respectively. For instance, 'AVAILABLE = 1' represents that the zebra is available, while 'MULE = 0' represents that the painted mule is not in the pen.

The second element in a causal model, *S*, is a set of structural equations that specifies the relationships of causal dependence between variables (events).[79] The causal dependence in play may be deterministic or indeterministic. I will focus on deterministic relationships for the time being. For each variable $V_i \in V$, *S* contains at most one structural equation of the following form:

$$V_i = f_i(PA_i).$$

PA_i stands for the set of V_i's *parents*, which is a subset of V; V_i is called the *child* of PA_i. Given that structural equations are supposed to specify the causal dependence between variables (events), the equation $V_i = f_i(PA_i)$ should be interpreted as stating that the value of V_i is causally determined by the values of the variables (events) in PA_i. The function f_i,

---

[79] In what follows, I have assumed the notion of causation in characterizing counterfactualization. I make no attempt to analyze the notion of causation in play. This assumption will not render the causal modeling semantics less desirable than the possible-worlds semantics, since the latter also need a similar assumption in order to be plausible (see Section 6.6).

hence, can be viewed as a function of causal dependence. In Z, it is natural to regard S as containing the following structural equations:

$$AVAILABLE = 1 - SHORTAGE$$

$$ZEBRA = \quad AVAILABLE$$

$$MULE = 1 - AVAILABLE$$

$$EXPERIENCE = max(ZEBRA, MULE).$$

Two clarifications. First, the equation, say, ZEBRA = AVAILABLE, does not express an identity relation. Rather, as specified, it states that the value of ZEBRA is causally determined by the value of AVAILABLE. Identity relation is symmetric but causal dependence, in general, is asymmetric. Second, the structural equation ZEBRA = AVAILABLE, strictly, should be formulated as ZEBRA = f_i(AVAILABLE) with f_i being the self-identity function, viz., ZEBRA = (AVAILABLE = AVAILABLE). For simplicity's sake, I will omit the self-identify function in the following discussion.

In words, 'AVAILABLE = 1 – SHORTAGE' means that whether or not the zebra is available causally depends on whether or not there is no shortage of zebras in zoos around the country. If there is a shortage of zebras, the zebra will not be available; otherwise, the zebra is available. 'ZEBRA = AVAILABLE' means that whether or not the zebra is in the pen causally depends on whether or not the zebra is available. 'MULE = 1 – AVAILABLE' means that whether or not the painted mule is in the pen causally depends on the negation of whether or not the zebra is available. If the zebra is available, the painted mule will not be in the pen. Otherwise, the painted mule is in the pen. 'EXPERIENCE = max(ZEBRA, MULE)' means that whether or not John has the visual experience of seeing a zebra in the pen causally depends on whether or not either the

190

zebra or the painted mule is in the pen. John will not have the visual experience of seeing a zebra in the pen if and only if neither a zebra nor a painted mule is in the pen.

Note that there is no structural equation for SHORTAGE. The value of SHORTAGE is not determined causally by the values of any other variable in the causal model, meaning that the cause of SHORTAGE is not specified in the model.

We can thus distinguish two types of variables: *exogenous* variables whose values are not specified by any structural equation in the model and *endogenous* variables whose values are specified by exactly one structural equation in the model. In Z, SHORTAGE is exogenous, while AVAILABLE, ZEBRA, MULE, and EXPERIENCE are endogenous. From a theoretical point of view, the values of exogenous are simply given to a causal model; it is presupposed in the model.

The third element of a causal model, *A*, is a function that assigns values to all variables in the model.[80] More precisely, for any $V_i \in V$, $A(V_i)$ assigns a value to V_i, which is consistent with the set of structural equations *S*. For instance, in Z, the value assignment is as follows:

$$A(\text{SHORTAGE}) = A(\text{MULE}) = 0, \text{ and}$$

$$A(\text{AVAILABLE}) = A(\text{ZEBRA}) = A(\text{EXPERIENCE}) = 1.$$

In words, in *Zoo*, there is no shortage of zebras around the country, the zebra is available to the zoo, there is no painted mule in the pen, there is a zebra in the pen, and John has the visual experience of seeing a zebra in the pen.

---

[80] Pearl's original formulation of the causal modeling semantics does not incorporate a value assignment as presented here. Here, I follow Hiddleston (2005) and Briggs (2012). One merit of having the value assignment is that it helps to make clear one dimension of the difference between counterfactual and actual situations. As noted in Section 6.10, (the causal model of) a counterfactual situation may be generated from altering the value assignment of (the causal model of) the actual situation.

It is common to use a *directed acyclic graph (DAG)* to represent the variables and the structural equations of a causal model. A DAG consists of a set of nodes, which stand for the variables, and a set of directed acyclic arrows, which capture (among other things) the parental relationship specified by the structural equations. If V_i is the parent of V_j (V_j is the child of V_i), then there is an arrow pointing from the former to the latter. Z's DAG is presented by Figure 6.1.

A DAG is a useful tool to visualize the relationships between variables. On Figure 6.1, we can easily see that SHORTAGE is the parent of AVAILABLE, that AVAILABLE is the parent of both ZEBRA and MULE, and that both ZEBRA and MULE are EXPERIENCE's parent. Following Richard Briggs (2012, 144), we can further define the *ancestor* relation: If one can get from V_i to V_j by following a series of arrows from tail to tip, then V_i is an ancestor of V_j. Conversely, if one can get from V_i to V_j by following a series of arrows from tip to tail, then V_i is an *descendant* of V_j. For any two different variables V_i and V_j, V_i is a *nondescendant* of V_j if and only if V_i is not a descendant of V_j. A variable is not a descendant, nor a nondescendant, of itself (*ibid.*). Notice that while a DAG is convenient in identifying the causal relationship (or kinship) between variables, it does not capture all the details about the relationship (the arrow does not specify the function f_i in play in the structural equations); one cannot figure out the structural equations just by looking at a DAG (cf. Pearl 2000).

A causal model captures the core of our mental representation (or understanding) of a particular circumstance. On the one hand, the set of structural equations S captures (our understanding of) the stable functional relationships, or, in Pearl's term, 'mechanism' (Pearl 2002, 96–7) in the circumstance. In *Zoo*, the events that there is no

shortage of zebras around the country, that the zebra is available, that there is no painted

mule in the pen, that there is a zebra in the pen, and that John has the visual experience of

seeing a zebra in the pen are all connected by certain mechanisms specified by the

structural equations. For instance, if there is a shortage of zebras around the country,

there will be no zebra available (by ZEBRA = AVAILABLE), or if neither a zebra nor a

pained mule is in the pen, John will not have the visual experience of seeing a zebra in

the pen, and so on (by EXPERIENCE = max(ZEBRA, MULE)).   On the other hand, the

assignment function A captures what actually happens in the circumstance. In *Zoo*, for

example, the assignment function specifies that there is no shortage of zebras around the

country (i.e., A(SHORTAGE) = 0), that there is a zebra in the pen (i.e., A(ZEBRA) = 1),

etc.
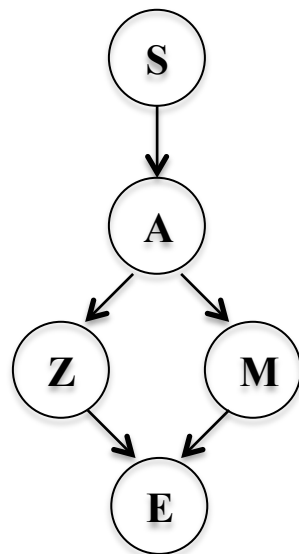
Figure 6.1
DAG of Z
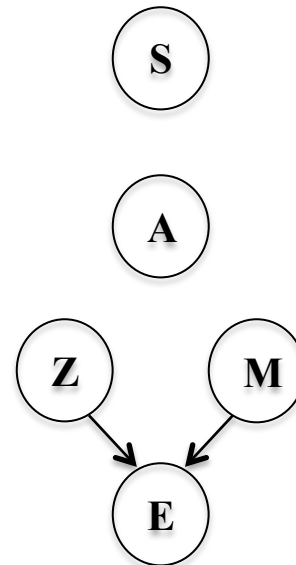
Figure 6.2
DAG of
$Z_{(AVAILABLE=0\&MULE=0)}$

With the notion of causal model specified, we are in a position to define the

notion of *intervention* as follows:

Let M ($=<V, S, A>$) be a causal model, B be a sentence of the form $V_1 = c_1 \wedge ... \wedge V_m = c_m,$[81] an *intervention on M with respect to B* generates a submodel $M_{B=1}$ ($= <V_{B=1}, S_{B=1}, A_{B=1}>$) of M such that:

(i) $V = V_{B=1}$.

(ii) $S_{B=1}$ is identical to S except that for each $V_i \in V_B$, $S_{B=1}$ replaces the structural equations $V_i = f_i(PA_i)$ of S with the structural equation $V_i = c_i$, if V_i is endogenous.

(iii) $A_{B=1}$ is identical to A except that (a) for each $V_i \in V_{B=1}$, $A_{B=1}$ sets the value V_i to c_i if V_i is exogenous, and that (b) for any $V_i \in V_{B=1}$ that is endogenous, $A_{B=1}(V_i)$ assigns a value to V_i, which is consistent with the set of structural equations $S_{B=1}$.

In words, an intervention on a model M with respect to B (B is the conjunction $V_1 = c_1 \wedge ... \wedge V_m = c_m$) is to replace the original structural equation of V_i (V_i belongs to the conjuncts of B) with the new structural equation $V_i = c_i$, if V_i is an endogenous variable. If V_i is an exogenous variables, an intervention simply set the value of when V_i to be c_i. The result of an intervention is a new causal model, which will be called a submodel of the original causal model. The value assignment of the submodel will assign values to the variables in ways that are consistent with the new structural equations (if V_j is a nondescendant of V_i, where V_i belong to the conjuncts of B, the value of V_j remains intact).

Given that a causal model M represents an *actual* situation, an intervention on M with respect to p generates a submodel M_p, which represents the *counterfactual* situation in which p is true.

To illustrate, recall that the model Z consists of a set of variables *V*:

---

[81] Pearl's (Galles and Pearl 1998) semantics deals only with a subset of sentences of the form 'A > C'. In particular, he consider conditional "A > C" of the form '(A_1 & ... & A_n) > (C_1 & ... & C_n)' while A_i and C_j have the form '$A_i = V_i$' and '$C_j = V_j$' respectively. Halpern (2000) has developed a semantics for "A > C" with A taking the form 'A_1 & ... & A_n' (like Pearl's), while C beng any Boolean combination of sentences of the form '$C_i = V_i$'. Briggs (2012) has further extended the semantics to deal with "A > C" with A to be any Boolean combination of sentences of the form '$A_i = V_i$'. For simplicity's sake, I will here focus on a language with less expressive power. That is, I will follow Pearl in assuming that the sentences involved in intervention consist only of conjunctions.

{SHORTAGE, AVAILABLE, ZEBRA, MULE, EXPERIENCE},

a set of structural equations S:

AVAILABLE = SHORTAGE

MULE = 1 – AVAILABLE

ZEBRA = AVAILABLE

EXPERIENCE = max(ZEBRA, MULE).

and an assignment function A as follows:

A(SHORTAGE) = A(MULE) = 0, and

A(AVAILABLE) = A(ZEBRA) = A(EXPERIENCE) = 1.

Suppose that we intervene in Z with respect to (AVAILABLE = 0 & MULE = 0). This intervention generates the submodel $Z_{(AVAILABLE=0\&MULE=0)}$ which has the same set of variables as Z, but its set of structural equations $S_{(AVAILABLE=0\&MULE=0)}$ is as follows:

AVAILABLE = 0

MULE = 0

ZEBRA = AVAILABLE

EXPERIENCE = max(ZEBRA, MULE).

Notice that the old mechanisms (AVAIALBLE = SHORTAGE) and (MULE = 1 – AVAILABLE) are both replaced by each conjunct of (AVAILABLE = 0 & MULE = 0). Moreover, the value assignment $A_{(AVAIALBE=0\&MULE=0)}$ is as follows:

$A_{(AVAILABLE=0\&MULE=0)}$(SHORTAGE) = $A_{(AVAILABLE=0\&MULE=0)}$(AVAILABLE) = $A_{(AVAILABLE=0\&MULE=0)}$(MULE) = $A_{(AVAIALBE=0\&MULE=0)}$(ZEBRA) = $A_{(AVAILABLE=0\&MULE=0)}$(EXPERIENCE) = 0.

Since SHORTAGE is a nondescendant of AVAILABLE (and MULE), its value should remain intact, i.e., A(SHORTAGE) = $A_{(\text{AVAILABLE}=0\&\text{MULE}=0)}$(SHORTAGE) = 0. The values of the rest of the variables are calculated in terms of S_B: Since AVAILABLE = 0 and ZEBRA = AVAILABLE, we get ZEBRA = 0. Since MULE = 0, ZEBRA = 0, and EXPERIENCE = max(ZEBRA, MULE), we deduce EXPERIENCE = 0.

If we draw a DAG for $Z_{(\text{AVIALABLE}=0\&\text{MULE}=0)}$, we get Figure 6.2. Comparing Figure 6.1 with Figure 6.2, we can visualize that the main function of intervention is to 'mutilate' (cf. Pearl 2002) the arrow in the original DAG such that some nodes cease to be connected to others—thereby altering the mechanisms of the original model. The visual effect for such a mutilation is that it changes the ancestral relationships between variables. For instance, by mutilating the arrow from SHORTAGE to AVAILABLE, whether the zebra is available no longer depends on whether there is no shortage of zebra in zoos around the country in $Z_{(\text{AVAILABLE}=0\&\text{MULE}=0)}$. An intervention also allows (though not necessarily) a different assignment of the value of the variables in $Z_{(\text{AVAILABLE}=0\&\text{MULE}=0)}$. According to $A_{(\text{AVAILABLE}=0\&\text{MULE}=0)}$, the value of AVAILABLE is set to be 0 (while according to A, AVAILABLE's value is 1).

On the causal modeling semantics, the truth condition of counterfactual conditionals can be defined in terms of intervention or submodel. Like the possible-worlds semantics, causal modeling semantics also adopts the standard form:

(S) "A > C" is true at a point a if and only if C is true at a selected set of points b in which A is true (more precisely, for all b $\in$ f(A, a), C is true in b).

While the possible-worlds semantics takes points to be possible worlds, causal modeling semantics regards them as causal models. While the possible-worlds semantics

characterizes the selection function *f* as the comparative similarity between worlds, causal

modeling semantics characterizes *f* as intervention, as follows:

> (CM) "A > C" is true in a causal model M if and only if C is true in the submodel M_A.

In words, to determine whether or not "If A had obtained, then C would have obtained"

(in symbol 'A > C') is true in M, we first generate the submodel M_A by intervening in M

with respect to the antecedent A, and then see whether or not the consequent C is true in

M_A.

To illustrate, consider *Zoo* again. In *Zoo*, the zebra is available and there is a

genuine zebra in the pen. Now, intuitively, "If the zebra were not available, there would

not have been a zebra in the pen" is true in *Zoo*. Can (CM) give the correct verdict? To

answer this question, we first construct a causal model Z to represent *Zoo*, and then

intervene in Z with respect to the antecedent *The zebra were not available* (in symbol

'AVAILABLE = 0'), to generate the submodel $Z_{AVAILABLE=0}$. After that, we determine

whether or not the consequence *There would not have been a zebra in the pen* (in symbol

'ZEBRA = 0') is true in $Z_{AVAILABLE=0}$.

Now, $Z_{AVAILABLE=0}$ consists of the set of variables *V* (={SHORTAGE,

AVAILABLE, MULE, ZEBRA, EXPERIENCE}), the set of structural equations

$S_{AVAILABLE=0}$ consists of (see Figure 6.3):

> AVAILABLE = 0
>
> MULE = 1 – AVAILABLE
>
> ZEBRA = AVAILABLE
>
> EXPERIENCE = max(ZEBRA, MULE).

The value assignment $A_{AVAILABLE=0}$ is as follows:

$A_{\text{AVAILABLE}=0}(\text{SHORTAGE}) = A_{\text{AVAILABLE}=0}(\text{AVAILABLE}) = A_{\text{AVAILABLE}=0}(\text{ZEBRA}) = 0$, and

$A_{\text{AVAILABLE}=0}(\text{MULE}) = A_{\text{AVAILABLE}=0}(\text{EXPERIENCE}) = 1$.

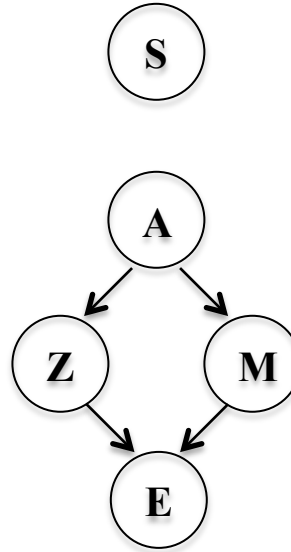Since ZEBRA = 0 is true in $Z_{\text{AVAILABLE}=0}$, (CM) predicts that "AVAILABLE = 0 > ZEBRA = 0" is true in Z, as desired.

Figure 6.3
DAG of
$Z_{\text{AVAILABLE}=0}$

It is worth noting that (CM) is capable of handling the aforementioned troubling cases against (PW). Consider, for instance, *Bomb*. (CM) predicts that "If the beetle had pushed the button, the universe would have been destroyed" is true in *Bomb*. Let us construct the causal model B that correctly represents our mental representation (understanding) of *Bomb*. Intuitively, B consists of the set of variables:

EXOGENOUS [82] represents (unspecified) indeterministic factors that cause PUSH to have the value 1.

PUSH represents whether the beetle pushes the button or not.

---

[82] In what follows, I will use 'EXOGENOUS' to represent the exogenous variables of a causal model.

SIGNAL represents whether a signal runs along a wire and out of the box or not.

DETONATE represents whether the mega-bomb is detonated or not.

BOX represents whether the black box and all of its contents are destroyed or not.

DESTROY represents whether the universe is destroyed or not.

As stipulated by *Bomb*, whether a signal has run along a wire and out of the box or not casually depends on whether the beetle pushes the button or not, whether the mega-bomb has been detonated or not causally depends on whether a signal has run along a wire and out of the box or not, whether the universe is destroyed or not causally depends on whether the mega-bomb has been detonated or not, and whether or not the black box and all of its contents are destroy or not causally depends on whether the mega-bomb has not been detonated. Hence, one may characterize the set of structural equations *S* of B as follows (see Figure 6.4):

PUSH = EXOGENOUS

SIGNAL = PUSH

DETONATE = SIGNAL

BOX = 1 – WIRE

DESTROY = DETONATE.

It should be noted that, by stipulation, whether the beetle will push the button or not is totally random. As a result, the value of EXOGENOUS may be assignment by a probability function. However, the rest of the mechanisms (i.e., SIGNAL, DETONATE, BOX, and DESTROY) are deterministic—e.g., if a signal has run along a wire and out of the box, the mega-bomb will be detonated. If the bomb is (not) detonated, the universe

199

(the black box and all of its contents) will be destroyed (to stipulate in this way simplifies the discussion below. But nothing significant hinges on this stipulation).
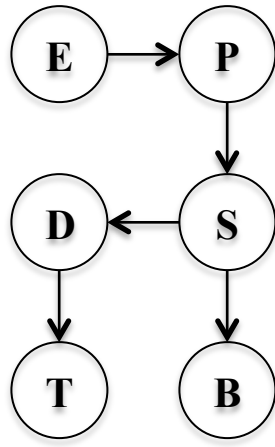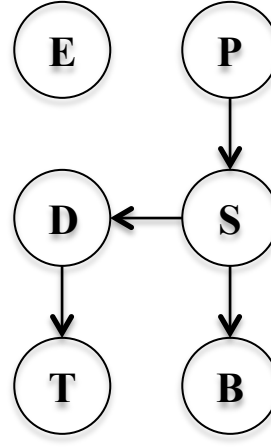

Figure 6.4
DAG of B

Figure 6.5
DAG of
$B_{PUSH=1}$

B's value assignment A is:

$A(EXOGENOUS) = A(PUSH) = A(SIGNAL) = A(DETONATE) = A(DESTROY) = 0$, and

$A(BOX) = 1$.

In words, in *Bomb*, the beetle does not push the button, there is no signal running along a wire and out of the box, the mega-bomb is not detonated, the universe is not destroyed, and the black box and all of its contents are destroyed.

Is "If the beetle had pushed the button, the universe would have been destroyed" (in symbol "PUSH = 1 > DESTROY = 1") true in B? To answer this question, according to (CM), we need to first intervene in B with respect to PUSH = 1.

To intervene in B with respect to PUSH = 1 generates the submodel $B_{PUSH=1}$ whose set of structural equations $S_{PUSH=1}$ consists of   (see Figure 6.5):

PUSH = 1

SIGNAL = PUSH

DETONATE = SIGNAL

BOX = 1 – SIGNAL

DESTROY = DETONATE.

The value assignment $A_{\text{PUSH}=1}$ is as follows:

$A_{\text{PUSH}=1}(\text{EXOGENOUS}) = A_{\text{PUSH}=1}(\text{DESTROY}) = 0$, and

$A_{\text{PUSH}=1}(\text{PUSH}) = A_{\text{PUSH}=1}(\text{SIGNAL}) = A_{\text{PUSH}=1}(\text{DETONATE}) = A_{\text{PUSH}=1}(\text{DESTROY}) = 1$.

As desired, "PUSH = 1 > DESTROY = 1" is true in B since DESTROY = 1 is true in $B_{\text{PUSH}=1}$.

The same maneuver can be used to account for cases such as *Bet* and *Bet** (I will let the readers to work out the details).

So far, so good. But acute readers may have already noticed that (CM), like (PW), also fails to take into account backtracking conditional. More precisely, (CM) is unable to respect the two kinds of counterfactualization manifested in *Ask**. To elaborate, let us construct an appropriate causal model K for *Ask**. Intuitively, K consists of the following events:

EXOGENOUS-**1** represents whether Jack is mad at Jim or not.

EXOGENOUS-**2** represents whether Jim is a prideful fellow or not.

ASK represents whether Jim asks Jack for help or not.

REJECT represents whether Jack will reject Jim or not.

According to *Ask**, whether Jim will ask Jack for help causally depends both on whether

201

Jack is mad at Jim and on whether Jim is a prideful fellow. If Jim is a prideful fellow and Jack is mad at him, Jim will not ask Jack for help. For the sake of discussion, we may suppose that Jim will ask Jack for help if either Jim is not prideful or Jack is not mad at him. Whether Jack will deny Jim or not causally depends on whether Jack is mad and whether Jim asks for help. If Jack is mad and Jim asks for help, Jack will reject Jim. We may also suppose that Jack will not reject Jim if either Jack is not mad or Jim does not ask for help. In sum, the structural equations S of K can be characterized as follows (see Figure 6.6):

ASK =   1 if EXOGENOUS-1 = 0 or EXOGENOUS-2 = 0

            0 if otherwise.

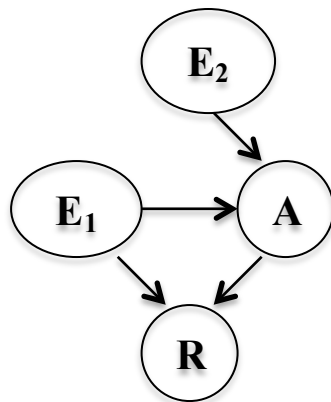REJECT =   1 if EXOGENOUS-1 = ASK = 1
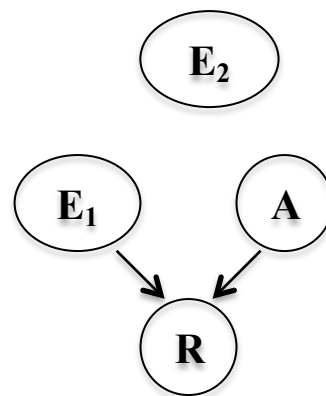
            0 if otherwise.

Figure 6.6
DAG of K

Figure 6.6
DAG of $K_{ASK=1}$

In *Ask**, Jack is indeed mad at Jim, Jim is a prideful fellow, Jim does not ask for help, and Jack does not reject Jim. The value assignment A of K is thus:

$A(\text{ASK}) = A(\text{REJECT}) = 0$, and

$A(\text{EXOGENOUS-1}) = A(\text{EXOGENOUS-2}) = 1$.

Now, if Jim were to ask Jack for help, would Jack reject him or not? To see how (CM) will answer this question, we need to first intervene in K with respect to ASK = 1. $K_{\text{ASK}=1}$'s set of structural equations $S_{\text{ASK}=1}$ consists of:

ASK = 1

REJECT =   1 if EXOGENOUS-1 = ASK = 1

              0 if otherwise.

The assignment $A_{\text{ASK}=1}$ is as follows:

$A_{\text{ASK}=1}(\text{EXOGENOUS-1}) = A_{\text{ASK}=1}(\text{EXOGENOUS-2}) = A_{\text{ASK}=1}(\text{ASK}) = A_{\text{ASK}=1}(\text{REJECT}) = 1$.

Hence, "ASK = 1 > REJECT = 1" is true in K. In words, on (CM), "If Jim were to ask Jack for help, Jack would have denied him" is true in *Ask**. As noted above, this is not an incorrect or implausible verdict. If one focuses on the fact that Jack is mad at Jim, it surely makes sense that to say that Jim would have been rejected if he were to ask Jack for help. However, this just shows that (CM) is insensitive to the alternative counterfactualization available in *Ask**. That is, if one focuses on the fact that Jim is a prideful fellow, who will never risk to be denied, it also makes sense to claim that "If Jim were to ask for help, he will not be rejected" is true in *Ask**, since if Jack was to be mad at him in the first place, Jim would not even ask for help.

In sum, *Ask** manifests two kinds of counterfactualization; yet (CM), like (PW), accounts for one of them but for not the other. As a semantics of counterfactual conditionals, both are thus incomplete. Yet, as I see it, the causal modeling semantics has an edge over the possible-worlds semantics in that, while the former has the resources to

explain these two kinds of counterfactualization, the same could hardly be said of the latter.

6.10 Interventional vs. Inferential Counterfactualization

According to (CM), "A > C" is true in M if, and only if, C is true in M_A which is generated by intervening in M with respect to A. Proponents of the causal modeling semantics have claimed that (CM) captures our intuitions regarding counterfactual conditionals in general (cf. Pearl 2000).[83] This claim, however, is undermined by cases like *Ask** which indicate that, while (CM) surely has captured a legitimate sense of counterfactualization, an alternative sense of counterfactualization is still at large. Fortunately, the causal modeling semantics has the resources necessary to explain the other kind of counterfactualization.

The causal model representing a particular situation enables us to extract certain important information about the situation. More precisely, it allows us to acquire what I will call the *interventional* as well as *inferential* information of the situation. The acquisition of the former consists in intervening in the causal model with respect to a certain proposition p. More precisely, an intervention on M with respect to p generates a submodel M_p which contains information about, given that a circumstance is presented by the causal model, what the circumstance would be if p were brought to obtain by manipulation (intervention). As shown in the aforementioned cases, interventional information is often what we want to know when using counterfactual conditionals. For

---

[83] Judea Pearl simply *defines* counterfactuals in terms of (CM). This is not totally dogmatic, since psychological studies have shown that (CM) has correctly predicted some of our ordinary counterfactual reasoning (cf. Sloman and Lagnado 2005; also see Rips 2010). The problem of (CM), rather, is that it does not exhaust the ordinary notion of counterfactualization.

instance, in *Bomb*, the question whether "PUSH > DESTROY" (i.e. PUSH = 1 > DESTROY = 1) is true in *Bomb* is actually a question whether DESTROY would obtain if PUSH was manipulated to have the value 1. The question thus can be answered by acquiring interventional information about *Bomb*.

Not only can we acquire interventional information, we can also acquire what I call 'inferential information' from a causal model. To explain, it is useful to introduce the distinction between what Judea Pearl and David Galles call *seeing* and *doing* (1998, 159) (note: these are Pearl and Galles's terminologies). As noted, a causal model contains a set of structural equations specifying the causal relationships among the events in the represented circumstance. Consider *Bomb* whose set of structural equations consists of (see Figure 6.4):

PUSH = EXOGENOUS

SIGNAL = PUSH

DETONATE = SIGNAL

BOX = 1 – DETONATE

DESTROY = DETONATE.

Suppose that you *see* that the universe is not destroyed (i.e., DESTROY = 0). This observation allows you to infer that the beetle did not push the button (i.e., PUSH = 0): From DESTROY = DETONATE and DESTROY = 0, we deduce DETONATE = 0. By DETONATE = 0 and DETONATE = SIGNAL, we deduce that SIGNAL = 0. By SIGNAL = 0 and SIGNAL = PUSH, we deduce that PUSH = 0.

By contrast, suppose that regardless of what happens in the black box, someone detonated the mega-bomb (i.e., DETONATE is manipulated to have the value 1). The

universe would thus be destroyed (i.e., DESTROY = 1) even though the beetle did not push the button. (This is so since BOX = 1 – DETONATE and DESTROY = DETONATE. If DETONATE = 1, it follows that DESTROY = 1 as well.)

The former is *seeing* while the latter is *doing*. Seeing does not intervene in the circumstance: We first observe some events of the circumstance, and then infer what happens for the other events via structural equations. Doing, by contrast, intervenes in the circumstance: We first manipulate some events of the circumstance, and then determine what will follow via structural equations. Put in another way, seeing involves finding out (inferring) the value of a certain variable via the set of *original* equations, while doing, replacing the original equations with *new* ones.

Intervention is a kind of doing. Essentially, an intervention on M with respect to p is an action of replacing the structural equation $p = f_i(PA_i)$ with $p = 1$; this captures the notion of manipulation (i.e., *doing*). Interventional information about M is essentially about what M would be like given a certain manipulation (intervention).

If we replace the value assignment of a causal model while leaving its structural equations intact, it is an act of *seeing*, not doing. For instance, in *Bomb* (i.e., B), the beetle does not push the button (i.e., PUSH = 0), the bomb is not detonated (i.e., DETONATE = 0), and the universe is not destroyed (i.e., DESTROY = 0). Suppose that we observe that the bomb is detonated (i.e., DETONATE = 1). We can thus infer that the beetle does push the button (by DETONATE = SIGNAL and SIGNAL = PUSH) and that the universe is destroyed (by DESTROY = DETONATE). Unlike intervention, however, our inference does not consist in changing the original structural equations. Now, if we

206

perform an act of seeing (observation) on a causal model M in this way, we acquire *inferential* information about M.

Corresponding to the distinction between interventional and inferential information is the distinction between interventional and inferential counterfactualization. Roughly, counterfactualization is about what it would be like if certain counterfactual events happened. When actual/counterfactual situations[84] are represented by causal models, counterfactualization consists in two conditions: The alteration of the value assignment of some events of the actual situation and/or the replacement of the structural equations. On interventional counterfactualization, viz. (CM), "A > C" is true in M if and only if C is true in M_A. A submodel represents the counterfactual circumstance different from the actual circumstance M in that the former has different structural equations (which in turn represent the stable mechanisms of the circumstance). The difference in structural equations often (but not necessarily) implies the difference in the assigned value of some variables. Inferential counterfactualization, by contrast, does not contain the replacement of structural equations, but only the alteration of the value assignment.

The distinction between inferential and interventional counterfactualization is the key to understanding the two kinds of counterfactualization manifested by cases like *Ask**. As mentioned above, there are two legitimate ways of counterfactualizing what would have happened if Jim had asked Jack for help (i.e., ASK = 1). On the one hand, one may reason, since Jack was mad at Jim, Jim would have been rejected if he had asked (i.e., ASK = 1 > REJECT = 1). This is, as should be clear now, a kind of interventional counterfactualization: We first intervene in the original circumstance (i.e., K) by

---

[84] Recall that 'actual' and 'counterfactual' are our terms of art. In particular, an 'actual' situation upon which a counterfactual situation is envisioned may not be real. See footnote 2.

supposing that Jim had asked Jack for help. We then ask whether Jim would be rejected under this supposition (i.e., $K_{ASK=1}$). On the other hand, one may reason that since Jim is a prideful fellow who would not risk being rejected, he would have asked Jim for help only if he would not be rejected (i.e., Jim is not mad at him). As a result, if Jim asked for help, it must be the case that Jack was not mad, and in this Jim would not be rejected (i.e., $ASK = 1 > REJECT = 0$). Obviously, this line of reasoning is a kind of inferential counterfactualization. We know that if Jack is mad at Jim, he will reject Jim and that Jim is a prideful fellow who will ask for help only if he will not be rejected. Now, Jim was mad at Jim and Jim had not asked Jim for help. But if Jim were to ask Jim for help, we can *infer* that Jim had not been mad at him.[85]

So far, I have been arguing that there are two kinds of counterfactualization, i.e., interventional and inferential. (CM) has captured the former but not the latter. Yet, the casual-modeling framework has resources necessary for explaining inferential counterfactualization. More importantly, while we can extract two kinds of information—interventional and inferential information—from a causal model, the same cannot be said of the possible-worlds semantics. It is unclear how the possible-world framework can deal with inferential information. Hence, the causal modeling semantics has an edge over the possible-worlds semantics in that the former, but not the latter, is in a position to respect our intuitions regarding counterfactualization.

This advantage of the causal modeling semantics over the possible-worlds semantics is particularly important to epistemologists. As I will argue in the next section, our intuitions regarding counterfactual conditionals in cases involving pseudo skeptical

---

[85] I conclude that the distinction between backtracking and forward-tracking counterfactualization is in fact a special case between the distinction between interventional and inferential counterfactualization.

hypotheses are actually driven by inferential counterfactualization *only*. Failure to distinguishing between inferential and interventional counterfactualization has led many epistemologists to mistakenly conclude that such cases are genuine counterexamples to the sensitivity condition—the fact that the possible-worlds semantics is widely adopted by epistemologists, I suspect, is responsible for this widespread mistaken conclusion. Epistemologists, thus, have an additional motivation to adopt the causal modeling semantics, which respects these two kinds of counterfactualization.

6.11 Cases Involving Pseudo Skeptical Hypotheses Reconsidered

My goal in this section is to argue for the following claims:

> (a)   Cases involving pseudo skeptical hypotheses usually involve inferential counterfactualization rather than interventional one.

> (b)   The sensitivity condition of knowledge requires only interventional counterfactualization.

(a) implies that our intuitions regarding counterfactual conditionals in cases involving pseudo skeptical hypotheses are driven by inferential counterfactualization. By contrast, according to (b), our intuitions regarding sensitivity in evaluating knowledge are driven by interventional counterfactualization. From (a) and (b), it follows that cases involving pseudo skeptical hypotheses could hardly be used to argue against the sensitivity condition of knowledge. In what follows, I will deal with (a) and (b) respectively.

6.11.1 Arguing for (a)

As the readers may have already noticed, the first example that we use to introduce the casual modeling semantics, i.e., *Zoo*, has the same structure as cases involving pseudo

skeptical hypotheses. Recall that such cases are stipulated in such a way that (i) the subject believes that p via R and that (ii) such cases contain some *pseudo* skeptical scenarios such that if p were false, the subject would still have R (see Section 5.6). In *Zoo*, John believes that there is a zebra in the pen via his visual experience of seeing a zebra in the pen and there is in fact a zebra in the pen. However, if the zebra were not in the pen, then a skeptical scenario, arguably, obtained, namely, a painted mule would be in the pen. In other words, whether the zebra is in the pen or not, John is going to have the visual experience of seeing a zebra.

Intuitively, "If there were no zebra in the pen, John would still have the experience of seeing a zebra in the pen" (i.e., "ZEBRA = 0 > EXPERIENCE = 1") is true in *Zoo*. It is important to notice that "ZEBRA = 0 > EXPERIENCE = 1" is not true in Z according to (CM); "ZEBRA = 0 > EXPERIENCE = 1" is not true under interventional counterfactualization. On (CM), "ZEBRA = 0 > EXPERIENCE = 1" is true in Z if and only if "EXPERIENCE = 1" is true in $Z_{ZEBRA=0}$. The set of structural equations of $Z_{ZEBRA=0}$ consists of (see Figure 6.8):

AVAILABLE = 1 – SHORTAGE

ZEBRA = 0

MULE = 1 – AVAILABLE

EXPERIENCE = max(ZEBRA, MULE).

Hence, the value assignment $A_{ZEBRA=0}$ is as follows:

$A_{ZEBRA=0}$(SHORTAGE) = $A_{ZEBRA=0}$(ZEBRA) = $A_{ZEBRA=0}$(MULE) = $A_{ZEBRA=0}$(EXPERIENCE) = 0, and

$A_{ZEBRA=0}$(AVAILABLE) = 1.

Hence, "ZEBRA = 0 > EXPERIENCE = 1" is false in Z. In other words, "ZEBRA = 0 > EXPERIENCE = 1" is not true under interventional counterfactualization.
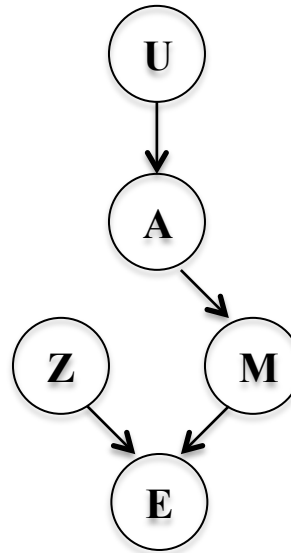

Figure 6.8
DAG of $Z_{ZEBRA=0}$

"ZEBRA = 0 > EXPERIENCE = 1", by contrast, is true under inferential counterfactualization. First, we know that the value of ZEBRA is 1 in Z. But, if the value of ZEBRA was 0, we can easily infer that the value of EXPERIENCE would still be 1: From ZEBRA = 0 and ZEBRA = AVAILABLE, we deduce that AVAILABLE = 0. From AVAILABLE = 0 and MULE = 1 − AVAILABLE, we deduce that MULE = 1. From MULE = 1 and EXPERIENCE = max(ZEBRA, MULE), we deduce that EXPERIENCE = 1.

The same can be said of other cases involving pseudo skeptical hypotheses mentioned at the end of last chapter. Consider, for instance, *Grandfather*:

> *Grandfather.* Suppose that Nancy's grandfather was born in 1945. Nancy believes that he was born in 1945 via his testimony. Suppose, for some reasons, her grandfather are obsessed to make others to believe that he was born in 1945 such that if he were not born in that year, he would still try to deceive everyone to

211

believe that he was born in 1945 (suppose he is very good at this, and there is a very high chance for him to successfully deceive everyone around him).

Intuitively, "If Nancy's grandfather were not born in 1945, Nancy would still have the testimony that he was born in 1945" seems true in *Grandfather*. What has not been realized, however, is that this counterfactual conditional is not true under interventional counterfactualization, but rather under inferential counterfactualization.

To elaborate, let us first construct a causal model G to represent *Grandfather*. If Nancy's grandfather were not born in 1945, then apparently his life would have taken a different causal path. For simplicity's sake, let us suppose Nancy's father was born right before the midnight of December the thirty-first, 1945. Suppose that if the doctor came in the delivery room late, her grandfather would have been born on January the first, 1946. The set of variables of G thus consists of:

> **E**XOGENOUS represents whether the doctor does not come in the delivery room late or not.
>
> 19**45** represents whether Nancy's grandfather is born in 1945 or not.
>
> 19**46** represents whether Nancy's grandfather is born in 1946 or not.
>
> **D**ECEIVE represents whether Nancy's grandfather has tried to deceive her to believe that he was born in 1945 or not.
>
> **T**ESTIMONY represents whether Nancy has her grandfather's testimony that he was born in 1945 or not.

On *Grandfather*, whether Nancy's grandfather has tried to deceive her believe that he was born in 1945 is determined by whether he is born in 1946—he will not deceive her unless he is born in 1946. Moreover, whether Nancy will have the same testimony that her grandfather was born in 1945 is determined both by whether he is born in 1945 and by whether he has tried to deceive everyone to believe that he was born in 1945. If he is

born in 1945, he will have told her so and Nancy will have the testimony that he was born

in 1945. If he is born in 1946, he will deceive her to believe that he was born in 1945. In

this case, Nancy will still have the testimony that he was born in 1945. G's set of

structural equations S, hence, consists of the following (see Figure 6.9):

1945 = EXOGENOUS

1946 = 1 – EXOGENOUS
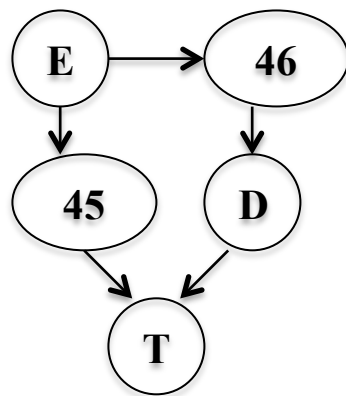
DECEIVE = 1946

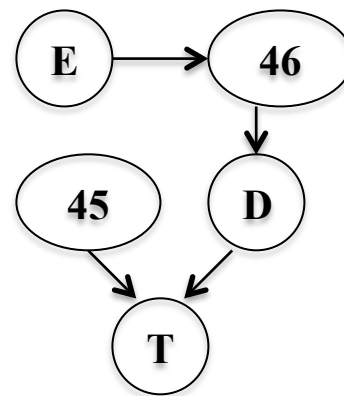TESTIMONY = max(1945, DECEIVE).

Figure 6.9
DAG of G

Figure 6.10
DAG of $G_{1945=0}$

In *Grandfather*, the doctor did not come into the delivery room late, Nancy's

grandfather was in fact born in 1945, he did not try to deceive Nancy to believe that he

was born in 1945, and Nancy has her grandfather testimony that he was born in 1945.

The value assignment A of G, hence, is as follows:

$A(1946) = A(DECEIVE) = 0$, and

$A(EXOGENOUS) = A(1945) = A(TESTIMONY) = 1$.

If Nancy's grandfather were not born in 1945, would she still have the testimony that he

was born in 1945, under interventional counterfactualization? In other words, is "1945 = 0 > TESTIMONY = 1" true in G according to (CM)?

Let us intervene in G with respect to 1945 = 0. The set of structural equations of $S_{1945=0}$ consists of (see Figure 6.10):

$$1945 = 0$$

$$1946 = 1 - \text{EXOGENOUS}$$

$$\text{DECEIVE} = 1946$$

$$\text{TESTIMONY} = \max(1945, \text{DECEIVE}).$$

The value assignment $A_{1945=0}$ is as follows:

$$A_{1945=0}(1945) = A_{1945=0}(1946) = A_{1945=0}(\text{TESTIMONY}) = A_{1945=0}(\text{DECEIVE}) = 0, \text{ and}$$

$$A_{1945=0}(\text{EXOGENOUS}) = 1.$$

As a result, "1945 = 0 > TESTIMONY = 1" is false in G. Therefore, "1945 = 0 > TESTIMONY = 1" is not true under interventional counterfactualization.

"1945 = 0 > TESTIMONY = 1", however, is true under inferential counterfactualization. That is, suppose that the value of 1945 is 0, we can easily infer that the value of TESTIMONY is 1: From 1945 = 0 and 1945 = EXOGENOUS, we deduce that EXOGENOUS = 0. From EXOGENOUS = 0 and 1946 = 1 – EXOGENOUS, we deduce that 1946 = 1. From 1946 = 1 and DECEIVE = 1946, we deduce that DECEIVE = 1. From 1945 = 0, DECEIVE = 1, and TESTIMONY = max(1945, DECEIVE), we deduce that TESTIMONY = 1.

To conclude, my contention is that our intuitions regarding "If p were false, the subject would still have the reason R for p" in cases involving pseudo skeptical

hypotheses are generally driven by inferential counterfactualization (rather than by interventional counterfactualization), i.e., (a) obtains.

(a) has an important impact on the study of knowledge, specifically, on the sensitivity condition. As noted, the fact that we generally find "1945 = 0 > TESTIMONY = 1" intuitive has been used to argue that Nancy's testimony that her grandfather was born in 1945 is *in*sensitive. This, however, is not the correct conclusion to be drawn, since "1945 = 0 > TESTIMONY = 1" is true only under *inferential* counterfactualization (by (a)), and, as I will argue, the sensitivity condition required by knowledge is related only to *interventional* counterfactualization, i.e., (b). This last point can be shown by a comparison between inferential and interventional counterfactualization.

6.11.2 Arguing for (b)

The main distinction between interventional counterfactualization and inferential one is manifested by the notion of counterfactual dependence. Call them *intervention-* and *inference-counterfactual dependence* respectively. An event C is intervention-counterfactually dependent on an event A if and only if C is causally dependent on A. In other words, C is intervention-counterfactually dependent on A if and only if A is the causal ancestor of C (i.e., C is at the causal downstream of A). Viewed in the light of DAG, there is a series of arrows running from A to C, while the direction of the arrows must be from tails to tips. Intervention-counterfactual dependence guarantees that the occurrence of A is a direct or indirect cause of the occurrence of B.

By contrast, an event C is inference-counterfactually dependent on an event A if and only if either C is causally dependent on A *or* C is causally dependent on some D,

which is a common cause of C and A. In other words, C is inference-counterfactually dependent on A if and only if either A is a causal ancestor of C (i.e., C is at the causal downstream of A) or, for some D, D is a causal ancestor of both C and A (i.e., C and A are at the causal downstream of D; in this case, C and A may be causally independent of each other). Viewed in the light of DAG, there is a series of arrows running from A to C, while the direction of the arrows does not have to be from tails to tips. Inference-counterfactual dependence does not guarantee that the occurrence of A is a direct or indirect cause of the occurrence of B.

The difference between intervention- and inference-counterfactual dependence gives us the principled reason to take the sensitivity condition to satisfy the former, but not the latter. Recall that the sensitivity condition is supposed to capture the intuitive notion of tracking crucial to the notion of knowledge—we want S's (sole) reason R for p to track the truth of p if S is to know that p at all. The tracking relation at issue requires that if A tracks the truth of B, then a significant alteration of B should also significantly alter A. For instance, suppose that someone walked along the sand and left a series of footsteps. Intuitively, the footprints on the sand tracks the truth of *A human walked along the sand*. It seems that a significant alteration of *A human walked along the sand* will have an effect on the footprints on the sand as well. If this person were prevented to walk along the sand, no footprints would be left on the sand. If a cat, as opposed to a human, walked along the sand, the footprints would have been very different. We tend to think that S knows that p when S's reason for p tracks the truth of p in this sense. For instance, suppose that S's sole reason for *A human walked along the sand* is that S's visual experience of the series of footprints left on the sand. Intuitively, it seems that S knows

that someone walked along the sand. By contrast, suppose that S guesses (without seeing the footprints) that a human walked along the sand. It seems that S's lucky guess does not track the truth of *A human walked along the sand*. If the person were prevented to walk along the sand, S would still guess that a human walked along the sand. If a cat, as opposed to a human, walked along the sand, S would still guess that a human walked along the sand. Intuitively, we think that a lucky guess does not give rise to knowledge.

Inference-counterfactual dependence fails to capture this sense of tracking. That is, if A is inference-counterfactually dependent on B, it is not necessary that A tracks the truth of B, namely, it is not necessary that a significant alteration of B will also significantly alter A. To elaborate, notice that if C is inference-counterfactually dependent on A, then either C is causally dependent on A, or C and A are causally dependent on D (A and C may not be causally dependent on each other). However, if A and C are causally dependent on D (while A and C are not causally dependent on each other), then it seems that the tracking relation does not hold between A and C—the significant alteration of A may not alter C; nor will the significant alteration of C always alter A. Consider the following case.

> *Smoker.* Suppose that the fact that S is smoker causes S to be a coffee drinker and have lung cancer (that S is a coffee drinker and that he has lung cancer are not caused by any other event). Moreover, being a coffee drinker, in general, is not causally dependent on being a patient of lung cancer, nor the other way around.

On *Smoker*, S being a smoker is a common cause of S being a coffee drinker and S having lung cancer. According to the definition of inference-counterfactual dependence, whether S is a coffee drinker or not is inference-counterfactually dependent on whether S has lung cancer or not, and whether S has lung cancer or not is also inference-counterfactually dependent on whether S is a coffee drinker or not. However, *S is a coffee*

217

drinker does not tracks the truth of *S has lung cancer*; a significant alteration of *S has lung cancer* may not alter *S is a coffee drinker*. Suppose we cure S's lung cancer. In this case, S might still be a coffee drinker. Likewise, *S has lung cancer* does not track the truth of *S is a coffee drinker*. Suppose we force S to stop being a coffee drinker. In this case, S might still have lung cancer.

Intervention-counterfactual dependence, by contrast, captures the tracking relation required by knowledge. That is, if A is intervention-counterfactually dependent on B, it is necessary that A tracks the truth of B, namely, it is necessary that a significant alteration of B will significantly alter A. Consider *Smoker* again. S being a smoker is a common cause of S being a coffee drinker and S having lung cancer. According to the definition of intervention-counterfactual dependence, whether S is a coffer drinker or not is intervention-counterfactually dependent on whether S is a smoker or not. By the same token, whether S has lung cancer or not is intervention-counterfactually dependent on whether S is a smoker or not. What is important is that *S is a coffee drinker* tracks the truth of *S is a smoker*; a significant alteration of *S is a smoker* will significantly alter *S is a coffee drinker*. Suppose that S had not been a smoker in the first place. In this case, S would not be a coffee drinker, since, according to *Smoker*, no other event would cause S to be a coffee drinker. Likewise, *S has lung cancer* tracks the truth of *S is a smoker*; a significant alteration of *S is a smoker* will significantly alter *S has lung cancer*. Suppose that S had not been a smoker in the first place. In this case, S would not have lung cancer, since, according to *Smoker*, no other event would cause S to have lung cancer.

In sum, intervention-counterfactual dependence captures the tracking relation required by knowledge, while inference-counterfactual dependence falls short of doing so.

As a result, if the sensitivity condition is supposed to capture the intuitive idea of tracking, which is crucial to knowledge, it must be characterized in terms of interventional counterfactualization as opposed to inferential counterfactualization, i.e., (b) obtains.

Given that (a) cases involving pseudo skeptical hypotheses usually involve inferential counterfactualization rather than interventional one, and that (b) the sensitivity condition of knowledge requires only interventional counterfactualization, it should be clear that cases involving pseudo skeptical hypotheses do not constitute genuine counterexamples to the sensitivity account.

In fact, to interpret the sensitivity condition in terms of interventional counterfactualization gives rise to the correct verdict that the subject has knowledge in cases involving pseudo skeptical hypotheses. For instance, in *Zoo*, We intuitively think that John knows that a zebra is in the pen. In fact, John's reason for *A zebra is in the pen*—John's visual experience that a zebra is in the pen—is sensitive according to (CM): If the zebra had not been in the pen (i.e., ZEBRA = 0), John would lack the visual experience that of seeing a zebra in the pen (i.e., EXPERIENCE = 0), since "ZEBRA = 0 > EXPERIENCE = 0" is true in Z by (CM) (see Section 6.7).

By the same token, we intuitively think that Nancy knows that her grandfather was born in 1945 in *Grandfather*. In fact, Nancy's reason for *Nancy's grandfather was born in 1945*—Nancy's grandfather's testimony that he was born in 1945—is sensitive according to (CM): If her grandfather were not born in 1945 (i.e., 1945 = 0), she would lack the testimony that he was born in 1945 (i.e., TESTIMONY = 0), since "1945 = 0 > TESTIMONY = 0" is true in G by (CM) (see above).

6.12 Taking Stock

The possible-worlds semantics of the counterfactual conditionals has been widely adopted by epistemologists. In this chapter, I have been arguing that epistemologists are better off to adopt the causal modeling semantics. There are two reasons for this claim. First, a close scrutiny of the possible-worlds semantics reveals that it should incorporate the notion of causal dependence into its account. Since one of the main differences between the orthodox possible-worlds semantics and the causal modeling semantic is that the latter but not the former presupposes the notion of causation (causal dependence) in analyzing counterfactuals, if one is to incorporate the notion of causation into the possible-worlds semantics, one may as well just shift to the causal modeling semantics. Second, and more importantly, the causal modeling semantics allows us to appreciate the distinction between interventional and inferential counterfactualization, which is crucial in understanding the notion of sensitivity involved in the notion of knowledge.

We have seen that sensitivity is a necessary property for one's (sole) reason for p if one is to know p (Chapter 4). Traditionally, the sensitivity condition has been under severe attacks, and is widely regarded as insurmountable. I disagree. In my opinion, the sensitivity account is still alive and kicking, as long as characterized properly. In sum, the proper characterization consists in the following claims:

(1)   Sensitivity is a property of one's reason.

(2)   The sensitivity of one's reason is to be characterized in terms of counterfactual conditionals.

(3)   As far as the notion of sensitivity is concerned, the truth condition of counterfactual conditionals is defined by (CM) and interventional counterfactualization.

The majority of this chapter has been dedicated to (3). The moral of (3) is that we should be sensitive to the distinction between interventional and inferential counterfactualization in understanding the sensitivity condition. The latter is a legitimate mode of counterfactualization; just that it is not the mode of counterfactualization requires by knowledge. This fact, I suspect, has failed to be recognized by people who find cases involving pseudo skeptical hypotheses as constituting genuine counterexamples to the sensitivity condition.

Once the sensitivity condition is understood properly, i.e., via (CM), we are in a position to see how practical (pragmatic) considerations affect our evaluation of the sensitivity of one's reason. Since sensitivity constitutes one's epistemic position, to say that sensitivity is dependent on practical consideration is tantamount to saying that one's epistemic position is dependent on practical consideration. In the next chapter, I will show how these are all possible.

Chapter Seven

The Skeptical Puzzle Resolved

7.1 Introduction

In the previous chapters, I have established an account of knowledge relying on the notion of sensitivity. It is my goal in this chapter to apply this account to construct a pluralistic solution—more precisely, a standing-variantist solution—of the skeptical puzzle.

On the pluralistic solution, roughly, the same statement, say, "S knows that op (or not-sh)" is true in the ordinary context but fails to be the case in the skeptical context ('op' stands for ordinary propositions about the external world we normally take ourselves to know, while 'sh' for skeptical hypotheses that are incompatible with op). As a specific type of the pluralistic solution, the standing-variantist solution claims that S's standing on knowing dimensions—or S's epistemic position—with respect to op (or not-sh) varies with the skeptical and ordinary contexts such that it satisfies the standards for knowledge in the former but fails to do so in the latter (Section 1.7).

I have proposed to characterize one's standing on knowing dimension—more precisely, one's epistemic position—with respect to p in terms of the sensitivity of one's reason for p. Understood in this way, standing variantism is the thesis that the sensitivity of one's reason for p varies with contexts.

The evaluation of the sensitivity of S's reason is related to a certain counterfactual situation: S's reason for p is sensitive if and only if S would not have the same reason if p had not been the case. The sensitivity of one's reason thus depends on the counterfactual situation in which p failed to obtain. So far, I have assumed that there is a unique counterfactual situation in which p had not been the case. In what follows, this assumption will further be discarded. As we will see, the evaluation of a counterfactual conditional, in general, depends on the counterfactual situation *in play*, which may vary with contexts. The variability of the sensitivity of one's reason, hence, is just the corollary of the variability of the counterfactual situation in play.

This feature of variability of sensitivity gives rise to the claim that S's reason for op (and not-sh) is sensitive in the ordinary context but fails to be the case in the skeptical one. The reason, I argue, is that the skeptic and the ordinary people are engaging in different kinds of *epistemic inquiries*. To shift from the ordinary context to the skeptical one is to shift from one inquiry to another. Each inquiry mandates different counterfactual situations to be in play in evaluating the sensitivity of S's reason. This is why the sensitivity of S's reason varies with such contexts, even holding fixed S's situation.

The following consists of eleven sections. In Section 7.2, I outline my main argument for the standing-variantist solution. I call my view 'the sensitivity-variantist solution'. Section 7.3 argues that the utterance "S does not know that not-sh" is true in the skeptical context. Section 7.4 argues that "S knows that op" is true in the ordinary context. Section 7.5 argues that "S does not know that op" is true in the skeptical context and that "S knows that not-sh" is true in the ordinary context. Section 7.6 points out that

the evaluation of sensitivity depends on the inquiry the evaluator is engaging in. In Section 7.7, I argue that the skeptic and the ordinary people engage in different inquiries, which gives rise to the fact that S's reason for op (and not-sh) is true in the ordinary context but fails to be the case in the skeptical one. However, just because two parties engaging in different inquiries, it does not follow that these inquiries are legitimate. Section 7.8, I argue that in we should evaluate whether the skeptic and the ordinary people are engaging legitimate *epistemic* inquiry. In Section 7.9, I reject Wittgenstein's claim that the skeptic is engaging in an illegitimate empirical inquiry. In Section 7.10, I argue that both the skeptic and the ordinary people are engaging in (distinct) legitimate epistemic inquiries. In Section 7.11, I consider and reply to three possible objections. The last section wraps up the chapter.

7.2 The Main Argument

The tracking-based account (of knowledge) is in a position to offer a pluralistic solution to the skeptical puzzle:

> 1.1   S knows that op.
>
> 1.2   S does not know that not-sh.
>
> 1.3   S knows that op only if S knows that not-sh.

The idea, roughly, is that the skeptic is right in claiming that 1.2 and not-1.1 in the skeptical context, and ordinary people are also right in claiming that 1.1 and not-1.2 in the ordinary context, since S's reason for op (or not-sh) is sensitive in the ordinary context but fails to be the case in the skeptical context. 1.1-1.3 thus cease to constitute a genuine paradox since in no context they are true altogether.

My argument consists in four steps. First, I argue that S's reason for not-sh is insensitive, evaluated in terms of the counterfactual situation in which sh is the case (Section 7.3), and that S's reason for op is sensitive, evaluated in terms of the counterfactual situation in which not-sh is the case (Section 7.4). Second, I argue that S's reason for op is insensitive, evaluated in terms of the counterfactual situation in which sh is the case, and that S's reason for not-sh is sensitive, evaluated in terms of the counterfactual situation in which not-sh is the case (Section 7.5). The first two steps give rise to the first premise of my main argument:

> 7.1   S's reason for op (or not-sh) is sensitive when evaluated in terms of the counterfactual situation in which not-sh is the case, but fails to be the case when evaluated in terms of the counterfactual situation in which sh is the case.

In the third step, I first argue that which counterfactual situation is in play when one is evaluating S's reason depends on the epistemic inquiry one is engaging in. Different inquires may call for different counterfactual situations. I then argue that the skeptic and the ordinary people are engaging in different inquiries. I distinguish what I will call the 'non-ordinary-world' evaluation and the 'ordinary-world' evaluation; they determine the counterfactual situation in play in the skeptical and ordinary contexts respectively. I argue for the second premise of my argument:

> 7.2   The skeptic engages in the non-ordinary-world evaluation according to which the counterfactual situation in play when evaluating the sensitivity of S's reason for op is the one in which sh is the case, while an ordinary subject engages in the ordinary-world evaluation according to which the counterfactual situation in play when evaluating the sensitivity of S's reason for op is the one in which not-sh is the case.

From 7.1 and 7.2, we thus can conclude that:

> 7.3   S's reason for op (and not-sh) is sensitive according to the
>   ordinary-world evaluation and fails to be the case according to the
>   non-ordinary-world evaluation.

7.3 constitutes the core of the pluralistic account that I want to defend in this thesis. I call it 'the sensitivity-variantist solution (to the skeptical puzzle)'. It is a kind of standing-variantist solution, since it presupposes the variability of S's standing on knowing dimensions with respect to p—more precisely, the variability of S's epistemic position with respect to p as opposed to the variability of the standards for knowledge. The sensitivity-variantist solution claims that S's reason for op (and not-sh) is sensitive in the ordinary context but fails to be the case in the skeptical context.

7.3 The Skeptic Is Right about "S does not know that not-sh"

Notice that the skeptic's denial of S's knowledge of not-sh does not consists in her denial of not-sh, nor her denial of S's belief that not-sh. The skeptic claims that S's epistemic position with respect to not-sh falls short of satisfying the standards for knowledge in play. We have seen that whether one's epistemic position with respect to p satisfies the standards for knowledge depends on whether one's reason R for p is sensitive or not (see Chapter 4). The skeptic's claim, hence, is tantamount to the contention that S's reason R for not-sh is *not* sensitive.

We have seen that the sensitivity of one's reason R for p should be characterized in terms of counterfactual conditionals (i.e., the counterfactual-condition formulation; see Section 4.2). One's reason R for p is sensitive if and only if S would not have R if p were not the case. Put more formally,

> (SEN) S's reason R for p is sensitive if and only if "not-p > not-$R_S(p)$" is
>   true in M,

where '>' stands for the connection of counterfactual conditional, '$R_S(p)$' stands for *S had the reason R for p*, and M represents the situation S is in.

(SEN) specifies the sensitivity of one's reason in terms of counterfactual conditionals. I have argued that the truth condition of counterfactual conditionals, as long as such conditionals are relevant to the characterization of knowledge, is determined by (CM):

>(CM)   "A > C" is true in M if and only if C is true in M_A,

where M is the causal model representing the actual situation (i.e., S's situation), while M_A is generated by intervening in M with respect to A (see Chapter 6). (CM) and (SEN) imply that:

>7.4   S's reason for p is sensitive if and only if "not-$R_S(p)$" is true in M_{not-p}.

7.4 specifies the (sufficient and necessary) conditions for the sensitivity of one's reason. For reasons that will become clear later, let us call 7.4 'the *special* sensitivity condition (SSC)'.

Now, the skeptic is claiming that S's reason for sh is insensitive. For the sake of discussion, let 'SH' be the skeptical hypothesis that S is a handless **brain in a vat** (**BIV**) (being stimulated to experience having hands).

The skeptic is claiming that S's reason for SH is insensitive, namely, S would still have the same reason R for not-SH even if SH had been the case. By SSC, S's reason for SH is insensitive if and only if "$R_S(not-SH)$" is true in M_{SH}. Hence, the skeptic's claim is tantamount to the following:

>7.1.1   "$R_S(not-SH)$" is true in M_{SH},

where M is the causal model represented the situation S is in, and M_{SH} is the causal model generated by intervening in M with respect to SH.[86] In what follows, I will argue for 7.1.1.

First, let us construct the causal model M representing S's situation. For simplicity's sake, let us assume that S's (sole) reason for believing that S is not a handless BIV is S's experiences of having hands. Intuitively, the set of variables of M consists in:

> **E**XOGENOUS represents the factors that determine S being a handless BIV.
>
> **B**IV represents whether the factors that determine S being a handless BIV or living in the real world (with the value "0" standing for the former, "1" for the latter).
>
> **H**AND represents whether S has both of his hands or not.
>
> HAND**L**ESS represents whether S has lost both of his hands or not.
>
> **E**XPERIENCE represents whether S has the experience of having hands or not.

We can all agree that whether S lives in the real world and has both of his hands or not depends on whether the factors that determine S being a handless BIV obtain or not. If S lives in the real world and has both of his hands, he will have the experiences of having hands. However, if S is a handless BIV, S will still have the (same) experiences of having hands. Hence, the set of structural equations *S* of M is as follows (see Figure 7.1):

> BIV = EXOGENOUS
>
> HAND = 1 – BIV

---

[86] Given that a submodel is generated by manipulating the original model (see Chapter 6), 7.1.1 expresses the idea that if we were to manipulate the actual situation M such that SH had been the case, S would still have the same reason R for not-SH as he does in the actual situation.
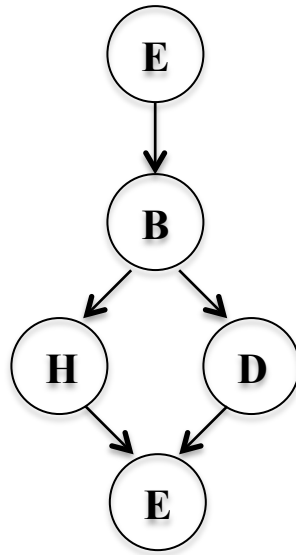
HANDLESS = BIV

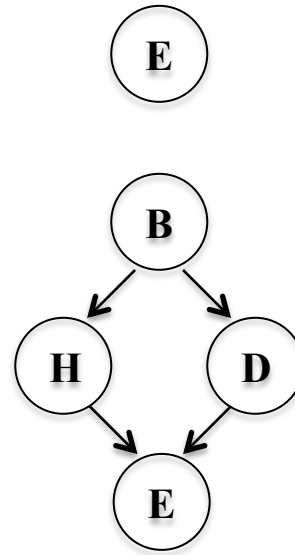EXPERIENCE = max(HAND, HANDLESS)

Figure 7.1
DAG of M

Figure 7.2
DAG of
$M_{BIV=1}$

Now, since the skeptic is not claiming that SH is true, the skeptic may grant that, for the sake of argument, S is not a handless BIV, S lives in the real world and has both of his hands, and S has the experiences of having hands. In other words, the value assignment A of M is thus:

A(EXOGENOUS) = A(BIV) = A(HANDLESS) = 0, and

A(HAND) = A(EXPERIENCE) = 1.

The skeptic claims that S's reason for not-SH—S's experiences of having hands—is insensitive. According to (SEN), the skeptic's claim amounts to the claim that "BIV = 1 > EXPERIENCE = 0" is false in M. By SSC, this is tantamount to the claim that "EXPERIENCE = 1" is true in $M_{BIV=1}$. Let us then intervene in M in terms of BIV = 1.

The set of structural equations $S_{BIV=1}$ consists of (see Figure 7.2):

BIV = 1

HAND = 1 – BIV

HANDLESS = BIV

EXPERIENCE = max(HAND, HANDLESS).

The value assignment $A_{BIV=1}$ is as follows:

$A_{BIV=1}$(EXOGENOUS) = $A_{BIV=1}$(HAND) = 0, and

$A_{BIV=1}$(BIV) = $A_{BIV=1}$(HANDLESS) = $A_{BIV=1}$(EXPERIENCE) = 1.

We can see that "EXPERIENCE = 1" is indeed true in $M_{BIV=1}$. Hence, the skeptic's claim that S's reason for not-SH is insensitive is indeed correct; 7.1.1 obtains. Given that sensitivity is a necessary condition for knowledge, the skeptic is thus right in denying that S knows that not-SH.

7.4 The Ordinary Subject Is Right about "S knows that op"

A similar story can be told about the commonsensical idea that S knows that op. Let 'OP' stand for the ordinary proposition *S has hands*, and 'the ordinary subject' for an ordinary, average cognizer who takes S as knowing that OP. If the ordinary subject is right, S's (sole) reason R for OP must be sensitive.

To say that S's reason R for OP is sensitive, according to SSC, is tantamount to saying that:

7.1.2   "Not-R_S(OP)" is true in M'$_{OP}$,

where M' is the causal model representing the situation S is in, and M'$_{OP}$ is the submodel generated by intervening in M' with respect to OP.[87] In this section, I argue for 7.1.2.

To determine whether 7.1.2 is true in M' or not, we need to first construct the causal model M'. For simplicity's sake, let us assume that S's sole reason for OP is that S has the experiences of having hands. We may further assume that the ordinary subject and the skeptic understand S's situation in the same way; they construct the same causal model to represent the situation S is in. In other words, M = M'. This assumption is not implausible since the skeptic and the ordinary subject do not have to disagree with each other about the factual matters. In particular, the skeptic may, for the sake of argumentation, agree with the ordinary subject about what happens in the real world (to me, skeptical arguments along this line are the most powerful ones). The skeptic's claim is that S's reason for OP is insensitive.

The ordinary subject contends that S's reason for OP—S's experience of having hands—is sensitive. According to SSC, the ordinary subject's claim is equivalent to the claim that "EXPERIENCE = 0" is true in M$_{HAND=0}$. Let us, then, intervenes in M in terms of HAND = 0.

The set of structural equations $S_{HAND=0}$ consists of (see Figure 7.3):

BIV = EXOGENOUS

HAND = 0

HANDLESS = BIV

EXPERIENCE = max(HAND, HANDLESS).

---

[87] Given that a submodel is generated by manipulating the original model (see Chapter 6), 7.1.2 expresses that if we were to intervene in (manipulate) the actual situation (represented by M') such that OP had not been the case, S would not have had the same reason R for OP in such a situation (represented by M'$_{OP}$).

The value assignment $A_{HAND=0}$ is as follows:

$$A_{HAND=0}(EXOGENOUS) = A_{HAND=0}(BIV) = A_{HAND=0}(HANDLESS) =$$
$$A_{HAND=0}(HAND) = A_{HAND=0}(EXPERIENCE) = 0.$$

Obviously, "EXPERIENCE = 0" is true in $M_{HAND=0}$. Hence, the ordinary subject's claim that S's reason for op is sensitive is indeed true; 7.1.2 obtains. No wonder we find the utterance "S knows that OP" intuitive.
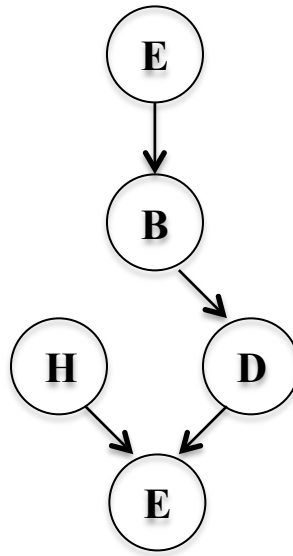

Figure 7.3
DAG of
$M_{HAND=0}$

Let me take stock. 7.1.1 implies that S's reason R for not-SH is insensitive when the counterfactual situation according to which the sensitivity of R is evaluated is $M_{BIV=0}$. 7.1.2, by contrast, implies that the same reason R, when used to support OP, is sensitive if the counterfactual situation in terms of which the sensitivity of R is evaluated is $M_{HAND=0}$. In brief, R, as the support of not-SH, is insensitive in M, while the same R, as the support of OP, is sensitive in M. One important distinction between $M_{BIV=0}$ and $M_{HAND=0}$,

however, is that the latter, but not the former, presupposes BIV = 0. That is, not-SH is the

case in $M_{HAND=0}$, while not-SH is not the case in $M_{BIV=0}$.

We can thus formulate the logical conclusion of 7.1.1 and 7.1.2 as follows:

7.1.3   S's reason R for OP is sensitive when evaluated from the counterfactual
         situation in which not-SH is the case, and the same reason R for not-SH is
         insensitive if evaluated from the counterfactual situation in which SH is
         the case.

7.1.3 gets us halfway to our destination—i.e., the idea that the skeptic's denial of

S's knowledge of op and not-sh and the ordinary subject's attribution of S knowledge of

op and not-sh are both correct—but we are not there yet. The remaining problem is that

while 7.1.3 deals with the sensitivity of R's reason for OP (and not-SH) when evaluated

in terms of the situation in which not-SH (and SH) is the case, it does not deal with the

sensitivity of S's reason R for OP when evaluated in terms of the situation in which *SH is*

the case. Nor does it account for the sensitivity of R for not-SH when evaluated in terms

of the situation in which *not-SH is the case*. The next section takes up this task.

7.5 The General Context-Dependence of Sensitivity

When evaluating the sensitivity of S's R for OP (or not-SH), we ask the question whether

S would still have the same R if OP (or not-SH) had not been the case. The idea is that

the sensitivity of R for OP (or not-SH) depends on the counterfactual situation according

to which R is evaluated. So far, I have assumed that the sensitivity of S's R for p is to be

characterized by SSC.

However, as the characterization of the evaluation of S's reason, SSC is *limited*.

In particular, it does not allow that the same counterfactual situation to be used to

evaluate the sensitivity of R for p *as well as* the sensitivity of R for q (where p ≠ q). More

precisely, according to SSC, we can evaluate the sensitivity of R for p in terms of the counterfactual situation in which p had not been the case, and we can evaluate the sensitivity R for q in terms of the counterfactual situation in which q had not been the case, but we *cannot* account for the sensitivity of R for q in terms of the counterfactual situation in which p had been the case. Nor can we evaluate the sensitivity of R for p in terms of the counterfactual situation in which q is the case.

For the present purposes, this limitation of SSC is unacceptable. To illustrate, notice that for any distinct propositions p and q, the comparative strength between one's epistemic position with respect to p and one's epistemic position with respect to q can often be tested by appealing to comparative conditionals such as "If S knows that p, then S knows that q" or "If S does not know that q, then S does not know that p". As DeRose points out, "the intuitive correctness of *If S knows that P, then S knows that Q* and *If S doesn't know that Q, then S doesn't know that P* can indicate that S is in at least as strong an epistemic position with respect to Q as she's in with respect to P" (1995, 31).

In the case of the skeptical puzzle, we surely find it intuitively correct that if S knows that OP, then S knows that not-SH, and that if S does not know that not-SH, S does not know that OP (i.e., 1.3). Arguably, this indicates that our epistemic position with respect to not-SH is at least as good as our epistemic position with respect to OP. This last claim is at any rate plausible. For instance, suppose that one's sole reason for OP is R_1. It follows that R_1 is also one's reason for not-SH. If one also has other reason R_2 for not-SH, then one's reason for not-SH is even better than one's reason for OP. In other words, one's reason for not-SH is as least as good as one's reason for OP.

The tracking-based account can easily account for the comparative conditional "If S knows that OP, then S knows that not-SH". On this account, if S knows that p, competently and adequately infers not-SH from OP, and thereby comes to believe that not-SH, while retaining knowledge of OP throughout, then S also knows that not-SH (see Section 4.6). The interesting question is whether the tracking-based account can also account for why S's epistemic position with respect to not-SH is as least as good as S's epistemic position with respect to OP.

The tracking-based account suggests that S's epistemic position with respect to p is to be characterized in terms of the sensitivity of S's reason R for p. If S's epistemic position with respect to not-SH is at least as good as S's epistemic position with respect to OP, it seems that the following must also hold:

(a) if S's reason for OP is sensitive, then S's reason for not-SH is also sensitive, and

(b) if S's reason for not-SH is insensitive, then S's reason for OP is insensitive.

Now, traditional proponents of the sensitivity account have characterized sensitivity along the line of SSC. The problem or limitation of SSC is that it cannot account for (a) and (b). One just needs to look at the results of the previous two sections to see this point. In Section 7.4, we showed that, according to SSC, S's reason R—S's experiences of having hands—for OP is sensitive. However, we also saw, in Section 7.3, that SSC predicts that the same R for SH is insensitive. Hence, it seems that SSC putatively leads to the denial of (a) and (b).[88]

---

[88] Perhaps one should not be too surprised by the fact that SSC has difficulties dealing with cases like (a) and (b), since it has been well aware of by early proponents of the sensitivity condition that sensitivity is in tension with principle of epistemic closure (cf. Nozick 1981; Dretske 1970). That is, many have pointed out

In this section, I want to propose a characterization of the sensitivity condition that can comfortably account for cases like (a) and (b). My proposal is that we should generalize SSC such that we can evaluate the sensitivity of one's reason for p and one's reason for q in terms of the same counterfactual situation (as will be explained below, this move is prohibited by SSC).

More precisely, I will generalize SSC in two steps. First, we may reformulate SSC as follows:

> SSC' S's reason R for p is sensitive in M if and only if "not-p $\supset$ not-$R_S(p)$" is true in M_{not-p},

where '$\supset$' stands for material implication. This move is uncontroversial since SSC and SSC' are equivalent (assuming M does not allow contradictions).

For the present purposes, SSC' is very useful since it illustrates the underlying thoughts that motivate SSC. Recall that knowing that p is incompatible with believing truly that p *accidentally* (see Chapter 4). To prevent accidentality, we want to avoid the situation such that S would have the same R had p not been the case. The fact that S would still have R while p had not been the case indicates that R is a 'false positive' of p—it falsely indicates the truth of p. Hence, what we are preventing here, in cases of knowledge, is that S's reason for p could have easily been a false positive.[89] In other

_____

that one of the natural implications of accounts like SSC is that it allows S's belief that (or reason for) p to be sensitive and S's belief that (or reason for) q to be insensitive, even if p implies q. However, it is possible for the sensitivity condition to account for the idea of epistemic closure. The idea, roughly, is that as long as the counterfactual situation in terms of which S's belief that (or reason for) p is sensitive is held fixed, S's belief that (or reason for) p will be sensitive as well. The following account can be seen as a development along this line.

[89] This idea is in line with Williamson's (2000) idea that knowing that p implies that one's belief that p could not be easily false (in my formulation: One's reason for p would not be easily lost). Notice, however, that Williamson may not agree that this idea can be captured by 'the general sensitivity condition (GSC)' mentioned below.

words, the underlying thought that motivates SSC is that we want to prevent R for p from being a false positive in the *relevant* counterfactual situation (according to SSC', the only relevant situation is the situation in which p had not been the case).

Cases of (a) and (b) suggest a natural way to generalize SSC'. According to SSC', the relevant counterfactual situation in evaluating the sensitivity of R for OP (or not-SH) *must* be the situation in which OP (or not-SH) had not been the case. However, given that the sensitivity of R for not-SH depends on the sensitivity of R for OP (i.e., (a)), and that the insensitivity of R for OP depends on the insensitivity of R for not-SH (i.e., (b)), it seems natural for us to give up the idea that one's reason R for p can only be evaluated in terms of counterfactual situations in which p failed to hold. That is, in the cases of (a), it seems natural that we should evaluate the sensitivity of R for not-SH in terms of the counterfactual situation that is relevant in evaluating the sensitivity of R for OP, since, in such cases, if R for OP is sensitive, R for not-SH is *thereby* sensitive. Likewise, in the cases of (b), it seems natural that we should evaluate the insensitivity of R for OP in terms of the counterfactual situation that is relevant in evaluating the insensitivity of R for not-SH, since, in such cases, if R for not-SH is insensitive, R for OP is *thereby* insensitive.

Hence, I propose that we should generalize SSC' (and hence SSC) as follows. Call it 'the *general* sensitivity condition (GSC)':

> GSC   S's reason R for p is sensitive in M if and only if "not-p $\supset$ not-$R_S(p)$" is true in M',

where M' is the causal model representing the counterfactual situation *in play* in evaluating the sensitivity of R for p (M' may not be the situation in which p were false).

Some clarifications are called for. First, while, strictly speaking, both SSC and GSC relativize the sensitivity of R for p in M to a submodel, only SSC, not GSC, implies that there is a unique submodel (i.e., $M_{p=0}$), determined by M, in evaluating the sensitivity of R for p. As a result, according to SSC, the sensitivity of R for p is held fixed once M is held fixed. GSC, to the contrary, allows different submodels to be adopted in evaluating the sensitivity of R for p; the sensitivity of p is not held fixed even if M is held fixed.

Second, SSC is in fact a special case of GSC. If the counterfactual situation in play in terms of which the sensitivity of R for p is evaluated is the situation in which p did not obtain, SSC and GSC are logically equivalent.[90]

Third, and more importantly, on GSC, there are two main ways S's reason R for p can be sensitive. On the one hand, S's R for p is sensitive if S would not have R in the counterfactual situation in which p did not obtain (i.e., "not-p $\supset$ not-$R_S(p)$" is true if "not-p" and "not-$R_S(p)$" is true) . On the other hand, S's R for p is sensitive if p did obtain (i.e., "not-p $\supset$ not-$R_S(p)$" is true if "p" is true). That is, in the counterfactual situation in which p obtained, R for p is automatically sensitive regardless of whether or not S would have R. For future reference, we may call the former 'substantial' sensitivity, and the latter 'trivial'. They are so called since in the cases of substantial sensitivity, S's reason R for p does not constitute a false positive of p since S would actually lack R in the situation in which p does not hold, while in the cases of trivial sensitivity, S's reason for p does not constitute a false positive since the situation is such that p does not fail to obtain (if p does not fail to obtain, S's R for p cannot be a false positive). If p were not false, then of course R for p *could not* be a false positive of p.

---

[90] When the counterfactual situation in play in terms of which the sensitivity of R for p is evaluated is the situation in which p did not obtain, GSC is just SSC'. And SSC' and SSC are logically equivalent.

In general, GSC is able to account for how the sensitivity of R for not-q may be determined by the sensitivity of R for p or the insensitivity of p be determined by the insensitivity of not-q (where p and q are incompatible). It can be easily proven that if R for p is substantially sensitive, then R for not-q is at least trivially sensitive. Likewise, if R for not-q is insensitive, then R for p is also insensitive.[91]

In particular, GSC is able to account for cases like (a) and (b) according to which the sensitivity (or insensitivity) of R for p depends on the sensitivity (or insensitivity) of R for q. The trick, roughly, is *to hold fixed the relevant counterfactual situation when evaluating the sensitivity of R for OP and the sensitivity of R for not-SH (or the insensitivity of R for not-SH and the insensitivity of R for OP)*. That is, in the cases of (a), the sensitivity of R for not-SH should be evaluated in terms of the counterfactual situation that is relevant in evaluating the sensitivity of R for OP. Similarly, in the cases of (b), the insensitivity of R for OP should be evaluated in terms of the counterfactual situation that is relevant in evaluating the insensitivity of R for not-SH (i.e., the counterfactual situation in which SH obtained). These moves are not be implausible since

---

[91] An informal proof: Suppose that M is the relevant counterfactual situation. R for p is sensitive if and only if "not-p $\supset$ not-R_S(p)" is true in M. There are three possible cases for R for p to be sensitive:

M1: "not-p & not-R_S(p)" is true in M. (Substantial Sensitivity)
M2: "p & R_S(p)" is true in M. (Trivial Sensitivity)
M3: "p & not-R_S(p)" is true in M. (Trivial Sensitivity)

R for p is insensitive if and only if the following holds:

M4: "not-p & R_S(p)" is true in M.

Suppose that p and q are incompatible. R for q is sensitive if and only if "not-q $\supset$ not-R_S(q)" is true in M. Likewise, there are three possible cases for R for q to be sensitive:

M5: "not-q & not-R_S(q)" is true in M. (Substantial Sensitivity)
M6: "q & R_S(q)" is true in M. (Trivial Sensitivity)
M7: "q & not-R_S(q)" is true in M. (Trivial Sensitivity)

R for q is insensitive if and only if the following holds:

M8: "not-q & R_S(q)" is true in M.

Now, if M1 is the case (i.e., S's R for p is substantially sensitive), then ether either M5 (i.e., S's R for not-q is substantially sensitive) or M7 (i.e., S's R for not-q is trivially sensitive) is the case. That is, if R for p is substantially sensitive, then R for q is (at least trivially) sensitive. Moreover, if M8 is the case, then M4 is the case. That is, if R for q is insensitive, then R for p is insensitive.

the sensitivity of R for not-SH depends on the sensitivity of R for OP and the insensitivity of R for OP depends on the insensitivity of R for not-SH.

To elaborate, recall that the relevant counterfactual situation in terms of which the sensitivity of S's R for OP is evaluated is represented by $M_{HAND=0}$, and that the relevant counterfactual situation in terms of which the insensitivity of S's R for not-SH is evaluated is represented by $M_{BIV=1}$ (see Section 7.3 and Section 7.4). Now, like SSC, GSC also predicts that S's R for OP is sensitive when evaluated in terms of $M_{HAND=0}$, and that S's R for not-SH is insensitive when evaluated in terms of $M_{BIV=1}$ (given that SSC is a special case of GSC, this should be obvious).[92]

As a result, on GSC, in cases of (a), we should evaluate the sensitivity of R for not-SH in terms of $M_{HAND=0}$ since $M_{HAND=0}$ represents the relevant counterfactual situation in terms of which the sensitivity of R for OP is evaluated. Clearly, in this case, R for not-SH is sensitive in terms of $M_{HAND=0}$, since "BIV = 1 $\supset$ EXPERIENCE = 0" is true in $M_{HAND=0}$—in fact, R for not-SH is trivially sensitive in terms of $M_{HAND=0}$, given that "BIV = 1" is false in $M_{HAND=0}$.

By the same token, in cases of (b), we should evaluate the insensitivity of R for OP in terms of $M_{BIV=1}$ since $M_{BIV=1}$ represents the relevant counterfactual situation in terms of which the insensitivity of R for not-SH is evaluated. Clearly, in this case, R for OP is insensitive in terms of $M_{BIV=1}$, since "HAND = 0 $\supset$ EXPERIENCE = 1" is true in

---

[92] To illustrate, recall that the value assignment $A_{HAND=0}$ is as follows:

$A_{HAND=0}(BIV) = A_{HAND=0}(HANDLESS) = A_{HAND=0}(HAND) = A_{HAND=0}(EXPERIENCE) = 0$,

and the value assignment $A_{BIV=1}$ is as follows:

$A_{BIV=1}(HAND) = 0$, and

$A_{BIV=1}(BIV) = A_{BIV=1}(HANDLESS) = A_{BIV=1}(EXPERIENCE) = 1$.

Clearly, on SSC', R for OP is sensitive in terms of $M_{HAND=0}$, since "HAND=0 $\supset$ EXPERIENCE = 0" is true in $M_{HAND=0}$. Similarly, R for not-SH is insensitive in terms of $M_{BIV=1}$, since "BIV = 1 $\supset$ EXPERIENCE = 1" is true in $M_{BIV=1}$.

$M_{BIV=1}$—in fact, R for OP is insensitive in terms of $M_{BIV=1}$, given that "EXPERIENCE = 1" is true in $M_{BIV=1}$.

In other words, GSC gives rise to the following results:

(a)'   if S's reason R for OP is sensitive, then S's reason R for not-SH is (trivially) sensitive, and that

(b)'   if S's reason R for not-SH is insensitive, then S's reason for OP is insensitive.

Both (a)' and (b)' are in line with our ordinary intuitions regarding OP and not-SH. On the one hand, we all realize that if S knows that OP, then S must somehow have already known not-SH, since knowledge of OP presupposes knowledge of not-SH. (a)' explains this intuitions pretty well: If S's reason R for OP is indeed sensitive, then S's reason R for not-SH is trivially sensitive—the latter's sensitivity follows directly from the former's sensitivity. On the other hand, we also admit that if one is not in a position to know that one is not a BIV via one's sensory experiences, one is also not in a position to know that one has hands via the same sensory experiences. This is exactly what (b)' delivers.

Now, if GSC correctly formulates the sensitivity condition, we have the following sub-premise:

> 7.1.4   S's reason R for not-SH is (trivially) sensitive when evaluated in terms of the counterfactual situation in which not-SH is the case, and the same reason R for OP is (trivially) insensitive if evaluated in terms of the counterfactual situation in which SH is the case.

From 7.1.4 and 7.1.3, i.e.,

> 7.1.3   S's reason R for OP is sensitive when evaluated from the counterfactual situation in which not-SH is the case, and the same reason R for not-SH is insensitive if evaluated from the counterfactual situation in which SH is the case,

we get the desired conclusion that:

7.1   S's reason for OP (or not-SH) is sensitive when evaluated in terms of the counterfactual situation in which not-SH is the case, but fails to be the case when evaluated in terms of the counterfactual situation in which SH is the case.

In the previous three sections, I have tried to argue for the first premise of my main argument. In the next three sections, I will then present my argument for the second premise of the main argument.

Before we get to that, it is worth noting an important theoretical feature of GSC. That is, GSC consists in a novel interpretation of the sensitivity condition:

The general context-dependence Thesis. Whether S's reason for p is sensitive or not depends on the counterfactual situation *in play*.

We have seen that if the counterfactual situation in play is such that not-SH was the case, then S's reason R for OP is substantially sensitive and the same R for not-SH is thereby trivially sensitive. By contrast, if the counterfactual situation in play is such that SH was the case, then S's reason R for not-SH is insensitive and the same reason R for OP is thereby insensitive. However, little has been said about how a counterfactual situation may be in play. This is the topic of the next section.

7.6 The Inquiry-Dependence of Sensitivity

The question, hence, is this: How do we know which counterfactual situation we should consider when evaluating the sensitivity of R for p? My proposal is that which counterfactual situation is in play in evaluating the sensitivity of S's reason depends on the *inquiry* the relevant parties—the people who are evaluating S's reason—are engaged in.

By inquiry, I mean the general search for certain type of information in a purposeful way. Inquiry differs from aimless gathering of information in that the former, but not the latter, is guided by certain considerations or purposes. Simply put, inquiry has a *topic*. Inquiries are individuated by their topics. To inquire whether *Smith* stole the book is different from inquiring whether Smith *stole* the book, or whether Smith stole a *book*. Not all inquiries aim at truth. A guilty person may inquire into information indicating that she is innocent. In what follows, however, I will limit myself to *epistemic* inquiry—inquiry that actually aims at pursuing truth and avoiding error.

The evaluation of S's reason for R depends on the counterfactual situation in play. Different inquiries may call for different counterfactual situations. Consider the following case:

> *Testimony.* Kenny was murdered last night. Adam testifies that John alone murdered Kenny. Suppose that Adam's testimony will be very reliable unless Kenny was murdered by his daughter Mary. If Mary were the murderer, Adam would still testify that John alone murdered Kenny, in order to spare Mary from being caught.

Suppose that Smith is inquiring into whether John or Johnson murdered Kenny. From Smith's perspective, Johnson will be the murderer if John is not. Intuitively, in Smith's context, the sensitivity of Adam's testimony for *John murdered Kenny* should be evaluated in terms of the situation in which Johnson were the murderer—Adam's testimony is sensitive only if Adam's would not so testify if Johnson were the murderer. It seems that Adam's testimony is sensitive in Smith's context. By contrast, suppose that Janet is inquiring into whether John or Mary murdered Kenny. From Janet's perspective, Mary will be the murderer if John is not. Intuitively, in Janet's context, the sensitivity of Adam's testimony for *John murdered Kenny* should be evaluated in terms of the situation

in which Mary were the murderer—Adam's testimony is sensitive only if Adam's would not so testify if Mary were the murderer. It seems that Adam's testimony is not sensitive in Janet's context (one might object that it cannot be the case that both Smith's and Janet's epistemic appraisals are correct. I will come back to this point).

Cases like *Testimony* indicates a special kind of the context-dependence thesis for sensitivity, namely, the idea that the counterfactual situation in play in one's evaluation the sensitivity is determined by the inquiry in which one engages. Call it 'the inquiry-dependence thesis (for sensitivity)':

> *The inquiry-dependence thesis*. The inquiry in which one is engaging determines the counterfactual situation in play when one evaluates the sensitivity of S's reason for p.

As I see it, the inquiry-dependence thesis explains why the skeptic is right in her context in claiming that S's reason R for OP (or not-SH) is insensitive while the ordinary subject is also right in claiming that S's reason R for OP (or not-SH) is sensitivity in her own context. My proposal is that the skeptic and the ordinary subject are engaging in different inquiries. The skeptic is inquiring into whether S's reason would be sensitive if S had been a handless BIV.[93] The ordinary subject, by contrast, is inquiring into whether S's reason would be sensitive if S had not had hands. It should not be surprising that different counterfactual situations are in play in the skeptic's and ordinary subject's contexts, since the skeptic is inquiring into what would have happened if S had been a handless BIV and the ordinary subject is inquiring into what would have happened if S had not had hands.

---

[93] Of course, to say that the skeptic inquires into whether or not S's R for not-SH is sensitive simplifies the story a bit. The complete story is that the skeptic is inquiring into whether or not S knows that OP. And this leads her to consider whether or not S knows that not-SH. She then argues that S does not know that not-SH since S's reason R for not-SH is insensitive.

Of course, just because different epistemic inquiries may give rise to different epistemic appraisals, it does not follow that all epistemic appraisals are equally correct or legitimate. Suppose that John claims that the movement of Mars will affect Mary's fortunes. Suppose that an astrologist thinks that John's reason for *The movement of Mars will affect Mary's fortunes* is sensitive, while an astronomer takes John's reason to be insensitive. In one sense, one can say that John's reason is sensitive, according to the astrological appraisal, while insensitive according to the astronomical appraisal. Yet, intuitively, we do not take both appraisals to be equally correct. There is a legitimate sense that the astrologist's evaluation of John's reason is incorrect, while the astronomer's, correct.

Likewise, in *Testimony*, is it not the case that either Adam's testimony for *John alone murdered Kenny* is either sensitive or not? More precisely, if Smith is right in presupposing that Mary is innocent, then Janet's presupposition that Mary *may* be the murderer is objectively incorrect. Hence, even if Adam's testimony is insensitive from Janet's point of view, it is in fact objectively sensitive. While these considerations do not falsify the inquiry-dependent thesis, they do shed an important light on the current issue. What we need is to have a better understanding of different types and levels of inquiries. In some local inquiries, some inquiries are considered correct or incorrect. But such a clear-cut correctness or incorrectness tends to disappear as we move on to a more general inquiry (I will get back to this issue in Section 7.8).

For the next three sections, I will construct my positive argument for the claim that both the skeptic's denials of knowledge and the ordinary subject's attributions of knowledge are correct according to the inquiry they are engaging in. Section 7.7 argues

that the skeptic and the ordinary subject are in fact engaging distinct epistemic inquiries. Section 7.8 to Section 7.10 argue that these two inquiries are both legitimate epistemic inquiries.

7.7 The Skeptic and the Ordinary Subject Are Engaging in Different Inquiries

On the present account, the skeptic and the ordinary subject have to engage in different inquiries. Since inquiries are individuated by their topics, the skeptic and the ordinary subject are inquiring into different topics. How can we tell whether the skeptic and the ordinary subject are inquiring into different topics or not? How do we identify the topic of an inquiry, or distinguish different topics of different inquiries?

My proposal is that the topic of an inquiry is determined (partly) by the propositions it *presupposes*. Inquiries require presuppositions; as Wittgenstein writes in *On Certainty* (1969), an inquiry cannot be exempted from presuppositions:

> 341. That is to say, the *questions* that we raise and our *doubts* depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.

> 342. That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted.

> 343. But it isn't that the situation is like this: We just *can't* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.[94]

Inquiry cannot get off the ground without a certain presupposition (what Wittgenstein calls 'assumption'), since one cannot inquire into everything *at once*. If I want to know whether the earth is round or not, I have to take for granted that the earth exits. If I am investigating whether the butler is the murderer, I have already assumed that someone is

---

[94] Unless specified otherwise, the quotations of Wittgenstein below are from *On Certainty*.

murdered. What Wittgenstein calls "the logic of scientific investigations" is in fact the logic of inquiry.

Differences in inquiries can be distinguished by their presuppositions. I am inquiring into whether *Smith* stole the book, while you are inquiring into whether Smith *stole* the book. We engage in different inquiries since my inquiry presupposes that the book is stolen and leaves it open whether Smith has taken the book or not, while your inquiry presupposes that Smith has taken the book but leaves it open whether Smith stole it or not.

Two senses of differences, however, are relevant here. We are engaging in the inquiry about the culprit of a crime, say, we both want to know who stole the book. I wonder whether *Smith* stole the book, while presupposing that Mary did not do it. You wonder whether *Mary* stole the book, while presupposing that Smith is innocent. In one sense, we are engaging in different inquiries—my inquiry is about Smith while yours, Mary. Yet, in another sense, it seems that we are still engaging in the same inquiry—we are both inquiring about the culprit of the crime. While we presuppose different people to be innocent, our inquiries share the same presupposition that the book was stolen. More importantly, while the former does not seem to determine the subject matter of our inquiries, the latter does—I can still engage in the inquiry about the culprit of the crime even if I discard the presupposition that Mary is innocent, but the same cannot be said if I discard the presupposition that the book was stolen.

Hence, there are two kinds of presuppositions. Some presuppositions are crucial in defining the topic of an inquiry—they *constitute* its topic—and some are not. Wittgenstein once uses an analogy of the river-bed to describe these two kinds of

presuppositions. The target of an inquiry is like the fluid in a river, while its presuppositions are like the river-bed. More importantly, the presuppositions of an inquiry fall into two kinds just like there are two kinds of components of the river-bed:

> 99. And the bank of that river consists partly of hard rock, subject to no alteration or only to an imperceptible one, partly of sand, which now in one place now in another gets washed away, or deposited.

Some presuppositions are not robust; they can be altered or discarded, just like the sand in the river-bed, which can easily be washed away or deposited in some other places. They do not affect the general topic of the inquiry. Some presuppositions are as robust as hard rock, and not subject to alteration. These are the presuppositions that define, or constitute, the inquiry.[95] For instance, in doing history, the presupposition that historical reports are not universally unreliable defines what counts as doing history:

> 163. We check the story of Napoleon, but not whether all reports about him are based on sense-deception, forgery and the like. For whenever we test anything, we are already presupposing something that is not tested.

Such presuppositions own a special normative status within a certain inquiry:

> 337. One cannot make experiments if there are not some things that one does not doubt. But that does not mean that one takes certain presuppositions on trust. When I write a letter and post it, I take it for granted that it will arrive—I expect this.

> If I make an experiment I do not doubt the existence of the apparatus before my eyes. I have plenty of doubts, but not that. If I do a calculation I believe, without

---

[95] Admittedly, the distinction between these two kinds of presuppositions is *vague*, *relative*, and *coming in different varieties*. Firstly, there is no clear-cut distinction between presuppositions that define the inquiry and presuppositions that do not. For instance, it is not sure whether the presupposition that the world is deterministic is a defining feature of empirical inquiry. Secondly, whether the presuppositions define an inquiry is related to a certain inquiry. For instance, the presupposition that the movements of planets affect human life defines astrological inquiry. But the same presupposition does not define empirical inquiry. Thirdly, some presuppositions that define an inquiry are standing certainties, such as the world does not come to exist five minutes ago. They are presupposed across many kinds of inquiries. Others are certainties in particular situations such as John's cognitive faculty is reliable. For simplicity's sake, I will not deal with such nuances here. But, as far as I can tell, it should be safe to say that nothing significant hinges on such nuances.

any doubts, that the figures on the paper aren't switching of their own accord, and I also trust my memory the whole time, and trust it without reservation.

Michael Williams adopts and elaborates Wittgenstein's ideas. He calls such presuppositions *methodological necessities*. Like Wittgenstein, Williams takes methodological necessities to play a normative role:

> A proposition that serves as a 'methodological necessity' has a special normative status within a particular type of inquiry, in the sense that exempting it from doubt is a precondition of engaging in inquiry, of that type. (Williams 2001, 227)

That methodological necessities are exempted from doubt is crucial to an inquiry:

> Methodological necessities are a source of default entitlements because they determine the *direction of inquiry*. For example, serious worries as to whether the Earth even exited five minutes ago, or whether every piece of documentary evidence is some kind of forgery, do not result in an especially scrupulous approach to historical investigation. Rather, they preclude any engagement in historical research… some questions have to be set aside for us to think historically at all. (Williams 2001, 160)

The "direction" of an inquiry determines what questions can be legitimately challenged when engaging in a certain inquiry. The presuppositions that determine the direction of inquiry enjoy a default entitlement not to be questioned, at least not within the inquiry. To engage in historical investigation at all, one must not question whether the earth even existed five minutes ago.

My contention is that the skeptic and the ordinary subject are engaging in the different *kinds* of inquiries since the methodological necessities of their inquiries are different. The skeptic is inquiring into what would have happened if S were a handless BIV. His inquiry does not presuppose that S is not a handless BIV. The ordinary subject, by contrast, is inquiring into what would have happened if S had not had hands. His inquiry does presuppose that S is not a handless BIV.

The presupposition that certain skeptical hypotheses do not obtain is crucial to the ordinary "world-picture"—the view about the nature of the world we inhabit. We take ourselves not to live in the kind of fancy worlds created by epistemologists. For instance, we are not living in the "Matrix" world in which everybody has been envatted. Nor are we living in the demon-evil world in which a powerful demon systematically and constantly deceives our senses. The ordinary world-picture is essentially an anti-skeptical-hypotheses view.

The ordinary subject is working under the ordinary world-picture. She is asking what would have happened if S had not had hands *in the ordinary world*. We may take the ordinary subject as trying to evaluate the sensitivity of S's reason R for p with respect to the ordinary world. Call the ordinary subject's inquiry 'the ordinary-world evaluation (of the sensitivity of S's reason)'. The skeptic, by contrast, is engaging in a different business; she is working under a world-picture that may *not*, at least not necessarily, be ordinary. The skeptic's inquiry does not presuppose the falsity of the skeptical hypotheses. It is necessary for the skeptic not to presuppose as such if she is to inquire into what would have happened if S were a handless BIV. What the skeptic is trying to evaluate is the sensitivity of S's reason R for p with respect to a world in which a certain skeptical hypothesis obtains. Call the skeptic's inquiry 'the non-ordinary-world evaluation (of the sensitivity of S's reason)'.

One of the methodological necessities of the ordinary-world evaluation is that S is not a handless BIV, while that S may be a handless BIV is a methodological necessity of the non-ordinary-world evaluation. One cannot engage in the ordinary-world evaluation

without presupposing that S is not a handless BIV; one cannot engage in the non-ordinary-world evaluation while presupposing that S is not a handless BIV.

The direction of inquiry of the ordinary-world and non-ordinary-world evaluation is thus different. Engaging in the ordinary-world evaluation, one cannot legitimately question whether the world S is in is an ordinary world. But the same question is legitimate in the pursuit of the non-ordinary-world evaluation. This explains why different counterfactual situations are in play in these two kinds of inquiry. That the skeptic engages in the non-ordinary-world evaluation explains why, in the skeptic's context, the counterfactual situation in which S is a handless BIV is in play in evaluating the sensitivity of S's reason. By the same token, that the ordinary subject engages in the ordinary-world evaluation explains why the counterfactual situation in which S is not a handless BIV is in play in evaluating the sensitivity of S's reason.

Moreover, that S's R for not-SH (OP) is trivially sensitive (insensitive) in the ordinary (skeptical) context is also explained. In the ordinary-world evaluation according to which the ordinary world is presupposed, it is impossible for not-SH to fail to obtain (it is the definition of the ordinary world that not-SH is the case). Hence, any reason that supports not-SH will be trivially sensitive. By contrast, in the non-ordinary-world evaluation according to which the non-ordinary world is presupposed (or the ordinary world is not presupposed), it is possible for S's R for OP to be insensitive, if R for not-SH is insensitive in the first place. The insensitivity of R for not-SH, however, is guaranteed when evaluated in terms of the non-ordinary world-picture (see Section 7.3).

7.8 The Legitimacy of Inquiries

Just because p is a methodological necessity, however, does not mean that the acceptance of p is correct. Astrology and astronomy, arguably, are two distinct inquiries, consisting of inconsistent methodological necessities. But many, I suppose, will deny some methodological necessities of astrology. In other words, even though one cannot question the methodological necessities *within* an inquiry, one may still want to question it *outside*. So it seems that the "default entitlement" Williams attributes to methodological necessities are limited at best. It does not (and should not) prevent them from being challenged in a broader context.

Likewise, just because the non-ordinary-world and ordinary-world evaluations are two distinct inquiries does not imply that their methodological necessities are equally correct. If S's R for OP (or not-SH) is both sensitive in the ordinary context and insensitive in the skeptical context, it must be that the methodological necessities of the ordinary-world and non-ordinary-world evaluations are equally correct—S's R for OP and not-SH being sensitive (insensitive) in the ordinary (skeptical) context is the direct result of the ordinary-world (non-ordinary-world) evaluation; if the latter is correct, so is the former.

A series of questions naturally arise: What do we mean by saying that a methodological necessity is correct? Under what kind of context do we challenge the methodological necessities of an inquiry? When we reject the methodological necessities of astrology, on what ground our rejections consists in? To answer such questions, it is important to notice that inquires can come in a variety of "levels". When I inquire into whether *Smith* killed Kenny and you inquire into whether *Mary* killed Kenny, it seems that, in one sense, we are engaging distinct inquiries. But, in another sense, we are also

engaging in the same inquiry, i.e., the inquiry of the murderer of Kenny. Moreover, the inquiry of the murderer of Kenny, in turns, falls into the category of empirical inquiry. Likewise, suppose I inquire into whether Plato is the teacher of Aristotle and you inquire into whether Socrates is the teacher of Plato. In one sense, we are engaging in different inquiries—my inquiry is about Plato, while yours, Socrates. On another sense, it seems that we are still engaging in the same kind of inquiry, i.e., historical inquiry. Historical inquiry, in turns, is also a kind of empirical inquiry. In brief, inquiries constitute a hierarch of kinds.

Astrology and astronomy fall into the category of empirical inquiry. Our judgment that the former's methodological necessities are incorrect is a judgment of empirical inquiry. Hence, a methodological necessity of an inquiry can be evaluated according to the standards of evaluation of a more general inquiry of which the original inquiry is a special case. Viewed in this light, the correctness of the methodological necessities of the non-ordinary-world (or ordinary-world) evaluation must be evaluated according to the standards of evaluation of a more general inquiry of which the non-ordinary-world (or ordinary-world) evaluation is a special case. Which general inquiry are we talking about here? My suggestion is that both the ordinary-world and non-ordinary-world evaluations are special cases of *epistemic inquiry*, and we should evaluate the correctness of methodological necessities according to the standards of evaluation of epistemic inquiry.

In what follows, I assume that epistemic inquiry aims at the goal of attaining truth and avoiding error. Other common epistemic virtues such as explanatory power,

simplicity, etc. will be regarded as subsidiary to these this epistemic goal.[96] Viewed in this light, my claim is that both the ordinary-world and non-ordinary-world evaluations are correct according to the standard of attaining truth and avoiding error.

Notice that, in general, the ordinary-world and non-ordinary-world evaluations are *empirical* inquiries, since they are *about the world* in the broadest sense. Empirical inquiry is a specific kind of epistemic inquiry—it is the kind of inquiry particularly concerned with attaining empirical truths and avoiding empirical errors. However, the term 'empirical' often has the implication of "related to or about the *ordinary* world-picture". In order to avoid being confused with the ordinary-world inquiry, I will continue using the more neutral term 'epistemic inquiry'.

7.9 Is the Non-Ordinary-World Evaluation Illegitimate Empirical Inquiry?

There has been a quite famous line of objection in the literature, which claims that the methodological necessities of the non-ordinary-world evaluation—particularly, the presupposition that the world may not be ordinary—are illegitimate from the perspective of empirical inquiry. This appears to be a prominent theme of Wittgenstein's *On Certainty*.[97] In this section, I will argue that this objection is in fact mistaken.

When talking about empirical inquiry, Wittgenstein writes:

> 94. But I did not get my picture of the world by satisfying myself of its correctness; nor do I have it because I am satisfied of its correctness. No: it is the inherited background against which I distinguish between true and false.

---

[96] At least this is the orthodox view (cf., e.g., Annis 1978; Alston 1989; David 2001; David 2005). Recent literatures, however, have seen an immerging tendency to reject the orthodox view (cf., e.g., Kvanvig 2003; Riggs 2002; Riggs 2003). In my opinion, the orthodox view is correct, but I will not deal with this issue here.

[97] Here, my main interest is not exegetic. I will be happy to leave it open whether my interpretation of Wittgenstein's view is correct or not.

95. The propositions describing this world-picture might be part of a kind of mythology. And their role is like that of rules of the game; and the game can be learned purely practically, without learning any explicit rule.

96. It might be imagined that some propositions, of the form of empirical propositions, were hardened and functioned as channels for such empirical propositions as were not hardened but fluid; and that this relation altered with time, in that fluid propositions hardened, and hard ones become fluid.

The idea, I take it, is that propositions constitute our world-picture are the "rules of the game", a notion roughly equivalent to Williams's methodological necessities.[98] The world-picture Wittgenstein has in mind, I take it, is the *ordinary* world-picture. We can only talk about (empirical) truths and falsities against the background of the ordinary world-picture. Empirical inquiry essentially consists in the ordinary world-picture. As a result, to challenge this picture is (empirically) illegitimate:

372. Only in certain cases is it possible to make an investigation "is that really a hand?" (or "my hand"). For "I doubt whether that is really my (or a) hand" make no sense without some more precise determination. One cannot tell from these words alone whether any doubt at all is meant—nor what kind of doubt.

To question whether this is a hand (or whether one has hands) cannot be a legitimate doubt unless in some very specific circumstances such as one is anesthetized and cannot feel one's hands. But if one tries to raise the philosophical question whether one has hands or not, one fails to raise a genuine doubt.

It is unclear how broad Wittgenstein wants to interpret the term 'empirical' in §95. If all he means by 'empirical' inquiry is the inquiry based on the ordinary world-picture, while admitting that there are broader kinds of inquiries based on the non-ordinary-world picture, then I think most of the aforementioned claims may well turn out

---

[98] Wright calls such methodological necessities 'cornerstones' for our ordinary beliefs about the external world. Wright points out that "it is a cornerstone for a large class of our beliefs that we are not cognitively disabled or detached from reality in a certain way" (Wright 2004, 168).

to be true. That is, we can only distinguish such empirical truths and falsities against the background of the ordinary world-picture, and to question the ordinary world-picture is illegitimate within such an empirical inquiry. Under this interpretation, Wittgenstein is arguing for the idea that philosophical (skeptical) inquiry, which is based on the non-ordinary world-picture, cannot be legitimately pursued within the inquiry based on the ordinary world-picture. This appears to be trivially true. The problem is that, while this last claim is very plausible, it does not, strictly, reject skepticism, since the skeptic's challenge does not focus on the ordinary world-picture (in other words, the skeptic does not engage in such an empirical inquiry).

In order to have an impact on skepticism, Wittgenstein has to interpret 'empirical' inquiry very broadly to the extent that it is equivalent to what we here call 'epistemic inquiry', viz. the general inquiry governed by the epistemic goals of which the ordinary-world and non-ordinary-world inquiries are instances. Under this interpretation, what Wittgenstein is suggesting is that the ordinary world-picture is a methodological necessity of epistemic inquiry, the rule of the game that renders epistemic inquiry possible. Understood in this way, one does not bear any epistemic relation to the ordinary-world picture. Rather, the ordinary world-picture must be in play in order for one to bear any epistemic relation to a proposition (cf. Coliva 2010).

If the ordinary world-picture is constitutive of epistemic inquiry, it cannot be challenged within epistemic inquiry. It follows that the non-ordinary-world evaluation, if it is to be a kind of epistemic inquiry at all, is illegitimate: It tries to question things that cannot be questioned. If the non-ordinary-world inquiry is illegitimate, then the skeptic is wrong in claiming that S does not know that OP (or not-SH), since the skeptic's

conclusion is derived from the non-ordinary-world evaluation. Thus, the problem of the non-ordinary-world evaluation is analogous to the problem of the astrological inquiry: They both contain methodological necessities that can be shown as incorrect when evaluated from a more general kind of inquiry.

The idea that the ordinary world-picture is constitutive of epistemic inquiry, however, is implausible. On the one hand, epistemic inquiry is governed by the epistemic goal, which, as I see it, is the goal of attaining truth and avoiding falsehood. To say that the ordinary world-picture is a methodological necessity, hence, is tantamount to saying that the goal of attaining truth and avoiding falsehood does not apply to the picture itself. This, in turn, implies that we cannot talk about the truth or falsehood of the ordinary world-picture, which is absurd. The ordinary world-picture is an empirical hypothesis (with 'empirical' interpreted very broadly) that either correctly represents the world or not. To say that, in principle, the ordinary world-picture, as an empirical hypothesis, cannot be challenged within empirical inquiry seems implausible.[99]

---

[99] There is a similar objection in close vicinity. That is, many have argued that the skeptic's denials of knowledge are never legitimate since they rely on "hyperbolic doubts" that are "systematically uneliminable" (Fogelin 1994, 91). Skeptical hypotheses are often so designed such that they are immune to rejection through sensory experiences or testimony. The hypotheses that I am not a handless BIV or that others do not have minds are immune to rejection through sensory experiences or testimony—if I were a BIV, I would still have the same sensory experience, and if others did not have minds, they would still tell me that they do, etc. Some have argued that hypotheses that are systematically uneliminable are illegitimate. Skepticism raises the questions that I might be a handless BIV or that others might not have minds. But these are not genuine questions, since they cannot possibly be answered. "If a question can be framed at all," as Wittgenstein puts it, "it is also *possible* to answer it" (Wittgenstein 1922, sec. 6.5 also see Fogelin (1994, 91)). A real doubt, we are told, has to be "occasioned by a real life situation" (Annis 1978, 214).

I agree that skeptical hypotheses are 'hyperbolic doubts', so to speak. But I think this only justifies taking the ordinary world-picture as a methodological necessity of the ordinary-world evaluation. The ordinary-world evaluation considers whether one's reason for p is sensitive from the point of view of the ordinary world-picture. There are many practical reasons for engaging in the ordinary-world inquiry. And, from the perspective of the ordinary-world inquiry, those hyperbolic doubts are not legitimate doubts in the first place. Yet, all these are consistent with the idea that idea that those hyperbolic doubts are legitimate doubts in yet another (perhaps more general) kind of inquiry.

On the other hand, if the ordinary world-picture is a methodological necessity of epistemic inquiry, we do not (have to) bear any epistemic relation to it. But intuitively it does seem that we bear (or at least can bear) a certain epistemic relation to the ordinary world-picture. For instance, we generally think that we can have good (or bad) reasons for believing the ordinary world-picture. In fact, I think Wittgenstein will admit that, in some extreme circumstances, one may need to have reason to accept the ordinary world-picture—imagine someone living in a world where many brains are envatted.

7.10 Why the Ordinary-World and Non-Ordinary-World Evaluations Are Both Legitimate Epistemic Inquiries

If what is said in Section 7.9 is correct, we have no reason to take the ordinary world-picture as a methodological necessity of epistemic inquiry. In this section, I want to further argue that from the point of view of epistemic inquiry, both the ordinary-world and non-ordinary-world evaluation are legitimate. That is, I want to argue that both the presuppositions of the ordinary world-picture—a methodological necessity of the ordinary-world evaluation—and the non-ordinary world-picture—a methodological necessity of the non-ordinary-world evaluation—are correct according to epistemic inquiry.

To see this, we need to be clear about the standards of evaluation of epistemic inquiry. As stated, epistemic inquiry aims at the goal of attaining truth and avoiding error. It has been widely accepted, however, that obtaining truth and avoiding error are two different goals that can pull us to different directions. As Williams James expresses in the following famous paragraph:

258

> There are two ways of looking at our duty in the matter of opinion—ways entirely different, and yet ways about whose difference the theory of knowledge seems hitherto to have shown very little concern. We must know the truth; and we must avoid error—these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws. … [H]e who says, "Better go without belief forever than believe a lie!" merely shows his own preponderant private horror of becoming a dupe.[100]

If one merely wants to attain truth, with no intention to avoid error, one of the best strategies is to accept every p, for doing so will guarantee that all truths are obtained. If one merely wants to avoid error, with no intention to attain truth, one of the best strategies is to refuse to accept any p, for doing so will guarantee that no error will be committed. Neither strategy, however, seems promising: Only a dupe accepts everything, and only a skeptic accepts nothing.

The moral is that both epistemic goals are indispensable.[101] The challenge is to find a way to fulfill these two goals that somehow conciliates the inherent tension between them. However, the tension between these two goals runs deep. In some cases, the more truths one obtains, the more risk one takes in gaining falsehoods, and, by the same token, the more falsehoods one avoids, the more truth one misses—assuming that

---

[100] See James (1912). The quotation is from Riggs (2003, 342).

[101] David has some qualms about this argument. He complains:
> It is said sometimes that the goal has to have both parts because taken separately they could be "trivialized": Believing everything would trivially satisfy the positive part; believing nothing would trivially satisfy the negative part. This strikes me as a curious consideration. Surely, believing everything is humanly impossible; it may even be impossible *tout court*, for it involves believing explicitly contradictory propositions. Moreover, I find even believing three things very hard, when one of them is, say, the proposition that I have no head. Although believing nothing is not quite as impossible as believing everything, those of us who are actually able to have beliefs (those of us who are not brain-dead, who are not fetuses, who are not trees) will not find it much easier. In any case, it is not clear what the consideration is supposed to show. (David 2001, 158)

The question whether it is psychologically possible to believe everything or to believe nothing at all is an interesting one. But for the present purposes, the question is a red herring. On the one hand, the present account is a normative discussion of the notion of sensitivity, and it is not uncommon for normative account to appeal to idealization of some sorts. On the other hand, David's concerns are off the mark since the attitude of believing is even not necessary for the formulation of epistemic goals. For instance, we can just apply to the attitude of 'acceptance' described in Wright (2004), which does not imply believing.

one is not omniscient and infallible. In general, one cannot set both as primary goals and try to achieve them at once.[102] If attaining truth is our primary goal, we should primarily tend towards acceptance and take the risk; if avoiding error is our primary goal, we should primarily tend towards withholding judgment and be cautious. If our primary goals are both attaining truth and avoiding error, we are torn. The only way out is to set one as the primary goal, while the other, secondary. Consider an analogy. We want to punish guilty people and avoid punishing innocent ones. If we want only to punish the guilties, one of the best strategies is to punish everybody. In this case, all guilty people will definitely be punished. If we want only to avoid punish the innocents, one of the best strategies is to punish no one. In this case, no innocent person will be punished. The problem, of course, is that we want them both. The tension between these two goals also runs deep. Given that we are not omniscient and infallible, the more guilty people we punish, the more innocent people we punish. Likewise, the fewer innocent people being punished, the fewer guilty people being punished. We can, however, take one as the primary goal, while the other, secondary. If we take punishing guilty people as the primary goal, while avoiding punishing innocent ones secondary, we should assume everybody to be guilty until proven innocent. Conversely, if avoiding punishing innocent people is the primary goal, while punishing guilty people the secondary goal, we should assume everybody to be innocent until proven guilty.

---

[102] An Example: Suppose that your perception is pretty reliable in recognizing objects of type-A, say, for every 100 objects of type-A, you will be able to recognize 95 of them as type-A objects. Now, if your primary goal is to recognize as many type-A objects as you can, then your goal compromise the goal of avoiding mistaken recognition of type-A objects, since the more type-A objects you recognize, the more non-type-A objects you mistakenly recognize as type-A objects. In brief, it is impossible to set both the goal of recognizing as many type-A objects as possible and the goal of avoiding mistaken recognition of type-A object as much as possible as one's primary goals. The achievement of one essentially compromises the achievement of the other.

By the same token, there are two ways one can satisfy the goals of attaining truth and avoiding error. On the one hand, if attaining truth is the primary goal and avoiding error secondary, one should accept p unless one has reason against it. On the other hand, if avoiding error is the primary goal and attaining truth secondary, one should not accept p unless one has reason for p. Let us call the former the *conservative goal*, and the latter, the *liberal goal*.[103]

The epistemic goal of attaining truth and avoiding error, thus, consist of two distinct and indispensible sub-goals. What is important, for the present purposes, is that each of these sub-goals explains why both the non-ordinary-world and ordinary-world evaluations are legitimate epistemic inquiries.

First, we need to notice that the hypotheses of the ordinary world-picture and non-ordinary world-picture are very special. As skeptical hypotheses show, the non-ordinary world-picture is compatible with all of our experiences, rendering the latter (alone) incapable of falsifying the former. For instance, the hypothesis that one is a BIV being stimulated to have the experiences of living in the real world is totally compatible with

---

[103] The terms 'conservative' and 'liberal' are from Pryor (2004). Pryor points out that a theory of justification can assign a hypothesis H two different roles:

> A theory treats H **conservatively** when it says that you need some justification to believe that H *in order to* have a given kind of prima facie justification to believe P. That is, the condition M that *make you* have that prima facie justification include you having this justification to believe H. (Pryor 2004, 354)
>
> A theory that treat H **liberally** *denies* that having prima facie justification to believe P require you to have antecedent justification to believe H. But it does count not-H as an *undermining* hypothesis; evidence against H undermines your prima facie justification to believe P. (Pryor 2004, 354)

While it is not hard to find similarities between what Pryor calls 'conservative' and 'liberal' justification and what I call the 'conservative' and 'liberal' goal here, it should be noted that Pryor does not intend the distinction to play a role in explaining the methodological necessities of epistemic inquiry (for simplicity's sake, I will not further discuss Pryor's theory of justification here). Williams also has a similar distinction in mind when he is talking about two possible models of justification. Williams's account will be discussed in Section 7.11 when I reply to the third possible objection to my account.

one's experiences. If empiricism is right that all knowledge starts from experiences, one has no (empirical) reason against the non-ordinary world-picture.

By contrast, we generally lack a non-circular reason for the ordinary world-picture as well. Crispin Wright's reconstruction of the skeptical argument has brought out this point vividly:

> Let P be any proposition purporting to express a routine observation about my local perceptible environment—say: that I have two hands—and consider the following trinity:
>
> I   My current experience is in all respect as if P.
>
> II   P
>
> III There is a material world. (Wright 2004, 170)

Now ask yourselves: What are your reasons for III? It takes just a moderate amount of reflection to recognize that propositions like II are our only reasons for III. Moreover, I is arguably our only reason for II. The problem is that, seemingly, I cannot offer adequate support for II unless we have already assumed III. This shows that our reasons for III (if any) are essentially circular (Wright 2004, 170ff.).

The claims that the non-ordinary world-picture is immune to (empirical) objections, and that there is no non-circular reason for the ordinary world-picture are the two sides of the same coin. They depend on skeptical considerations, and give rise to skeptical conclusions. Worried about skepticism, it is not surprising that some philosophers try to reject them. For instance, some had argued that ordinary perceptual beliefs in fact falsify the non-ordinary world-picture (cf. Moore 1962a; 1962b). Some argued that type-I propositions of Wright's argument indeed offer (immediate) justification for type-II propositions; without the help of type-III propositions, type-I

propositions themselves are enough to justify type-II propositions (cf. Pryor 2000; 2004). Others had argued that type-II propositions are not our only reasons for type-III propositions (cf. Wright 2004).

To offer a thorough discussion of this issue, however, goes beyond the scope of this dissertation. For simplicity's sake, I will just say this: The main motivation against the ideas that the non-ordinary world-picture is immune to objections, and that there is no non-circular reason for the ordinary world-picture is that they lead to skeptical conclusions, which threaten our ordinary knowledge attributions. However, it is mistaken to think that these two ideas threaten ordinary knowledge attributions, or so I shall argue. At any rate, if the sensitivity-variantist solution is correct, one can grant that we do not have objections to the non-ordinary world-picture, and that we do not have non-circular reason for the ordinary world-picture without sacrificing the ordinary knowledge attributions at all (in fact, this is just the beauty of the pluralistic solution). If so, the incentives for the aforementioned accounts are undermined.

Hence, suppose that we have no objections to the non-ordinary world-picture, and that there is no non-circular reason for the ordinary world-picture. We are now in a position to see how both the presupposition of the ordinary world-picture and the presupposition of the non-ordinary world-picture satisfy the standards of evaluation of epistemic inquiry. More precisely, they satisfy the liberal goal and conservative goal respectively.

On the one hand, the presupposition that the ordinary-world picture may not hold, which is a methodological necessity of the non-ordinary-world evaluation, is correct according to the conservative goal, since we generally do not have any non-circular

263

reason to support the ordinary picture, though we certainly do not have good reason against it as well. On the other hand, the presupposition of the ordinary world-picture, which is a methodological necessity of the ordinary-world evaluation, is correct according to the liberal goal, since we generally do not have any good reason to deny the ordinary world-picture, even though we do not have any non-circular reason for it as well. As a result, both the non-ordinary-world and ordinary-world evaluations are legitimate epistemic inquiries.

To sum up, I have argued in the previous sections that the skeptic and the ordinary subject engaging in different kinds of inquiry. The skeptic engages in the non-ordinary-world evaluation whose methodological necessities do not include the presupposition that the ordinary world-picture holds, while the ordinary agent engages in the ordinary-world evaluation whose methodological necessities include the ordinary world-picture. Both the ordinary-world and non-ordinary-world evaluations are legitimate epistemic inquiry: The former satisfies the liberal goal, while the latter, the conservative goal. If what has been said is correct, we have established the following:

> 7.2 The skeptic engages in the non-ordinary-world evaluation according to which the counterfactual situation in play when evaluating the sensitivity of S's reason for p is the one in which SH is the case, while an ordinary subject engages in the ordinary-world evaluation according to which the counterfactual situation in play when evaluating the sensitivity of S's reason for p is the one in which not-SH is the case.

7.1 and 7.2 imply the sensitivity-variantist solution:

> 7.3 S's reason for OP (and not-SH) is sensitive according to the ordinary-world evaluation and fails to be the case according to the non-ordinary-world evaluation.

On the sensitive-variantist solution, S's standing on knowing dimensions with respect to op (and not-sh) varies with the skeptic and ordinary contexts, since S's standing on

knowing dimension is characterized in terms of the sensitivity of S's reason. S does not know that op (and not-sh) in the skeptical context since his standing on knowing dimensions with respect to op (and not-sh) fails the satisfies the standards for knowledge in play, while S knows that op (and not-sh) in the ordinary context since his standing on knowing dimension with respect to op (and not-sh) satisfies the standards for knowledge. This explains why the skeptical puzzle arises in the first place: Our intuitions that "S knows that op" (i.e., 1.1) is true are not mistaken since the statement is indeed true in the ordinary context; our intuitions that "S does not know that not-sh" (i.e., 1.2) is true are also not mistaken since the statement is also true in the skeptical context; our intuitions that "S knows that op only if S knows that not-sh" (i.e., 1.3) is true are correct since the statement is true in all contexts. Moreover, the skeptical puzzle does not constitute a genuine paradox since in no context 1.1-1.3 are true altogether.

The skeptical puzzle arises from the fact that we fail to see that 1.1 and 1.2 are true in different contexts, or not-1.1 and not-1.2 are also true in some contexts. Once we see that 1.1 and 1.2 are not true in the same context, the sense of paradox regarding 1.1-1.3 disappears.

7.11 Possible Objections and Replies

In this section, I will consider and reply to three possible objections to the sensitivity-variantist solution.

Objection 1: The sensitivity-variantist solution is too concessive.

The sensitivity-variantist solution claims that the skeptic is *right* in her own context. To some, this appears to be too concessive to skepticism. Skepticism, some might argue, cannot be correct. The task of epistemology is to reveal its flaw, not to surrender to it. However, it is not clear to me that we should always presuppose that the skeptic could not be right in the first place. As far as I can tell, there is no principled reason for this presupposition.

A source of the general worry about skepticism seems to be that if skepticism is right, then ordinary knowledge will be threatened. As Hume has famously argues, there is no satisfactory theoretical response to skepticism. Skepticism is theoretically correct, but practically trivial. The only cure of skepticism is to stop theoretical thinking and focus on practical considerations instead. This worry, however, does not bother the pluralistic solution. The pluralistic solution claims that both skepticism and the ordinary notion of knowledge are right in a certain context. That said, the ordinary notion of knowledge is theoretically correct in the ordinary context.

Perhaps, the general worry is that skepticism could not be right since it is *absurd*. Skepticism claims that we do not know, say, that we have hands, and this is at any rate unacceptable. Well, it is unacceptable from the point of view of the ordinary notion of knowledge. The sense of absurdity is understandable since most of us engage in the ordinary-world evaluation most of the time. "S knows that op" is obviously true in the ordinary context, and its denial is absurd in this context. However, the fact that "S does not know that op" is absurd in the ordinary context is not a reason against the claim that the same knowledge attribution is not true in the skeptical context. The skeptic context

can easily be motivated by considering the truth value of, say, "S knows that not-sh". It is also "natural" to think that "S does not know that op" in the skeptical context.

Could it be that skepticism could not be right since it is trivial? It is easy to realize that skepticism does not have any practical merits. Some would even go so far to say that skepticism is theoretically insignificant. I am inclined to think that skepticism is not theoretically insignificant, since I think we can learn something important about the general structure of knowledge and justification from skeptical arguments. But even if skepticism is indeed insignificant, both practically and theoretically, still, this does not show that it is incorrect. We have to distinguish being trivial from being wrong. A theory can be trivial and yet still be right. The pluralistic solution does not imply that skepticism is (practically) significant. Accepting this view does not imply that you have to take the skeptical puzzle seriously.

Objection 2: The sensitivity-variantist solution implies implausible error theory.
The pluralistic account to the skeptical puzzle has it that 1.1 and 1.2 are true *only* in different contexts. This claim explains why 1.1 and 1.2 are initially plausible and avoids taking 1.1-1.3 to constitute a genuine paradox. However, it seems that we normally do not have explicit recognition that 1.1 and 1.2 are true in different contexts, and that we have the tendency to regard the truth of 1.1 and 1.2 as leading to contradiction—this is why 1.1-1.3 appear paradoxical in the first place. If the pluralistic account is right, we have mistakenly regarded 1.1 and 1.2 as being true in the same context, or at least failed to recognize that 1.1 and 1.2 are true in different contexts. In other words, the pluralistic account seems to imply a certain error theory—that we systematically fail to grasp the

truth values of 1.1 and 1.2. Many, however, think that this kind of error theory is implausible.

The sensitivity-variantist solution is a kind of pluralistic solution. As a result, it also appeals to a certain error theory in accounting for the facts that we do not normally take 1.1 and 1.2 to be true in different contexts, and that 1.1 and 1.2 seem to lead to contradiction. The idea, roughly, is that we somehow fail to recognize that the skeptic and the ordinary subject are engaging in different inquiries such that different counterfactual situations are in play in evaluating S's epistemic position.

This kind of error theory, however, is not totally unacceptable. In general, it is not always implausible to reject our intuitions, since our intuitions are not infallible, but vague and imprecise in many cases. In Chapter Six, I have shown that our ordinary intuitions regarding counterfactualization fail to distinguish, at least not explicitly, the distinction between inferential and interventional counterfactualization. We evaluate counterfactual conditionals such as "If Nancy's grandfather were not born in 1945, she would still have his testimony that he was born in 1945" (Section 6.9) and "If the zebra were not available, there would not have been a zebra in the pen" (Section 6.7) without realizing, at least not explicitly, that our evaluations are driven by different kinds of counterfactualization. That our ordinary intuitions conflate, or tend to conflate, different kinds of counterfactualization has given rise to the widespread, yet mistaken, idea that the sensitivity account is bounded to face a certain kind of counterexample (Section 6.9). Our intuitions can and should be corrected in some cases. Error theory is not always an implausible option.[104]

---

[104] As Daly and Liggins (2010) point out, philosophers tend to underestimate the prospects for error theory.

The sensitivity-variantist solution claims that we tend to ignore the fact that the skeptic and the ordinary subject are engaging in different kinds of inquiries. Perhaps this claim is not totally unacceptable after reflection. In my opinion, ordinary, not philosophically trained, people tend to grant the skeptic that we do not know that we are not a BIV—a common reaction: How could you *know* this kind of thing? The underlying reason that drives this kind of reaction is the consideration of sensitivity: If I were to be a BIV, I would still have all these sensory experiences and would still be led to believe that I am not a BIV. But ordinary people also balk at full-blooded skepticism. They do not want to grant that they do not know that they have hands—also a common reaction: How could I not know that these are hands? If pressed, many may grant that their "mindset" of claiming that they know that they have hands is different from the one of claiming that they do not know that they are not a BIV. The difference between "mindsets" is subtle, vague, and somehow hard to explain, but few, I suspect, will admit that they are downright inconsistent. The sensitivity-variantist solution tries to explain this puzzling phenomenon by suggesting that we are engaging in different inquiries when torn between the skeptic's and the ordinary subject's conclusions. We tend to agree with the skeptic in evaluating S's reason for not-sh in terms of the non-ordinary-world, hence engaging in the non-ordinary-world evaluation. We also tend to agree with the ordinary subject in evaluating S's reason for op in terms of the ordinary-world, hence finding ourselves doing the ordinary-world evaluation. We sometimes feel that the skeptic is changing the subject, or the skeptic and the ordinary subject talk pass each other, indicating that we are not totally blind to the fact that the skeptic and the ordinary subject do not engage in the same inquiry.

How empirically plausible this theory is is still an open question. So far, I have seen no conclusive evidence against it.

Objection 3: The skeptic has committed to implausible epistemological realism

When talking about the general structure of justification, Williams distinguishes two kinds of model, which he calls the 'prior grounding' model and the 'default and challenge' model. According to the former, "epistemically responsible believing is believing on the basis of adequate evidence" (Williams 2001, 36). According to the latter, "a person can believe responsibly without basing his beliefs self-consciously on evidence. Nevertheless, he commits himself to providing evidence, should the adequacy of his epistemic performance be challenged" (2001, 36). Williams goes on to point out that the main difference between these two models is that:

> The difference between the 'Prior Grounding' and 'Default and Challenge' conceptions of justification is like that between legal systems that treat the accused as guilty unless proved innocent and those that do the opposite, branding presumptive innocence and throwing the burden of proof onto the accuser. Adopting the second model epistemic entitlement is the default status of a person's beliefs and assertions. One is entitled to a belief or assertion (which, remember, is an implicit knowledge-claim, unless clearly qualified) in the absence of appropriate 'defeaters': that is, reasons to think that one is *not* so entitled. (Williams 2001, 149)

So far, it seems that what Williams has said is in line with the distinction between the conservative and liberal goals discussed in Section 7.10. However, unlike the present account, Williams rejects the prior grounding model. He argues that this model presupposes the implausible thesis:

> *Claimant-Challenger Asymmetry*. Whenever knowledge is claimed, the burden of justification lies with the claimant. If I represent myself as knowing that P, I invite you to ask me how I know. There is nothing you have to do, or no way that things have to be, in order for you to have the right to enter a challenge. (2007, 99)

Assuming claimant-challenger asymmetry, "a request to back up a belief or assertion need no justification: in conceding an unrestricted commitment to produce grounds, the claimant grants the sceptic's entitlement to request them. The sceptic acquires the right to issue naked challenges" (2001, 150). However, naked challenges, according to Williams, are epistemically illegitimate. One does not raise a genuine challenge if one has no justification for one's question. Questioning, like asserting, is subject to the justification requirement. It is because the skeptic's question of the ordinary world-picture is not justified that skeptical challenges are is not legitimate. Williams thus claims that we should reject skepticism on the basis of the denial of claimant-challenger asymmetry.

There are some obvious problems for objections along this line, though. Firstly, it is not the case that the prior grounding model necessarily implies claimant-challenger asymmetry. In general, if one makes a claim of knowledge, either one has justification for one's claim or one does not. It seems that if one does not have justification for one's claim, others are thus having good reasons to challenge one's claim. Viewed in this light, the skeptic's challenge of one's claim of knowledge can be seen as the question as to whether one makes a claim of knowledge without any justification. The skeptic is not making naked challenges—the skeptic just want to make sure that you are justified in attributing knowledge.

Secondly, even if the prior grounding model does imply claimant-challenger asymmetry, it is not clear that we should thus reject the model. As should be noted, claimant-challenger asymmetry is indeed legitimate in many ordinary cases. For instance, suppose that Kenny was murdered, that all evidence strongly suggests that Smith is the murderer, and that no evidence suggests otherwise. Yet, before concluding that Smith is

271

the murderer, the detector considers whether Smith is indeed innocent or not. In this case, it seems that the detector's consideration is legitimate even if he does not have particular reason for the hypothesis that Smith is innocent.

While Williams does not deal with the first problem, he does have something to say about the latter. According to Williams, claimant-challenger asymmetry is an implausible thesis since it, eventually, presupposes epistemic realism that postulate that ordinary knowledge constitutes "a natural epistemological kinds", the idea that "beliefs can be sorted by some broad features of their propositional content into epistemic kinds ordered by relative epistemic priority" (Williams 2013, 26). Elsewhere, Williams also call this the 'totality condition'—the totality condition assumes "'knowledge of the world' picks out a natural or theoretically significant (natural) kind of knowledge, the sort of thing of which we might expect to give a unified account, and not just an aggregate of more or less unrelated cases" (Williams 1988, 423–4). Williams argues against totality condition and thus epistemic realism. Knowledge, according to Williams, does not form a natural kind (or a theoretically interesting domain of discourse). Hence, so long as skepticism presupposes epistemic realism (or imposes totality condition), it should be rejected.

Now, it is all too controversial whether the Totality Condition and/or epistemic realism are incorrect or not. To deal with this issue will go beyond the scope of the present study. However, it is enough to point out that, according to the sensitivity-variantist solution, the skeptic does not need to appeal to the Totality Condition. Nor does the skeptic presuppose epistemic realism. What the skeptic tries to do, on this account, is to evaluate S's reason for p from the point of view of the non-ordinary world-picture. The

skeptic is engaging in the non-ordinary-world evaluation, and this inquiry does not require the skeptic to deal with an abstract object called 'the knowledge of the world as such'—the skeptic is interested in the sensitivity of S's reason for p (where p stands for concrete propositions) (also see Greco 2000).

Surely, one may wonder whether the skeptical evaluation is worth considering (I suspect most will think that it is not). One may even argue that the skeptical evaluation is illegitimate. But the latter needs further arguments. The failure of the Totality Condition (or epistemological realism), if at all, is not enough to undermine the legitimacy of the non-ordinary-world evaluation.

7.12 Taking Stock

In this chapter, I argue for a kind of standing-variantist solution to the skeptical puzzle, which claims that:

> 7.3   S's reason for OP (and not-SH) is sensitive according to the ordinary-world evaluation and fails to be the case according to the non-ordinary-world evaluation.

In order to support 7.3, I argue for to the following premises:

> 7.1   S's reason for OP (or not-SH) is sensitive when evaluated in terms of the counterfactual situation in which not-SH is the case, but fails to be the case when evaluated in terms of the counterfactual situation in which SH is the case.

> 7.2   The skeptic engages in the non-ordinary-world evaluation according to which the counterfactual situation in play when evaluating the sensitivity of S's reason for p is the one in which SH is the case, while the ordinary subject engages in the ordinary-world evaluation according to which the counterfactual situation in play when evaluating the sensitivity of S's reason for p is the one in which not-SH is the case.

The argument 7.1-7.3 is valid. The task is to show that 7.1 and 7.2 are true.

With respect to 7.1, I show that the following hold:

7.1.3   S's reason R for OP is sensitive when evaluated from the counterfactual situation in which not-SH is the case, and the same reason R for not-SH is insensitive if evaluated from the counterfactual situation in which SH is the case.

7.1.4   S's reason R for not-SH is (trivially) sensitive when evaluated in terms of the counterfactual situation in which not-SH is the case, and the same reason R for OP is (trivially) insensitive if evaluated in terms of the counterfactual situation in which SH is the case.

7.1.3 is supported by:

SSC   S's reason for p is sensitive if and only if "not-$R_S(p)$" is true in M_{not-p}.

To argue for 7.1.4, I argue that we should generalize SSC such that we can talk about a counterfactual conditional being true (or false) in terms of the counterfactual situation in play. That is:

GSC   S's reason R for p is sensitive in M if and only if "not-p $\supset$ not-$R_S(p)$" is true in M',

where M' is the causal model representing the counterfactual situation *in play* (M' may not be M_{not-p}). 7.1.4 is true according to GSC.

With respect to 7.2, I argue that the counterfactual situation in play when evaluating the sensitivity of one's reason is determined by the epistemic inquiry one is engaging in. That is:

The inquiry-dependence thesis. The epistemic inquiry in which one is engaging determines the counterfactual situation in play when one evaluates the sensitivity of S's reason for p.

I then argue that the skeptic and the ordinary subject engage in different kinds of inquiries: The skeptic engages in the non-ordinary-world evaluation, while the ordinary subject, the ordinary-world evaluation. The counterfactual situation in play in the non-

274

ordinary-world evaluation is the one in which not-SH is not presupposed; the counterfactual situation in play in the ordinary-world evaluation, by contrast, is the one in which not-SH is presupposed. The difference between these two kinds of inquiries thus supports 7.2.

Just because the skeptic and the ordinary subject are engaging in different inquiries does not mean that both of their epistemic appraisals are correct. I argue that both appraisals are correct in that each of them satisfies a sub-goal of epistemic inquiry.

Appendix 1:

Stakes-Shifting Cases Explained

In Chapter Two, I argue that, despite widely accepted, the idea that stakes-shifting cases offer the best kind of evidence for standards variantism is essentially flawed. I have dealt with three prominent types of standards variantism—the linear account, the spherical account, and the particular view—in depth, and found that all are untenable. In this appendix, I will show that the sensitivity variantism developed in this dissertation in fact offers a quite satisfactory account to stakes-shifting cases.

Recall that stakes-shifting cases are (usually) a pair of cases such that (a) the same knowledge-attributing sentence, say, 'S knows that p', seems true when uttered in LOW but not in HIGH, and that (b) LOW and HIGH are similar in all significant aspects except that S's (or the speaker's, or the evaluator's) stakes in whether p are high in HIGH but low in LOW.

The presentation of stakes-shifting cases normally starts from LOW. Conceiving LOW, we are inclined to agree that S knows that p. HIGH is then introduced, and a certain hypothesis q (which is incompatible with certain of S's *presuppositions*) is raised, and S is explicitly (sometimes implicitly) specified to be unable to rule out q (or S's reason does not favor p over q to a satisfactory degree). Consider the following case again:

> *Airport.* Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the

flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes, I know—it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen 1999, 58)

Smith's context is LOW, while Mary and John's, HIGH. When we first consider Smith's situation, we tend to think that Smith knows that the flight stops in Chicago. Then we are introduced to HIGH—that Mary and John have a very important business contact—and the hypothesis that the itinerary contains a misprint is raised (Smith presumably presupposes that the itinerary does not contain a misprint). Implicitly, Smith is supposed to be unable to rule out *The itinerary contain a misprint* (or Smith's total reason does not favor *The itinerary does not contain a misprint* over *The itinerary contain a misprint* to a satisfactory degree).

Let us explain this dialectical structure of *Airport* in terms of the tracking-based account. At first, when consider LOW, we construct a causal model L that represents it. Naturally, the set of variables *V* of L consists in the following variables:

EXOGENOUS represents the factors that determine the flight to be schedule to stop in Chicago.

SCHEDULE represents whether the flight is scheduled to stop in Chicago or not.

ITINERARY represents whether the itinerary prints that the flight stops in Chicago or not.

READING represents whether Smith has read that the flight stops in Chicago or not.

Intuitively, at least when considering LOW, we tend to think that whether the itinerary prints that the flight stops in Chicago or not depends on whether the flight is in fact scheduled to stop in Chicago or not. Moreover, whether Smith has read that the flight

stops in Chicago or not depends on whether the itinerary actually prints that the flight

stops in Chicago or not. Hence, the structural equations S of L seem to be as follows (see

Figure A.1):

SCHEDULE = EXOGENOUS
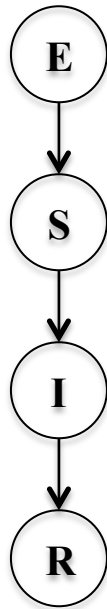
ITINERARY = SCHEDULE
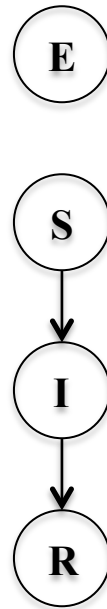
READING = ITINERARY

Figure A.1
DAG of L

Figure A.2
DAG of
$L_{SCHEDULE=0}$

According to *Airport*, we know that the flight is actually scheduled to stop in Chicago,

and the itinerary contains correct information of the flight schedule and Smith has correct

read the itinerary. That is, the value assignment A of L is as follows:

A(EXOGENOUS) = A(SCHEDULE) = A(ITINERARY) = A(READING) = 1.

It seems plausible to think that L represents our understanding of *Airport*.

Moreover, it seems that Smith's sole reason for *The flight stops in Chicago* is his reading

of the itinerary. If so, it is not surprising that we tend to think that Smith knows that the flight stops in Chicago, since Smith's reason, arguably, is sensitive. That is, if the flight were not schedule to stop in Chicago, then Smith would not have read that the flight stops in Chicago.

To see this, let us first specify the counterfactual situation in play in evaluating Smith's reason. Naturally, the counterfactual situation in play is the one in which the flight was not scheduled to stop in Chicago—we are inquiring whether Smith's would have the same reason if the flight was not scheduled to stop in Chicago. To consider what would have happened in such a situation, we intervene in L in terms of SCHEDULE = 0. The set of structural equations of $L_{SCHEDULE=0}$ consists of (see Figure A.2) :

SCHEDULE = 0

ITINERARY = SCHEDULE

READING = ITINERARY

The value assignment $A_{SCHEDULE=0}$ is:

A(SCHEDULE) = A(ITINERARY) = A(READING) = 0, and

A(EXOGENOUS) = 1.

Obviously, "SCHEDULE = 0 ⊃ READING = 0" is true in $L_{SCHEDULE=0}$, and as a result, "SCHEDULE = 0 > READING = 0" is true in L, since $L_{SCHEDULE=0}$ is the counterfactual situation in play.

Since we naturally think that Smith's reason for *The flight stops in Chicago* is sensitive, it is thus natural that we intuitively think that Smith knows that the flight stops in Chicago.

Now, consider Mary's challenge. Mary argues that it is possible that the itinerary contained a misprint such that while the itinerary printed that the flight stops in Chicago, the flight was not actually scheduled to stop in Chicago. What Mary is doing, at its core, is challenging the presupposition of our original representation of Smith's situation. When considering LOW, we represent Smith's situation as L. In particular, L presuppose that whether the itinerary prints that the flight stops in Chicago depends (only) on whether the flight is scheduled to stop in Chicago (i.e., ITINERARY = SCHEDULE). Mary is challenging this presupposition, claiming that whether the itinerary prints that the flight stops in Chicago or not depends not only on the actual schedule of the flight, but also on whether the itinerary contains a certain misprint or not.

In other words, Mary suggests that *an alternative causal model H should represent Smith's situation*. For the sake of discussion, suppose that Mary claims that the computer according to which Smith's itinerary was printed may contain incorrect information about the flight schedule. Hence, H consists of the following set of variables *V'*:

EXOGENOUS-**1** represents events that determine the flight to be scheduled to stop in Chicago.

EXOGENOUS-**2** represents events that determine the computer to have a piece of incorrect information about the flight schedule.

SCHEDULE represents whether the flight is scheduled to stop in Chicago or not.

COMPUTER represents whether the computer contains a piece of information about the flight such that regardless of whether the flight is scheduled to stop in Chicago or not, the computer will still list the flight as having a layover in Chicago.

ITINERARY represents whether the itinerary prints that the flight stops in Chicago or not.

READING represents whether Smith has read that the flight stops in Chicago or not.

It seems that Mary is suggesting that the structural equations S' of H should be specified as follows (see Figure A.3):

SCHEDULE = EXOGENOUS-1

COMPUTER = EXOGENOUS-2

ITINERARY = max(SCHEDULE, COMPUTER)

READING = ITINERARY

In words, Mary is suggesting that whether the itinerary will list that the flight stops in Chicago depends both on whether the flight is scheduled to stop in Chicago, and on whether the computer based on which the itinerary is printed contains incorrect information about the flight schedule. If flight is scheduled to stop in Chicago, then the itinerary will print that it does stop in Chicago. Similarly, if the computer contains incorrect information about the flight schedule, then regardless of whether the flight is scheduled to stop in Chicago, the itinerary will print that the flight will stop in Chicago.
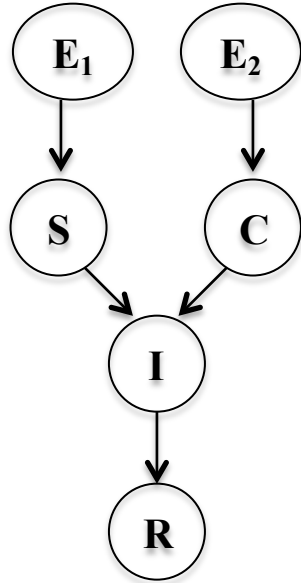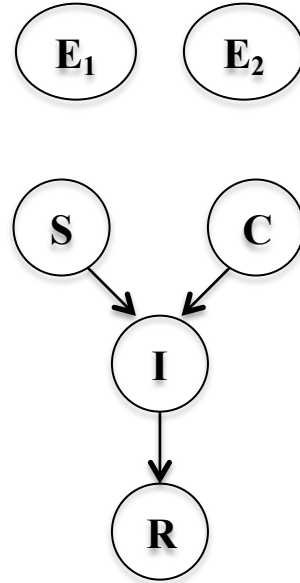
Figure A.3
DAG of H

Figure A.4
DAG of
$H_{\text{SCHEDULE}=0\&\text{COMPUTER}=1}$

We know that the flight was in fact scheduled to stops in Chicago, that the computer does not contain incorrect information about the flight schedule that the itinerary lists that the flight stops in Chicago, and that Smith has read the itinerary correctly. That is, the value assignment A' of H is as follows:

A'(EXOGENOUS-2) = A'(COMPUTER) = 0, and

A'(EXOGENOUS-1) = A'(SCHEDULE) = A'(ITINERARY) = A'(READING) = 1.

Mary is suggesting that Smith's situation should be represented by H as opposed to L. Mary argues that Smith's reason for *The flight stops in Chicago* is in fact insensitive, since if the flight was not schedule to stop in Chicago, but the computer contains incorrect information regarding the flight schedule, then Smith would still have read that the flight stops in Chicago. The counterfactual situation in play in evaluating Smith's reason, according to Mary, is the one in which the flight was in fact not schedule

to stop in Chicago, and the computer contained incorrect information about the flight schedule. This counterfactual situation can be represented by the submodel $H_{SCHEDULE=0\&COMPUTER=1}$, whose set of structural equations $S'_{SCHEDULE=0\&COMPUTER=1}$ is as follows (see Figure A.4):

SCHEDULE = 0

COMPUTER = 1

ITINERARY = max(SCHEDULE, COMPUTER)

READING = ITINERARY

The following is thus the value assignment $A'_{SCHEDULE=0\&COMPUTER=1}$:

$A'_{SCHEDULE=0\&COMPUTER=1}$(EXOGENOUS-2) =
$A'_{SCHEDULE=0\&COMPUTER=1}$(SCHEDULE) = 0, and

$A'_{SCHEDULE=0\&COMPUTER=1}$(EXOGENOUS-1) =
$A'_{SCHEDULE=0\&COMPUTER=1}$(COMPUTER) = $A'_{SCHEDULE=0\&COMPUTER=1}$(ITINERARY)
= $A'_{SCHEDULE=0\&COMPUTER=1}$(READING) = 1.

Whether Smith's reason for *The flight stops in Chicago* is sensitive or not depends on whether "SCHEDULE = 0 > READING = 0" is true in H or not. Given that $H_{SCHEDULE=0\&COMPUTER=1}$ represents the counterfactual situation in play, this amounts to the question whether "SCHEDULE = 0 ⊃ READING = 0" is true in $H_{SCHEDULE=0\&COMPUTER=1}$. Since "SCHEDULE = 0 ⊃ READING = 0" is obviously false in $H_{SCHEDULE=0\&COMPUTER=1}$, it is thus not surprising that if we agree with Mary that Smith's situation should be presented by H, then we intuitively think that Smith's reason for *The flight stops in Chicago* is sensitive. This further explains why we tend to think that Smith does not know that the flight stops in Chicago in HIGH.

The present explanation of stakes-shifting cases has an edge over the contextualist (or standard-variantist) explanation discussed in Chapter Two. I have pointed out that cases like the following have posted serious problems for the contextualist account (and standards variantism in general) (see Chapter Two):

> *Airport'*. Everything is like *Airport* except that Mary and John and Smith are friends and Mary and John knows that Smith is a Lakers fan. Suppose that Mary knows that the Lakers won yesterday's game, and that Smith believes truly that the Lakers won yesterday's game since she saw him reading today *The New York Times*. Suppose that, in LOW, Smith's standing on knowing dimensions with respect to *The flight stops in Chicago* is epistemically equivalent to his standing on knowing dimensions with respect to *The Lakers won yesterday's game*. Now, on their way to find the airline agent to check for the flight schedule, Mary and John are chatting. John asks Mary whether Smith knows that the Lakers have won yesterday's game.

Consider the following two possible situations:

> *Airport(a)*. Mary says to John, "Yes, Smith knows the Lakers won yesterday's game."

> *Airport(b)*. Mary says to John, "No, he doesn't. Smith believes truly that the Lakers won yesterday's game, but he does not know."

It seems that *Airport(a)*, rather than *Airport(b)*, is the natural reply in *Airport'*, and yet the contextualist account cannot account for this phenomenon. The problem, roughly, is that it is very hard for the contextualist account to predict that Smith knows the Lakers won yesterday's game while at the same time preserving the verdict that Smith does not know that the flight stops in Chicago, since Smith's standing on knowing dimensions with respect to *Lakers won yesterday's game* is on epistemic par with his standing with respect to *The flight stops in Chicago*, and if the former does not satisfy the standards for knowledge in play, so does the latter.

The present account can easily sidestep this problem. On this account, Smith does not know that the flight stops in Chicago since his standing on knowing dimensions with

respect to *The flight stops in Chicago* fails to satisfy the standards for knowledge in play. The present account can grant that Smith's standing with respect to *The Lakers won yesterday's game* satisfies the standards for knowledge in play. The evaluation of Smith's standing on *The flight stops in Chicago* is independent of the evaluation of Smith's standing on *The Lakers won yesterday's game*. (In terms of causal models, to determine the sensitivity of Smith's reason for *The flight stops in Chicago* and Smith's reason for *The Lakers won yesterday's game* requires two distinct causal models whose variables are not causally connected.)

Another piece of evidence that the present account is a better explanation is that the present account correctly predicts that, when LOW and HIGH are considered together, we tend to think that either "Smith knows that the flight stops in Chicago" is true in both LOW and HIGH, or it is false in both contexts. As I have noted (Section 2.4, footnote 27), the contextualist account fails to explain this phenomenon—the contextualist account predicts that we will tend to grant that "Smith knows that the flight stops in Chicago" is true in LOW and false in HIGH.

On the present account, when one considers the sensitivity of Smith's reason for *The flight stops in Chicago*, either one thinks that the counterfactual situation in which the itinerary contains a misprint is in play or one does not think so. If one is more concerned with the high-stakes situation Mary is in (as Mary herself clearly is), one is inclined to think that the counterfactual situation in which the itinerary contains a misprint is in play. Otherwise, one can just presuppose the counterfactual situation in play is the one in which the itinerary does not contain a misprint. What is at issue here is that one is facing two inquiries here: The normal inquiry—inquiry in which we tend to

engage when not facing high-stakes situation—according to which the itinerary is presupposed to be correct, and the inquiry, suggested by Mary, that does not take for granted that the itinerary is correct. Now, which inquiry one chooses to engage in may depend on a variety of factors, including the conversational topics, one's practical interests, etc. I do not have a complete account of how one's choice is affected by these various kinds of factors. Yet, this much seems clear: No matter which inquiry one decides to engage in, one will not think that "Smith knows that the flight stops in Chicago" is both true (in LOW) and false (in HIGH). Rather, one will either think that the statement is true (in LOW and HIGH) or that it is false (in LOW and HIGH). This seems to be the correct prediction.

Bibliography

Adams, Fred, and Murray Clarke. 2005. "Resurrecting the Tracking Theories." *Australasian Journal of Philosophy* 83 (2): 207–21.

Alston, William P. 1985. "Concepts of Epistemic Justification." *The Monist* 68 (1): 57–89.

———. 1989. *Epistemic Justification*. Ithaca: Cornell University Press.

Annis, David B. 1978. "A Contextualist Theory of Epistemic Justification." *American Philosophical Quarterly* 15 (3): 213–9.

Armstrong, David M. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.

Bennett, Jonathan. 1984. "Counterfactuals and Temporal Direction." *The Philosophical Review* 93 (1): 57–91.

———. 2003. *A Philosophical Guide to Conditionals*. Oxford: Clarendon Press.

Blaauw, Martijn. 2008. "Subject Sensitive Invariantism: In Memoriam." *The Philosophical Quarterly* 58 (231): 318–25.

Black, Tim. 2008. "Defending a Sensitive Neo-Moorean Invariantism." In *New Waves in Epistemology*, edited by Vincent F. Hendricks and Duncan Pritchard, 8–27. New York: Palgrave Macmillan.

Blome-Tillmann, Michael. 2007. "Contextualism and the Epistemological Enterprise." *Proceedings of the Aristotelian Society* 107: 387–94.

———. 2009a. "Knowledge and Presuppositions." *Mind* 118 (470): 241–94.

———. 2009b. "Contextualism, Subject-Sensitive Invariantism, and the Interaction of 'Knowledge'-Ascriptions with Modal and Temporal Operators." *Philosophy and Phenomenological Research* 79 (2): 315–31.

BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.

Briggs, Rachael. 2012. "Interventionist Counterfactuals." *Philosophical Studies* 160 (1) (August 1): 139–166.

Brogaard, Berit. 2008. "In Defence of a Perspectival Semantics for 'Know.'" *Australasian Journal of Philosophy* 86 (3): 439–59.

287

Cohen, Stewart. 1988. "How to Be a Fallabilist." *Philosophical Perspectives* 2: 91–123.

———. 1999. "Contextualism, Skepticism, and the Structure of Reason." *Philosophical Perspectives* 13: 57–89.

Coliva, Annalisa. 2010. *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense*. Palgrave Macmillan.

Daly, Chris, and David Liggins. 2010. "In Defence of Error Theory." *Philosophical Studies* 149: 209–30.

David, Marian. 2001. "Truth as the Epistemic Goal." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 151–69. New York: Oxford University Press.

———. 2005. "Truth as the Primary Epistemic Goal: A Working Hypothesis." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 296–312. Malden, MA: Blackwell.

DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52 (4): 913–29.

———. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104 (1): 1–52.

———. 2009. *The Case for Contextualism*. Oxford: Clarendon Press.

Descartes, Rene. 1984. *The Philosophical Writings of Descartes*. Translated by John Cottingham, Robert Stoothoff, and Dugald Murdoch. Vol. 2. Cambridge: Cambridge University Press.

Downing, P. B. 1958. "Subjunctive Conditionals, Time Order, and Causation." *Proceedings of the Aristotelian Society* 59 (January 1): 125–140.

Dretske, Fred. 1970. "Epistemic Operators." *The Journal of Philosophy* 67 (24): 1007–23.

———. 1971a. "Conclusive Reason." In *Perception, Knowledge, and Belief*, 3–29. New York: Cambridge University Press.

———. 1971b. "Two Conceptions of Knowledge: Rational vs. Reliable Belief." In *Perception, Knowledge, and Belief*, 80–93. New York: Cambridge University Press.

Edgington, Dorothy. 2004. "Counterfactuals and the Benefit of Hindsight." In *Cause and Chance: Causation in an Indeterministic World*, edited by Phil Dowe and Paul Noordhof, 12–27. New York: Routledge.

———. 2011. "Causation First: Why Causation Is Prior to Counterfactuals." In *Understanding Counterfactuals, Understanding Causation: Issues in Philosophy and Psychology*, edited by Christoph Hoerl, Sarah R. Beck, and Teresa McCormack, 230–41. New York: Oxford University Press.

Elga, Adam. 2001. "Statistical Mechanics and the Asymmetry of Counterfactual Dependence." *Philosophy of Science* 68 (3) (September 1): S313–S324.

Fantl, Jeremy, and Matthew McGrath. 2002. "Evidence, Pragmatics, and Justification." *The Philosophical Review* 111 (1): 67–94.

———. 2009a. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

———. 2009b. "Critical Study of John Hawthorne's Knowledge and Lotteries and Jason Stanley's Knowledge and Practical Interests." *Noûs* 43 (1): 178–92.

———. 2011. "Pragmatic Encrouachment." In *The Routledge Companion to Epistemology*, edited by Sven Bernecker and Duncan Pritchard, 558–68. New York: Routledge.

Feldman, Richard. 1995. "In Defence of Closure." *The Philosophical Quarterly* 45: 487–94.

———. 1999. "Contextualism and Skepticism." *Philosophical Perspectives* 45: 487–94.

Feldman, Richard, and Earl Conee. 2008. "Evidentialism." In *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 2nd. ed., 310–21. Malden, MA: Blackwell Pub.

Fine, Kit. 1975. "CRITICAL NOTICE." *Mind* LXXXIV (1): 451–458.

Fogelin, Robert J. 1994. *Pyrrhonian Reflections on Knowledge and Justification*. New York: Oxford University Press.

———. 2000. "Contextualism and Externalism: Trading in One Form of Skepticism for Another." *Philosophical Issues* 10: 43–57.

Galles, David, and Judea Pearl. 1998. "An Axiomatic Characterization of Causal Counterfactuals." *Foundations of Science* 3 (1) (January 1): 151–182.

Gendler, Tamar Szabó, and John Hawthorne. 2005. "The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies." *Philosophical Studies* 124 (3): 331–52.

Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–3.

Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73 (20): 771–91.

———. 1988. "Strong and Weak Justification." *Philosophical Perspectives* 2: 51–69.

———. 2011. "Reliabilism." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2011. http://plato.stanford.edu/archives/spr2011/entries/reliabilism/.

Greco, John. 2000. *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge University Press.

Halpern, Joseph Y. 2000. "Axiomatizing Causal Reasoning." *Journal of Artificial Intelligence Research* 12 (1) (May): 317–337.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.

———. 2005. "Chance and Counterfactuals." *Philosophy and Phenomenological Research* LXX (2): 396–405.

Heller, Mark. 1999a. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy* 77 (2): 196–208.

———. 1999b. "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives* 13: 115–29.

Hiddleston, Eric. 2005. "A Causal Theory of Counterfactuals." *Noûs* 39 (4): 632–57.

James, William. 1912. *The Will to Believe and Other Essays in Popular Philosophy*. Longmans, Green, and Co.

Kaplan, David. 1989. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals." In *Themes from Kaplan*, edited by Joseph Almog, John Perry, and Wettstein, 481–566. Oxford: Oxford University Press.

Klein, Peter. 2000. "Contextualism and the Real Nature of Academic Skepticism." *Philosophical Issues* 10: 108–16.

Kornblith, Hilary. 2000. "The Contextualist Evasion of Epistemology." *Philosophical Issues* 10: 24–32.

Kripke, Saul. 2011. "Nozick on Knowledge." In *Philosophical Troubles*, 162–224. New York: Oxford University Press.

Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.

Lewis, David. 1973a. *Counterfactuals*. Malden, MA: Blackwell.

———. 1973b. "Causation." *The Journal of Philosophy* 70 (17) (October 11): 556–567.

———. 1979a. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339–59.

———. 1979b. "Counterfactual Dependence and Time's Arrow." *Noûs* 13 (4) (November 1): 455–476.

———. 1986. "Postcripts to 'Counterfactual Dependence and Time's Arrow'." In *Philosophical Papers: Volume II*, 52–66. Oxford: Oxford University Press.

———. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–67.

MacFarlane, John. 2005. "The Assessment Sensitivity of Knowledge Attributions." In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne, 1:197–233. Oxford: Clarendon Press.

———. 2009. "Nonindexical Contextualism." *Synthese* 166: 231–50.

Martin, Raymond. 1975. "Empirically Conclusive Reasons and Scepticism." *Philosophical Studies* 28 (3): 215–7.

McGinn, Colin. 1984. "The Concept of Knowledge." *Midwest Studies in Philosophy* 9: 529–54.

Moore, G. E. 1962a. "Certainty." In *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 2nd. ed., 31–4. Malden, MA: Blackwell Pub.

———. 1962b. "Four Forms of Scepticism." In *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 2nd. ed., 29–30. Malden, MA: Blackwell Pub.

Neta, Ram. 2003. "Contextualism and the Problem of the External World." *Philosophy and Phenomenological Research* 66 (1): 1–31.

Northcott, Robert. 2009. "On Lewis, Schaffer and the Non-Reductive Evaluation of Counterfactuals." *Theoria* 75 (4): 336–343.

Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.

Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

———. 2002. "Reasoning with Cause and Effect." *AI Magazine* 23 (1): 95–111.

Pollock, John L., and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. 2nd ed. Rowman & Littlefield.

Pryor, James. 2000. "The Skeptic and the Dogmatist." *Noûs* 34 (4): 517–549.

———. 2004. "What's Wrong with Moore's Argument?" *Philosophical Issues* 14 (1): 349–378.

Rabinowitz, Dani. 2011. "The Safety Condition for Knowledge." *Internet Encyclopedia of Philosophy*. http://www.iep.utm.edu/safety-c/.

Richard, Mark. 2004. "Contextualism and Relativism." *Philosophical Studies* 119: 215–42.

Riggs, Wayne D. 2002. "Beyond Truth and Falsehood: The Real Value of Knowing That P." *Philosophical Studies* 107 (1) (January 1): 87–108.

———. 2003. "Balancing Our Epistemic Goals." *Noûs* 37 (2): 342–352.

Rips, Lance J. 2010. "Two Causal Theories of Counterfactual Conditionals." *Cognitive Science* 34: 175–221.

Roush, Sherrilyn. 2005. *Tracking Truth: Knowledge, Evidence, and Science*. Oxford: Clarendon Press.

Royall, Richard. 1997. *Statistical Evidence: A Likelihood Paradigm*. Boca Raton, Fla.: Chapman and Hall.

Schaffer, Jonathan. 2004a. "Skepticism, Contextualism, and Discrimination." *Philosophy and Phenomenological Research* 69 (1): 138–55.

———. 2004b. "Counterfactuals, Causal Independence and Conceptual Circularity." *Analysis* 64 (4) (October 1): 299–309.

———. 2005. "What Shifts? Thresholds, Standards, or Alternatives?" In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by Gerhard Preyer and Georg Peter, 115–130.

Schiffer, Stephen. 1996. "Contextualist Solutions to Scepticism." *Proceedings of the Aristotelian Society* 96: 317–33.

Sloman, Steven A., and David A. Lagnado. 2005. "Do We 'Do'?" *Cognitive Science* 29: 5–39.

Slote, Michael A. 1978. "Time in Counterfactuals." *The Philosophical Review* 87 (1) (January 1): 3–27.

Sober, Elliott. 2008. *Evidence and Evolution: The Logic behind the Science*. New York: Cambridge University Press.

Sosa, Ernest. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13: 141–53.

———. 2000. "Skepticism and Contextualism." *Philosophical Issues* 10: 1–18.

Spellman, Barbara A., Alexandra P. Kincannon, and Stephen J. Stose. 2005. "The Relation between Counterfactual and Causal Reasoning." In *The Psychology of Counterfactual Thinking*, edited by David R. Mandel, Denis J. Hilton, and Patrizia Catellani, 28–43. New York: Routledge.

Stalnaker, Robert. 1968. "A Theory of Conditional." In *Ifs: Conditionals, Belief, Decision, Chance, and Time*, edited by William L. Harper, Robert Stalnaker, and Gleen Pearce, 41–55. Boston: D. Reidel Publishing Company.

Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Clarendon Press.

Stine, Gail. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29 (4): 249–61.

Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.

Vogel, Jonathan. 1987. "Tracking, Closure, and Inductive Knowledge." In *The Possibility of Knowledge: Nozick and His Critics*, edited by Steven Luper-Foy, 197–215. New Jersey: Rowman & Littlefield.

———. 1990. "Are There Counterexamples to the Closure Principle?" In *Epistemology: An Anthology*, edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath, 2nd. ed., 290–301. Malden, MA: Blackwell Pub.

———. 1999. "The New Relevant Alternatives Theory." *Philosophical Perspectives* 13: 155–80.

———. 2000. "Reliabilism Leveled." *The Journal of Philosophy* 97 (11): 602–23.

———. 2012. "The Enduring Trouble with Tracking." In *The Sensitivity Principle in Epistemology*, edited by Kelly Becker and Tim Black, 122–51. New York: Cambridge University Press.

Wasserman, Ryan. 2006. "The Future Similarity Objection Revisited." *Synthese* 150 (1) (May 1): 57–67.

Williams, Michael. 1988. "Epistemological Realism and the Basis of Scepticism." *Mind* 97 (387): 415–39.

———. 1996. *Unnatural Doubts*. New Jersey: Princeton University Press.

———. 2001. *Problems of Knowledge*. New York: Oxford University Press.

———. 2013. "Skepticism, Evidence and Entitlement." *Philosophy and Phenomenological Research* 87 (1): 36–71.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Clarendon Press.

Wittgenstein, Ludwig. 1922. *Tractatus: Logico-Philosophicus*. Translated by D. F. Pears and B. F. McGuinness. London: Routledge & Kegan Paul Ltd.

———. 1969. *On Certainty*. Translated by Denis Paul and G. E. M. Anscombe. Oxford: Basil Blackwell.

Wright, Crispin. 2004. "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society Supplementary Volume* 78 (1): 167–212.

Zalabardo, José L. 2012. *Scepticism and Reliable Belief*. Oxford: Oxford University Press.

VITA

Kok Yong Lee grew up in Kuala Lumpur, Malaysia. After high school, he went to Southern College (now Southern University College) to study Chinese Literature in 1998. He became an undergraduate of National Taiwan University majoring in philosophy in the fall of 2001. He earned his BA in Philosophy in the spring of 2004. In the fall of 2004, he became a graduate student of the Department of Philosophy at National Taiwan University. He earned his first MA in Philosophy in the spring of 2006. The title of his MA thesis is "First-Person Authority and Privileged Access" supervised by Dr. Chin-Mu Yang. He became a graduate student of the Department of Philosophy at University of Missouri-Columbia in the fall of 2008. He earned his second MA in the spring of 2010. He wrote two of his MA papers ("On the Distinctive Value of Knowledge" and "The Norm of Assertion and Communication") with Dr. Matthew McGrath. He successfully defended his doctoral dissertation on December 3, 2013. The title of his dissertation is "A Solution to Skeptical Puzzles" supervised by Dr. Matthew McGrath.