Skepticism Avoided

by

Patrick Hawley

Submitted to the Department of Linguistics and Philosophy
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2003

Author . . . . / . . . .
Department of Linguistics and Philosophy
September 1, 2003

Certified by . . . .
Robert Stalnaker
Professor
Thesis Supervisor

Accepted by . . . .
Vann McGee
Chairman, Department Committee on Graduate Students

Skepticism Avoided

by

Patrick Hawley

Submitted to the Department of Linguistics and Philosophy
on September 1, 2003, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

I evaluate three replies to skepticism, drawing conclusions about the meaning of "justified", the viability of foundationalism, the value of knowledge, and the role of belief in rational action.

In the first chapter, I examine the following skeptical argument: Something is justified only if justified by a justified thing; circular and infinite chains of justification are illegitimate; therefore, no belief is justified. A linguistic investigation reveals that this argument contains two ambiguities not yet noticed by epistemologists. The linguistic observations favor foundationalism about justification, showing how the foundationalist can maintain his view, while explaining away the force of the skeptical argument. However, in the second chapter, I argue that foundationalism is unsatisfactory, for non-skeptical reasons. If a foundationalist tries to explain how some things can be basic, then she must endorse a certain kind of circularity. But a foundationalist should not endorse this circularity.

Dissecting a single skeptical argument is an interesting although limited endeavour. In the third chapter, I argue that an entire class of skeptical arguments can be avoided. Distinguishing rational action from rational belief change, I claim that certain changes in belief cannot occur during a rational act. In particular, I argue, some skeptical conclusions cannot be accepted while performing an ordinary rational act. The main conclusion of this chapter is: to avoid acting irrationally, it is rational to avoid certain skeptical arguments.

Sometimes it is better to concede to skepticism than to flee. In the fourth chapter, I argue that knowledge is no more valuable than stable true belief. This surprising claim supports the conclusion that skepticism about knowledge is harmless. Even if we cannot know anything about the external world—or even if we cannot know anything at all—we may have something just as valuable as knowledge: stable true beliefs.

Thesis Supervisor: Robert Stalnaker
Title: Professor

Acknowledgments

Most of what I have learned about Philosophy and Linguistics, I have learned from my fellow students at MIT. No community of colleagues could be more helpful and supportive. I could not name anyone without naming everyone. Thank you all.

This dissertation would not exist without the patient and supportive assistance of Robert Stalnaker. I am also grateful for the help of committee members Richard Feldman (who kindly aided from afar) and Alex Byrne.

Federica La Nave's generous labor gave me time. Elia Hawley's generous laughter gave me joy. Thank you, both, very much.

Contents

Part I

Skepticism and foundationalism

Chapter 1

What justifies that?

Abstract

I clarify and defuse an argument for skepticism about justification with the aid of some results from recent linguistic theory. These considerations shed light on discussion about the structure of justification.

1.1 Introduction

If the house is surrounded, then the house is surrounded by something. If your glass is filled, then your glass is filled by something. If an eye is opened, then an eye is opened by someone. If my belief is justified, is my belief justified by something? It is difficult to resist saying yes; in general, it is difficult to resist the following claim:

For any B, if B is justified, then B is justified by something.

Yet, resistance matters. The claim plays a role in a central argument for skepticism about justification. The claim also figures importantly in arguments against foundationalist accounts of justification.

In this chapter, I examine whether the claim can be successfully resisted. As a test case, I consider a skeptical argument that relies on the claim. With the help of a little linguistics, I clarify the skeptical argument, arguing that we can both reveal flaws in the argument, and explain its apparent compellingness. I then draw some conclusions about discussion of the structure of justification.

Let us now see the claim at work in a skeptical argument.

1.2 The Agrippan argument

One venerable argument for skepticism about justification goes like this. Suppose that one of my beliefs is justified. Since my belief is justified, it is justified by something. Now, if what justifies my belief is not itself justified, then my belief is not justified. So, what justifies my belief is justified. But if what justifies my belief is justified, then it is justified by something. And so on. This process either moves in a circle, or goes on infinitely. Since neither a circle nor an infinite chain suffices to justify my belief, my belief is not justified. Thus we have a reductio of the claim that my belief is justified. This argument can, at any time, be given by anyone, about any belief. So, no belief can ever be justified.[1]

I suspect that few philosophers are driven to skepticism about justification by the Agrippan argument on its own, for there are many ways to try to avoid the conclusion. According to coherentism, a circular chain of justification can be legitimate. A foundationalist thinks that each chain eventually ends. Infinitism says that an infinite chain of justification can be acceptable. These views about the structure of justification appear in various forms and combinations.

Although few philosophers are swayed by the Agrippan argument, many philosophers rely on parts of the Agrippan argument. The best argument for coherentism relies on the Agrippan argument (minus the premise that circular chains of justification are always illegitimate). The best argument for foundationalism relies on the Agrippan premise that neither circular nor infinite chains suffice for justification. Thus, plenty of non-skeptics should be interested in the Agrippan argument. An examination of the Agrippan argument might shed light on discussions about the structure of justification.

---

[1]Of ancient lineage, such skeptical reasoning appears in Sextus Empiricus, *Outlines of Pyrrhonism* and *Against the Mathematicians*. Diogenes Laertius (IX 88) attributes it to one Agrippa, a philosopher otherwise unmentioned in ancient manuscripts. J. Barnes discusses the Agrippan reasoning helpfully in *The Toils of Scepticism*.

This rest of this chapter is organized as follows. In section 3, I make the premises of the Agrippan argument explicit. Then, in sections 4 and 5, I try to show that attention to linguistic issues shows that the argument is unsound, and explains why it appears sound. Finally, in section 6, I draw some conclusions.

1.3 The Agrippan premises

The Agrippan argument, as just presented, is a reductio of the hypothesis that some belief is justified, relying on the following three premises:[2]

Premise I: For any A, if A is justified, then A is justified by something.

Premise II: For any A and B, if A is justified, and A is justified by B, then B is justified.

Premise III: Neither circular nor infinite chains suffice to confer justification.

So let me now turn to Premise I and Premise II.

1.3.1 Premise I and Premise II

Premise I is very compelling:

Premise I: For any A, if A is justified, then A is justified by something.

To any claim that A is justified, it seems one can coherently ask "What justifies A?" And not only can "what justifies A?" be coherently asked, the following response is inadequate: "Nothing justifies A, it just is justified." That the question "what justifies A?" can be asked, demanding a substantial response, is evidence that, for any A, "A is justified" entails "Someone or something justifies A" which, in turn, leads to "A is justified by someone or something."

Moreover, two analogies motivate Premise I. 1. Suppose Alfred was infected with a nasty cold. Clearly, Alfred was infected by someone (Ethelred perhaps) or something (like a dirty glass). If becoming justified is relevantly like getting infected, then: if

---

[2]In presenting the premises, I make no mention of belief. In Premises I and II, the quantifiers could range over beliefs, collections of beliefs, or something else. For present purposes, assume that the quantifiers range over anything epistemically relevant. (Although I am discussing the argument as if the conclusion is that no belief can be justified, it should be noted that similar reasoning may be used to argue that no action, nor person, nor anything else can be justified.)

15

belief A is justified, then A is justified by something. 2. To be justified is to be supported. But if something is supported, it has a support. If a book is supported, it is supported by something. If the analogy holds, then: if a belief is justified it is justified by something.

According to Premise II,

Premise II: For any A and B, if A is justified, and A is justified by B, then B is justified.

This makes sense. Certainly not just *anything* will serve as a justifier. In particular, it's difficult to see how a B lacking justification could be a justifier. Consider an example. Suppose your friend tells you that he believes that the Queen of England is in San Francisco. "That's interesting," you say. "Why do you believe that?" "Because all the heads of state in the world have moved to San Francisco." "Really?" you reply. "Why do you believe that?" "Oh, no reason at all. I just believe it," he answers. Assuming that your friend is speaking sincerely, your friend's belief that the Queen of England is in San Francisco seems to lack justification. And that belief, apparently, does not justify his belief that the Queen of England is in San Francisco. To reject Premise II is to embrace the possibility that a chain of justification can legitimately end with something which lacks justification.

Notice that Premise II can also be motivated by the support analogy. If a book is supported, and it is supported by a shelf, then the shelf is supported too. If the shelf were not supported– for example, by the wall– then the book would not be supported.

So, it appears that Premise I and Premise II are both plausible. Now a few brief remarks about Premise III.

1.3.2 Premise III

Premise III: Neither circular nor infinite chains suffice to confer justification.

Two analogies will reveal the plausibility of Premise III. Fully defending Premise III would require arguing against various forms of coherentism. Although important, this is not my present task.

First Analogy

Think of a belief's being justified by something as like a book's being physically supported by something. For example, my copy of *The Great Chain of Being* is supported by a shelf. Moreover, my book is part of a *structure of support*: my book is supported by the shelf; the shelf is supported by the wall; the wall is supported by the foundation of the building; and so on.

Now, plausibly, if a book is supported, then it is part of an acceptable support structure. And, clearly, some structures of support are not acceptable. Self-supporting structures, for example, are not acceptable: the book cannot be physically supported by itself. In addition– if the support relation is transitive– no circular structures are acceptable. You see the point: if, as seems plausible, "A is justified by B" describes a relation between A and B that is anti-symmetric and transitive, then A's being part of a circular chain does not suffice to make A justified.[3]

We can rule out infinite structures of physical support with empirical premises. Infinite structures of physical support require an infinite amount of matter. If there is a finite amount of matter in the universe, then infinite structures of physical support are impossible. Similarly we may look for empirical premises to rule out infinite chains of justification. Since humans are finite, they cannot understand infinite chains of justification. So infinite chains do not suffice to make human beliefs justified.[4]

An analogy is not an argument. Still, if being justified by something is relevantly like being physically supported by something, then we have reason to accept Premise III.

Second Analogy

Think of a belief's being justified as like a person's catching a cold. A healthy person can be infected by coming into contact with a person who is *already* infected. Consider, for simplicity, a cold caused by a virus that cannot survive outside a human

---

[3] J. Barnes makes this and other helpful points in *The Toils of Skepticism*.

[4] A thought like this goes back to at least Aristotle, *Posterior Analytics* I.3. Much more ought to be said to turn this thought into an argument. Recently, Peter Klein has defended the viability of infinite chains. See [38].

body. And suppose that a person can be infected with this virus only once in his lifetime.

Now if Arnold was infected by Ziggy, and Ziggy was infected by Leonard, then it does not follow that Arnold was infected by Leonard. At least, it is not obvious that the relation of *being infected by* is transitive. Still epidemiological circles are ruled out because we are assuming that you can only be infected by this virus once in your lifetime. So it cannot be that you infected yourself; and it cannot be that Arnold was infected by Ziggy, and Ziggy was infected by Arnold. Moreover, since there have only been finitely many people existing in the world, a chain of infection cannot trace back through an infinite number of different people.

According to this analogy, A gets justified by being "infected" by some B that is already "infected". If A is justified by B, then A is justified only if B is *already* justified. A can't "infect" itself with justification, nor can A get "infected" by being part of a circular chain. If there are only finitely many justifiable things then we can rule out infinite chains too.

If being justified is relevantly like catching a cold, then it seems that we have some reason to accept Premise III.

So, it appears that Premise III, as well as Premise I and Premise II, are plausible. And it also appears that accepting each of Premise I, II and III leads to skepticism about justification. However, despite appearances, skepticism about justification does not follow from Premise I, II and III. To see this we will need some conceptual sharpening. A little linguistics will help to get clearer about what is the case when something is justified.

1.4 Some ambiguity

Sentences of the form "X is justified" are syntactically ambiguous. They have at least three distinct readings. Call this the *ambiguity thesis*.

If the ambiguity thesis is correct, then Premise I and Premise II are syntactically ambiguous too. If so, then the soundness of the Agrippan argument is suspect. For

18

the argument is sound only if those disambiguations of Premise I and II that are true, are sufficient (with Premise III) for the conclusion.

I will defend the claim that the Agrippan argument has an unfilled gap: the plausibly true readings of Premise I and II are not sufficient (with Premise III) for the conclusion.

I will also defend the claim that the ambiguity thesis helps explain the plausibility of the Agrippan argument: Premise I and II are plausibly true on *different* disambiguations of "A is justified", and we fail to notice the shift between readings.

I will now present evidence to support the ambiguity thesis. After that, I will argue for the two further claims.

1.4.1 An old worry?

Some philosophers have worried that the meaning of "justify", and "justified" varies with context, and that this causes confusion in epistemological discussion. These are concerns that, I think, are on the right track. However, in the literature, these concerns are presented unclearly and confusedly. Here are two examples.

1. Some philosophers observe that a sentence like "My belief is justified" can apparently be understood merely to say that my belief has a certain status, that of being justified. But on a different understanding, the sentence seems to describe the activity of a person showing that a belief has that status.[5]

2. Some philosophers have noticed that justification sometimes seems to come in degrees, and sometimes not. On the one hand, one belief may be more justified than another. On the other hand, it seems that having a status does not come in degrees; a belief either has the status of being justified or not.[6]

---

[5]W. Alston remarks on the "pervasive confusion between the activity of *justifying* a belief- *showing* the belief to be reasonable, credible, or justified- and a belief's *being* justified, where this is some kind of epistemic state or condition of the believer." ("What's Wrong with Immediate Knowledge?" p 70). See also ("Level Confusions in Epistemology" p 166). R. Audi makes a similar point in *The Structure of Justification* (pp 25-26, also chapters 4 and 10), as does J. Pryor ("The Skeptic and the Dogmatist", p 535), who claims that the status of being justified is "epistemologically primary". R. Chisholm conspicuously avoids forms of the word "justify" in his epistemic principles and definitions in *Theory of Knowledge* (pp 135-140).

[6]A. Goldman, for example, notices this in "What is Justified Belief?" and in chapter 4 of *Epistemology and Cognition*.

19

In other work, I plan to address the second worry, and its relevance to skepticism about justification. My claims here are connected to the first worry.

1.4.2   Eventive and stative readings of passive sentences

[7]An analysis of sentences like "My belief is justified" is rather complicated, so let me begin with some simpler examples:

1 *That door was closed*

2 *My house was surrounded*

Passive sentences like 1 and 2 have two easily distinguishable readings. On one reading, 1 describes an event or a process, namely someone (or something) closing a door.[8] On another reading, 1 describes a state, namely the state of the door after having been closed. Similarly, 2 can describe the event of someone (or something) surrounding my house; or it can describe the state of my house after having been surrounded. Many– perhaps most– English verbs form passive sentences with these two readings.

This difference, between a *stative* reading of 1 and an *eventive* reading of 1 can be even more clearly seen if we add a time modifier:

3 *That door was closed yesterday at 3pm*

Was the door in a closed state just before 3pm yesterday? On an eventive reading, the door was in an unclosed state just before 3pm, and then someone or something closed it at 3pm. On a stative reading, the door may or may not have been in a closed state just before 3pm. So, the two readings of "That door was closed yesterday

---

[7]The discussion in this section and the next is indebted to Embick (2000), Marantz (2000) and Marvin (2000). Many of the example sentences are from Embick (2000). Thanks to Tatjana Marvin and Karlos Arregi.

[8]From now on I will just say "event" because I am not here examining the differences (if any) among events, processes and other non-states. Also, for convenience, I will sometimes speak loosely, saying that a sentence S entails the existence of some event. What I mean is that S entails S', where S' is a sentence that is often taken by philosophers to describe an event. Whether or not *you* take S' to commit one to the existence of an event depends on your own ontological scruples.

at 3pm" have different entailments. Using a superscript E to indicate an eventive reading, and a superscript S for a stative reading, we can say that "That door was closed yesterday at 3pm"^{E} entails

4 *That door was not in a closed state yesterday just before 3pm*

while "That door was closed yesterday at 3pm."^{S} lacks this entailment. Similarly,

5 *My house was surrounded by Federal agents yesterday at 3pm*

can describe the state of my house yesterday at 3pm, or, instead, an untoward event at that time.

1.4.3 Another stative reading

Consider the following sentences:

6 *That door was built closed*

7 *My house was built surrounded*

8 * *That door was built opened*

9 * *That door was built smashed*

10 * *That kettle was forged cleaned*

8, 9 and 10 are very odd.[9] However, 6 and 7 make perfect sense. The problem with 8 seems to be that "opened" indicates not only that the door was in an open state, but also that there was an earlier door opening. But, if the door was built that way, there was no earlier door opening. Similarly, the problem with 9 seems to be that "smashed" indicates not only that the door was in a smashed state, but also

---

[9]The '*' indicates that most native English speaking informants judge that sentence to be ungrammatical. Some informants do find some of the starred examples acceptable, but I have found no one who finds all such examples acceptable. (J. Pryor and J. Wilson helped think of examples.)

that there was an earlier door smashing. But, if the door was built that way, there was no earlier door smashing.

The data suggest the following. "That door was opened" (on a stative reading) entails the existence of an earlier event of opening that door. "That door was smashed" (on a stative reading) entails the existence of an earlier event of smashing that door. "That kettle was cleaned" (on a stative reading) entails the existence of an earlier event of cleaning that kettle.

On the other hand, "That door was closed" (on a stative reading) need not entail that there was an earlier event of closing that door. And "My house was surrounded" (on a stative reading) need not entail that there was an earlier surrounding event.

In short, it appears that there are two stative readings available for passive sentences: an *eventive stative* reading, which entails the existence of a certain earlier event and a *pure stative* reading, which lacks this entailment. I'll indicate the eventive stative reading with a superscript ES, reserving the superscript S for the pure stative reading. Thus, if "X was V'ed" has both an eventive stative and a pure stative reading, then "X was V'ed"^{ES} entails "Someone (or something) V'ed X",[10] while "X was V'ed"^{S} lacks this entailment.

Given that two stative readings are available for passive sentences, one might ask whether any single passive sentence exhibits both stative readings. I believe that the answer is *yes*, although I am unable to show this by simple and intuitive means. Still there is the thought from the last section: "On another reading, 1 describes a state, namely the state of a door after having been closed." Did you find that peculiar when I said it? If not, you have a reason to think that "That door was closed"^{ES} is available.

Despite the lack of a compellingly clear example, I will maintain that some passive sentences (like "That door was closed") have at least three readings: the eventive, the eventive stative and the pure stative. I adopt this hypothesis to help simplify the exposition of my argument; the argument survives (in a modified form) without

---

[10]Where this is a (disambiguated) sentence describing an event. Note that in this generalization 'ed' represents the appropriate passive or past tense ending. For example, the passive "The window was broken" entails the past tense "Someone (or something) broke the window".

it. Sufficient for a version of my argument is a weaker claim: some passive sentences have both an eventive and an eventive stative reading; some passive sentences have both an eventive and a pure stative reading. I believe that the weaker claim is well supported by the intuitive examples I have given.

1.4.4 Summary

This table summarizes the three readings of some passive sentences:

Eventive	S^E	describes the occurrence of an event
Eventive Stative	S^{ES}	a state resulting from an earlier event
Pure Stative	S^S	a "pure" state: no earlier event implied

1.4.5 A syntactic ambiguity

I have just suggested that some passive sentences have at least three readings. The argument was rough, based on intuitive judgments about meaning and entailment.

That there are three readings of sentences like "That door was closed" is compatible with recent linguistic theory. Some generative linguists theorize that sentences like "That door was closed" and "My house was surrounded" are syntactically ambiguous. According to them, such English sentences correspond to at least three distinct syntactic structures.[11]

One might try to explain the differences without positing a syntactic ambiguity. What is important here is that, however the differences between S^E, S^{ES}, and S^S are explained, they may have different entailments. I assume that this claim has been made reasonable. My main claims in this chapter neither depend on the claim that the ambiguity is syntactic, nor depend on a particular syntactic theory. My claim so far is that some passive sentences have three readings, and these readings have different entailments.[12]

---

[11]See, for example, Embick (2000).

[12]Note that Kratzer (2000) develops a slightly different distinction (due to Parsons (1990)) between *target state passives* and *resultant state passives*: "If I throw a ball onto the roof, the target state of this event is the ball's being on the roof, a state that may or may not last for a long time. What I am calling the Resultant-state is different; it is the state of my having thrown the ball onto the

However, that there is a reasonably well developed syntactic theory, attempting to explain the three readings as a case of syntactic ambiguity, supports what I will say in at least two ways. First, the notion of syntactic ambiguity seems to me to be one that is both less controversial and easier to understand than, for example, various mysterious claims that philosophers are prone to make about context dependence. Second, the syntactic theory is intended to explain other kinds of data. It has a degree of empirical support, both by English and by other languages. So if a syntactic ambiguity is posited by such a theory, we have reason, beyond the rough intuitive evidence above, to believe that it is present. In particular, we have a reason to believe that a syntactic ambiguity is present that is independent of philosophical exigencies.

1.4.6 "justified", finally

Like "That door was closed" and "My house was surrounded",

11 *My belief was justified*

has at least three readings. As before, it is easy to distinguish an eventive reading from a stative reading by adding an indication of time:

12 *My belief was justified yesterday at 3pm*

For an eventive reading, one can imagine that the evidence arrived yesterday at 3pm, thereby justifying my belief. On a stative reading, the sentence describes the state of my belief yesterday at 3pm.[13]

As before, we can test for a pure stative reading. The sentence

13 *My belief was formed justified*

---

roof, and it is a state that cannot cease holding at some later time." (Parsons (1990), p 234 quoted by Kratzer (2000), p 2.) I don't believe that this way of drawing the distinction is correct because it mistakenly predicts that "That door is opened" can be true if someone opened the door, closed it again, and now the door remains firmly shut. In any event, Kratzer agrees that some passive sentences are ambiguous between a reading which is event entailing and a reading which is not.

[13]Both "Rico's belief was justified by him" and "Rico's belief was justified by the evidence" have eventive readings. So the eventive/stative distinction I am drawing is not the one mentioned earlier between a *person's* performing an activity, and a belief's having a certain status. (See footnote 5.)

is as innocuous as "That door was built closed", and "That house was built surrounded". Moreover, "My belief was always justified" is acceptable too. Thus, it is reasonable to conclude, as before, that there is a stative reading of "My belief was justified" that entails that there was a justifying event, and a stative reading that lacks this entailment.[14]

1.4.7 A further complication

Before I can connect these observations to the Agrippan argument, I need to clarify the entailments of different passive sentences.

So far, I have concluded that, where S is a passive sentence with all three readings, both S^E and S^{ES} entail the existence of a certain event, while S^S does not. "X was V'ed"^{E} and "X was V'ed"^{ES} both entail "Someone (or something) V'ed X". "X was V'ed"^{S} does not. For example, "That door was closed"^{E} entails "Someone (or something) closed the door", as does "That door was closed."^{ES} "That door was closed"^{S} does not entail "Someone (or something) closed the door". "That door was closed"^{S} describes the door as being in a certain state. "That door was closed"^{ES} describes the door as being in a certain state as the result of a closing event.

However these conclusions need to be qualified. Sometimes "X was V'ed"^{S} *does* entail "Someone (or something) V'ed X". To see this, note that "surround" differs from "closed". If the door was built closed, then it is not the case that someone closed the door. Even if my house was built surrounded, then something surrounded

---

[14]The data is confusing. Some native English speakers find "My belief was formed justified" odd. Thus, according to them, "My belief was justified" resembles "The door was opened", in having only the eventive and the eventive stative readings. (Alec Marantz (p.c.) also suggests that there are some theoretical reasons to expect that "justified" is like "opened", and thus that "My belief was justified" lacks the pure stative reading.) Other native speakers find that they cannot hear anything like an eventive stative reading of "My belief was justified". According to them, "My belief was justified" only has the pure stative and eventive readings.

If there are indeed two rather than three readings of "My belief is justified", the claim I will make about arguments that use sentences like "My belief is justified" survives in a modified form. Even if no passive sentence exhibits both stative readings, as long as some passive sentences have an eventive stative reading, and some passive sentences have a pure stative reading, one may claim that people can be unsure or confused or even disagree about whether a non-eventive reading of "My belief is justified" is an eventive stative or a pure stative. This is enough for the kind of argument I am making.

my house. "My house was surrounded"^{S} entails "Something surrounded my house".

Noting a further ambiguity shows how to repair the conclusions. The active sentence "Something surrounded my house" has both an eventive and a stative reading. One can see the contrast by comparing "Someone closed the door at 3pm yesterday" with "Something surrounded my house at 3pm yesterday". The former has one, eventive, reading. The latter has both stative and eventive readings. That suggests the following generalizations. "X was V'ed"^{S} does not entail an *eventive* "Someone (or something) V'ed X". However, for some, but not all, choices of V, "X was V'ed"^{S} does entail a *stative* "Someone (or thing) V'ed X".

What about "My belief was justified"^{S} ? This is not an easy question. It is not obvious whether being justified is like being closed or like being surrounded. The active sentence "Something justified my belief" has both eventive and stative readings. But if my belief was always justified, did something justify my belief? That is not clear to me. If "My belief was justified"^{S} is true, did something justify my belief? I am not sure, but I hope to find out.

The cogency of the Agrippan argument hinges on the answer.[15]

1.4.8 *A* is justified by *B*

Before returning to the Agrippan argument, I must discuss different readings of passive sentences including a *by-phrase*: "That door was opened by John", or "My belief

---

[15]One further complication deserves mention. While I have discussed past tense passive sentences, the Agrippan argument is given in the present tense. Pure stative, eventive stative, and non-stative readings are also available in the present tense.

Non-stative readings are easily recognized. "That door is closed" has the *habitual* reading that someone usually closes the door. There is also the so-called "sportscaster" reading. Imagine an announcer narrating the action in a play. "She walks toward the open door. Suddenly, the door is closed. 'Must've been the wind' she says." "My belief is justified" has these non-stative readings too.

Now that I have identified them, in what follows, I will ignore non-stative readings of "A is justified". I do so because non-stative readings are easy to distinguish from stative readings. For example, neither the habitual reading nor the sportscaster reading is likely to cause confusion in our understanding of the Agrippan argument. Perhaps confusion between stative and non-stative readings plays a role in other epistemological discussions. I just doubt that that is the case here.

More importantly, the pure stative/eventive stative distinction is also seen in the present tense. Note the contrast between "The door is closed" and "The door is opened." While "The door is closed" may be true because the door was built closed, "The door is opened" is, on a stative reading, only true if there was an earlier door opening event.

is justified by the evidence".

In "The rabid dog was killed by the laws of Texas", the word "by" means "according to" or "in virtue of". But in "The ball was kicked by John", "by" does not mean "in virtue of". Moreover, in "The rabid dog was killed by the police by the laws of Texas", only the second "by" has the "in virtue of"/"according to" meaning. Now suppose that A is justified by B by induction. If so, then A is justified by B, and A is justified by induction. But this does not mean that induction bears the same relation to A as B bears to A.

For simplicity, I will initially assume that the "by" in "A is justified by B" does not mean "in virtue of" or "according to". I will later dispense with this assumption.

As far as whether passive sentences with a by-phrase show an eventive/stative ambiguity, some facts are clear. In some cases, the by-phrase forces an eventive reading. "That door was smashed" has both eventive and stative readings; however, "That door was smashed by John" only describes an event in which John smashed the door. Many– perhaps most– verbs follow this pattern. Other verbs form unambiguously stative active sentences ("John owned the house") and unambiguously stative passive sentences ("The house was owned by John"). There is also a third class of verbs. As just observed, verbs like "surround", and "justify" form active sentences that have both stative and eventive readings. In the passive, with a by-phrase, such verbs yield both stative and eventive readings. "My house was surrounded by Federal agents at 3pm yesterday" can either describe the state of my house at 3pm or an untoward event at 3pm.

I am not sure whether sentences like "My house is surrounded by Federal agents" which have at least one stative reading, have more than one stative reading. For one thing, it is not clear what an eventive stative reading would be. Perhaps we may find different stative readings by searching for different entailments. It may be that, on one reading, "A is justified by B" entails "A is justified"^{S}, while, on another reading, this entailment is lacking. Or, perhaps, "A is justified by B" can either mean that A is *fully* justified by B, or that A is *partly* justified by B. That is, "A is (fully) justified by B" and "B is justified"^{S} together entail "A is justified"^{S}, while "A is (partly)

justified by B" and "B is justified"^{S} do not.

The situation is confusing. In order to chart a clear path, I will ignore this possible plethora of readings, and assume for purposes here that "A is justified by B" has only one reading, a stative reading that lacks these entailments. I make this simplifying assumption because at this point I unfortunately do not know how to clear things up.

1.5 Back to the Agrippan argument

We are now in a position to examine the Agrippan argument. One goal of this section is to defend the thesis that the Agrippan argument has an unfilled gap: the plausibly true readings of Premise I and II are not sufficient (with Premise III) for the conclusion. A second goal is to explain how that thesis helps a non-skeptic.

1.5.1 Eventive Stative Version

Let me now clarify the Agrippan argument by disentangling eventive statives from pure statives. Recall the premises:

Premise I: For any A, if A is justified, then A is justified by something.

Premise II: For any A and B, if A is justified, and A is justified by B, then B is justified.

Premise III: Neither circular nor infinite chains suffice to confer justification.

Since Premise III is not my central focus here, let us accept Premise III. Premises I and II contain as parts sentences of the form "X is justified". Since such sentences are ambiguous, plausibly, Premises I and II are ambiguous too. So we can try to disambiguate Premises I and II. First, take the ambiguous parts as eventive statives:[16]

Premise Ia: if A is justified^{ES}, then A is justified by something.

Premise IIaa: if A is justified^{ES}, and A is justified by B, then B is justified^{ES}.

These premises (coupled with Premise III) are sufficient for the skeptical conclusion. Should we accept them?

---

[16]I suppress the initial quantifiers, and also abuse notation slightly by attaching the superscript to "justified", but the meaning should be clear.

There is reason to accept Premise Ia. As claimed earlier, "A was justified"^{ES} entails "Something justified A". This is the entailment that distinguishes the eventive stative from the pure stative. So, apparently, "A is justified"^{ES} entails "Something justifies A". Thus, "A is justified"^{ES} arguably entails "A is justified by something".[17]

However, Premise IIaa is in doubt. Recall the example given in section 1.3.1 in support of Premise II. An unjustified belief that all the heads of state have moved to San Francisco does not justify a belief that the Queen of England is in San Francisco. This thought provides a reason to accept that only a B in a justified state can serve as a justifier. This is not, however, a reason to accept that only a B in a justified state *as a result of an earlier justifying event* can serve as a justifier. This thought does not adequately support Premise IIaa.

Thus, I conclude that it is not clear whether or not to accept Premise IIaa. Further argument is needed. Perhaps someone could provide further argument. I merely point out that Premise IIaa lacks the evident plausibility that we found in Premise II. This version of the Agrippan argument is incomplete, as it stands.

1.5.2 Pure Stative Version

Perhaps disambiguating in favor of the pure stative holds promise:

Premise Ib: if A is justified^S, then A is justified by something.

Premise IIbb: if A is justified^S, and A is justified by B, then B is justified^S.

These two premises (coupled with Premise III) are also sufficient for the skeptical conclusion.

We have reason to accept Premise IIbb. Unlike Premise IIaa, Premise IIbb contains "B is justified"^{S} rather than "B is justified"^{ES}. The earlier thought that there are no unjustified justifiers does support Premise IIbb.

Premise Ib, however, is in doubt. As I urged in section 1.4.7, although it is clear that if that house was built surrounded, then it was surrounded by something, it is not obvious that if a belief was formed justified that it was justified by something.

---

[17] There is a leap here. "A is justified"^{ES} entails "Something justified A" (on an eventive reading). So, plausibly, it entails "Something justifies A" (on a stative reading).

No matter how that house got into a state of being surrounded, it was surrounded by something; "That house was surrounded"^{S} entails "That house was surrounded by something". However, "A is justified"^{S} may be true because A was formed justified. And if A was formed justified, it is not clear that A was (or is) justified by something. It is not clear whether or not to accept Premise Ib.

Thus, this version of the argument is also incomplete, as it stands.

1.5.3 Mixed Version

So far, we have seen that disambiguating in favor of the pure stative makes Premise II look good, at the expense of casting doubt on Premise I. And disambiguating in favor of the eventive stative makes Premise I looks good, at the expense of casting doubt on Premise II. Since each of the two premises has a plausible disambiguation, we can try combining them together:

Premise Ia: if A is justified^{ES}, then A is justified by something.

Premise IIbb: if A is justified^S, and A is justified by B, then B is justified^S.

Despite their plausibility, Premise Ia and Premise IIbb (coupled with Premise III) are not sufficient for the skeptical conclusion. Drawing the skeptical conclusion would be to equivocate between the pure stative and the eventive stative. The argument needs something more to link these premises together. These two supplementary premises will do the job:

Premise ES-S: if A is justified^{ES}, then A is justified^S,

Premise S-ES: if A is justified^S, then A is justified^{ES},

Premise ES-S is clearly true. "A is justified"^{ES} entails "A is justified"^{S}: it is impossible for A to be in a state (and be in that state as a result of a justifying event), and fail to be in that state.

Premise S-ES is a substantial claim in need of argument. It needs to be argued that any justified A is in that state as the result of a justifying event. Or at least, for a restricted version of the Agrippan argument, every justified belief is in that state as the result of a justifying event.

Perhaps a philosopher could defend Premise S-ES, or find some other way to make

this version of the argument valid. But, as it stands, this version of the argument is, like the others, incomplete.

1.5.4 Other Versions

We have examined three versions of the argument so far. There are eight versions of the argument in all, because there are two disambiguations of Premise I, and four disambiguations for Premise II. For the sake of completeness, here are all of them:

Premise Ia: if A is justified^{ES}, then A is justified by something.

Premise Ib: if A is justified^{S}, then A is justified by something.

Premise IIaa: if A is justified^{ES}, and A is justified by B, then B is justified^{ES}.

Premise IIab: if A is justified^{ES}, and A is justified by B, then B is justified^{S}.

Premise IIba: if A is justified^{S}, and A is justified by B, then B is justified^{ES}.

Premise IIbb: if A is justified^{S}, and A is justified by B, then B is justified^{S}.

We need not tarry long on other versions of the argument. Premise IIab is plausible because, given Premise ES-S, IIab follows from the plausible IIbb. However, like IIbb, IIab cannot be combined with Ia without something like the controversial Premise S-ES. Premise IIba can be successfully combined with Ia. However, IIba is at least as controversial as the controversial IIaa. For, like IIaa, IIba has an eventive stative in the consequent. Finally, although three further versions of the argument include Ib they are not worth examining; if Ib can be defended than the Pure Stative Version can be defended.

1.5.5 The gap

The plausible premises– Ia, IIbb and IIab– cannot be combined with Premise III to produce an argument sufficient for the Agrippan conclusion. I conclude that the argument has an unfilled gap.

Each of the following three premises will fill the gap: Premise Ib, Premise IIaa, Premise S-ES. The last premise yields the other two. For suppose that Premise S-ES is true. Then Ib follows from the Ia. And (given Premise ES-S) IIaa follows from

IIbb.

1.5.6 Shifty by-phrases

I have been assuming that "A is justified by B" has a univocal, stative meaning. But, earlier I pointed out an ambiguity in passive sentences with by-phrases. In a sentence like "A is justified by B by induction", the second "by" means something like "in virtue of" or "according to", while the first does not. Perhaps shiftiness in the meaning of the by-phrase also plays a role in the Agrippan argument. Let's have a look.

The controversial Premise Ib,

Premise Ib: If A is justified^S, then A is justified by something,

seems trivially true if "by" here means "in virtue of". Surely "A is justified"^{S}, if true, is true in virtue of something; *something* makes it the case that A is in that state. If, instead, "by" takes on its usual meaning then Premise Ib is a controversial claim.

Moreover, any version of Premise II looks implausible on the "in virtue of" reading. For if A is justified, and A is justified in virtue of B, why should we think that B is justified? For example, I may be justified in believing some things in virtue of my experience. But why think that my experience is something that needs to be— or even can be— justified?

Shiftiness in the meaning of the by-phrase is one more source of confusion in discussions of justification. And, perhaps, it is one more reason why the Agrippan argument appears compelling. On the "in virtue of" reading, Premise Ib looks trivially true. On the other reading, Premise IIbb is plausible.

1.5.7 Two small steps forward for the non-skeptic

The thesis that there is an unfilled gap helps a non-skeptic in two ways.

First, the non-skeptic can explain away the force of the Agrippan argument. Premise I appears true because Premise Ia is true. Premise II appears true because Premise IIbb is true. And the Agrippan argument appears valid because, on

some disambiguations of the premises, the argument is valid. If it is one part of the non-skeptic's task to explain why this skeptical argument appears plausible, then the non-skeptic has strengthened her position.

Second, the Agrippan argument poses no threat to the non-skeptic. The Agrippan argument is incomplete without an additional premise like Ib, IIaa or S-ES, and an incomplete argument does not threaten the non-skeptic's position.

These two points may not seem strong. One might wish, for example, for an argument that Ib, IIaa and S-ES are false. Still, even these two limited points are interesting, at least for the foundationalist about justification. Consider the following dialectic between skeptic and foundationalist. The skeptic presents Premises I, II and III to argue that no belief can be justified. The foundationalist, assuming that some beliefs are justified, argues that Premise I is false. Is the burden of proof on the skeptic or the foundationalist? Is one side begging the question? If the skeptic's argument has an unfilled gap the foundationalist is free to offer her argument without responding to the skeptic's argument. Worries about question beggingness and burden of proof can be set aside, at least until some other skeptical argument rears its head.[18]

1.6 Concluding remarks

I began this chapter with the following difficult to deny claim:

Premise I: *For any B, if B is justified, then B is justified by something.*

I displayed the central role Premise I plays in the Agrippan argument. Examining English passive sentences revealed that this premise, as well as Premise II of the Agrippan argument, are ambiguous. I concluded that the argument has an unfilled gap, and that this helps the non-skeptic.

What further conclusions can be drawn from this exercise?

---

[18]Curiously, the foundationalist who argues in favor of Premise III may run into trouble with the Agrippan argument. For the support analogy also favors Premise Ib. And the infection analogy favors the entailment from "A is justified"^{S} to "A is justified"^{ES}. On either analogy, the gap is filled! Thus, if the thought behind either of the analogies is what drives the foundationalist to accept Premise III, then the foundationalist's reasons in favor of Premise III may undermine her attempts to reject the Agrippan argument.

1. As I noted earlier, few philosophers are driven to skepticism by the Agrippan argument, but the argument is interesting nonetheless. The premises of the Agrippan argument form a common core around which turns the debate about the structure of justification. Foundationalists, coherentists, and infinitists all accept one or more of the premises of the Agrippan argument. They just disagree about which premises to accept. For example, coherentists and infinitists reject Premise III and accept Premises I and II, while foundationalists accept Premise III and reject Premise I or II.

What I have shown here helps clarify this debate. To successfully argue against foundationalism, coherentist/infinitists should argue for Ia, IIbb or ES-S. Foundationalists must resist all three of Ia, IIbb and ES-S. In resisting, foundationalists can add that Premise I and II appear true because they are true on some readings. For instance, is Premise I true? If a belief is justified is it justified by something? A foundationalist might say that this seems so because Premise Ib is true.

2. The present linguistic investigation provides distinctions useful for philosophers. Here are some examples. As noted earlier, Alston (and others) draw a distinction between a person's performing the activity of justifying a belief and a belief's having the status of being justified. Now we can gain more precision. First, there is the distinction between a state and an event. A belief may be in a certain state, or an event may occur in which a belief enters that state. That event need not involve an activity of "justifying a belief" performed by a person. Second, there is a difference between being in a state as a result of a certain event, and merely being in that state. These distinctions are blurred by the English passive.

In addition, English passives with by-phrases are a source of unclarity. Some foundationalists (e.g. Steup[60], Pryor[54], Feldman[24]) like to say that a belief can be justified by experience but experience is not the kind of thing that needs justification. I suspect that this is loose talk. As noted earlier, by-phrases can take on an "in virtue of"/"according to" reading. In "That bar was closed by the police by the laws of Massachusetts," only the second by-phrase means "in virtue of" or "according to". Such foundationalists appear to mean that some beliefs can be justified in virtue

of experience. But, the usual regress argument for foundationalism begins with a premise like "some beliefs are justified by other beliefs." Here the "in virtue of" reading is not what is called for. Apparently, some foundationalists confusingly mix together different readings of the by-phrase. (I say more about this in Chapter 2.)

3. Finally, the investigation begun here about passive sentences may bear fruit beyond epistemology. If an event was caused, was it caused by something? If a property was exemplified, was it exemplified by something? Getting clear about passive sentences like "X was caused" and "Y was exemplified" may help us clarify these questions, and the discussions where they are asked. It may sound unfashionable, but I am encouraged, by the results here, to believe that patient and careful study of language will help move philosophical discussion forward.

1.7 Onward

A foundationalist need not succumb to the Agrippan argument. Does that mean that foundationalism is the best theory of the structure of justification? In Chapter 2, I argue that the answer is no. There are non-skeptical reasons to think foundationalism unsatisfactory.

Dissecting the Agrippan argument may be interesting, but avoiding the conclusion of one skeptical argument is a limited endeavour. There are many skeptical arguments, and many skeptical conclusions. In Chapters 3 and 4, I ambitiously argue that we can respond to a large class of skeptical arguments at once.

Chapter 2

Obscure Foundationalism

Abstract

Foundationalism about justification is unsatisfactory. If a foundationalist tries to explain how some things can be basic, then she must endorse a certain kind of circularity. However, a foundationalist should not endorse this circularity.

2.1 A puzzle for the foundationalist?

The foundationalist, noting that some justified beliefs are justified by other justified beliefs, argues that not all justified things are justified by other justified things. For if all justified things are justified by other justified things, then either circular or infinite chains of justification are legitimate. But neither circular nor infinite chains of justification are legitimate. Furthermore, nothing is justified by an unjustified thing. So, the foundationalist concludes, since some beliefs are justified, some things are *basic*. They are justified, but not justified by other things.

Such is a version of the so-called *regress argument*, the main argument for foundationalism about justification. It is not my present goal to examine the merits of the regress argument in detail. Instead, I intend to study the viability of the foundationalist picture which emerges from the regress argument: some things are basic; all justified beliefs are either basic, or ultimately derive their justificatory status from basic things; circular and infinite chains of justification are not legitimate. In Chap-

ter 1, I suggested that the foundationalist can evade the Agrippan argument, while explaining away its appeal. If this is correct, then the foundationalist can successfully resist the Agrippan argument. Now, setting skepticism aside, I will raise a different kind of problem for foundationalism.

I will argue that the best foundationalist solution is to decline to offer an explanation why or how basic things are justified. Call this *obscure* foundationalism. More popular views are rather less obscure. Most foundationalists attempt to explain why or how basic things are justified. For example, some philosophers think that perceptual beliefs can be basic, and they try to explain this by alluding to a special relationship that perceptual beliefs have to our experiences that can make perceptual beliefs justified. I believe that such explanatory attempts face a serious difficulty. I will argue that only obscure foundationalism avoids the difficulty. If we are to be foundationalists, we should be obscure foundationalists. I will conclude that, to the extent that obscure foundationalism is unsatisfactory, foundationalism is unsatisfactory.

Next I will present the foundationalist framework with some care. After that I will present a difficulty for those foundationalists who attempt to explain why or how basic things are justified.

2.2 Vanilla foundationalism

2.2.1 Some symbolism

I will next state with more precision the theses that foundationalists share. For brevity, I will help myself to some symbolism. "*Ja*" abbreviates "*a* is justified". "*a < b*" abbreviates "*a* is justified by *b*". The symbol "<" can be repeated. For example, "*a < b < c*" is "a is justified by b, and b is justified by c."

Recall that in Chapter 1, I argued that sentences like "*a* is justified" and "*a* is justified by *b*" are ambiguous. So something must be said to isolate the intended readings of "*Ja*" and "*a < b*".

Among the readings of "*a* is justified" are two ways of describing *a*'s state. An *eventive stative* reading describes *a* as being in a certain state as the result of a justifying event. A *pure stative* reading merely describes *a* as being in a certain state. Sentences like "*Ja*" are to be taken, unambiguously, in the pure stative sense. (Readers may consult Chapter 1 for a more detailed explanation of the intended reading.)

"*a* is justified by *b*" also has more than one reading. In order to focus on the intended reading for "$a < b$", it will help to first set aside some extraneous readings.

As noted in Chapter 1, on one reading of "*a* is justified by *b*," "by *b*" means "in virtue of *b*" or "according to *b*". On other readings, "by *b*" means something else. For example, in "This belief was justified by that belief by induction," only the second "by" means "in virtue of" or "according to". I will set aside the in-virtue-of/according-to reading; "$a < b$" is NOT to be understood in the in-virtue-of/according-to sense.

Another reading to set aside is one in which "*a* is justified by *b*" describes an event. This is not the intended reading of "$a < b$"; "$a < b$" should be taken to describe a state. (It is easy to see the difference in the past tense: "*a* was justified by *b* at noon yesterday" can either describe *a*'s state yesterday at noon, or it can describe an event, which occurred at noon, in which *a* entered a justified state.)

Now that we have set these extraneous readings aside, let me delineate the intended reading of "$a < b$". It is possible to read "*a* is justified by *b*" so that "*Ja*" follows. I intend this to be true of "$a < b$". Thus, if $a < b$ then Ja. Also, it appears possible to read "*a* is justified by *b*" to mean either that *a* is *fully* justified by *b* or that *a* is *partly* justified by *b*. I intend the latter sense. It may be that $a < b$, although *b* is not the full justification for *a*.

Here is a paradigm case to display the intended reading of $a < b$. This is a case where a belief is justified because it is inferred from other justified beliefs. Suppose I have a belief b_1 that my son is either sleeping in his room or having a snack in the kitchen. I glance into the kitchen and, seeing he is not there, form a belief b_2 that he is not having a snack in the kitchen. Assume that Jb_1 and Jb_2. From b_1 and b_2 I

infer a belief b_3 that he is sleeping in his room. Let us also assume that, because b_3 is inferred in this way, Jb_3. Then in this case $b_3 < b_1$, and $b_3 < b_2$.

Hopefully this is enough to focus on an unambiguous reading of "a is justified by b". But it should be noted that I have chosen this reading for "$a < b$" somewhat arbitrarily. Other readings of "a is justified by b" could be used to state the claims of the foundationalist. In addition, it should be kept in mind, as I said in Chapter 1, that there is some doubt as to exactly which readings are had by "a is justified by b".

To avoid clutter, from now on I will only explicitly disambiguate sentences containing the word "justified" when failing to do so may cause confusion.

2.2.2 Vanilla foundationalism

According to *vanilla foundationalism*,

1. For some x, Jx and there is no y such that $x < y$.

2. For any x and y, if $x < y$ then Jy.

3. For any x, it is not the case that every true sentence beginning with the fragment "$x <$" is entailed by a true sentence of the form

      (i) "$x < x_1 < x_2 < ...$" containing an infinite number of terms, or

      (ii) "$x < x_1 < ... < x_n < x$", for some $n > 0$, or

      (iii) "$x < x$".[1]

Vanilla foundationalist 3 may look obscure. The idea is that the vanilla foundationalist thinks that x, if justified, is not justified solely in virtue of being part of a circular chain of justification, an infinite chain of justification or by being self-justifying. The vanilla foundationalist permits circular chains of justification, infinite chains of justification and cases in which $x < x$. But the vanilla foundationalist thinks that these are not sufficient to make it the case that Jx.

The scope of vanilla foundationalism may be somewhat surprising. Vanilla foundationalism encompasses views that some might hesitate to call "foundationalism". For example, a vanilla foundationalist may claim that a belief is justified just in case it

---

[1](ii) could be omitted because this case is subsumed by (i). If some sentence like (ii) is true, then there is a true sentence like (i) which entails it. (iii) could be omitted for the same reason.

is part of a coherent set of beliefs. According to this vanilla foundationalist, every justified belief is basic. In addition, vanilla foundationalism excludes views that some call "foundationalist". Some philosophers say that basic beliefs are justified solely because they justify themselves. Vanilla foundationalism excludes this view. Thesis 3 is intended to rule out this possibility. I am motivated to present foundationalism in this somewhat unusual way in order to mark out the class of views susceptible to coming argument.

Some additional remarks about these three theses are in order.

First, some foundationalists who appear to reject 1 and 2 are best seen as vanilla foundationalists. For example, some philosophers say that when a perceptual belief is justified, it is justified by experience, but experience is not the kind of thing that can be justified.[2] While vanilla foundationalists think that basic things are justified but not justified by anything, these foundationalists say that basic things are experiences which neither are nor can be justified. This is an interesting view; however, it is misleadingly described. These foundationalists say that some justified beliefs are justified by other justified beliefs. These foundationalists also say that some justified beliefs are justified by experience. But clearly these philosophers don't think that what makes a belief justified, when justified by a belief, is the same as what makes a belief justified, when justified by an experience. For one thing, when justified belief b_1 is justified by belief b_2 that is thought to be, in part, because belief b_2 is justified. But when belief b_3 is justified by experience e, the thought is *not* that this is because experience e is justified. Yet, confusingly, the same locution "is justified by" appears in describing both the belief-belief case and the belief-experience case. It seems to me that these philosophers are helping themselves to one of the ambiguities discussed in Chapter 1. On one reading of "a is justified by b," "by b" means "in virtue of b" or "according to b". On another reading, "by b" means something else. For example, in "This belief was justified by that belief by induction," only the second "by" means "in virtue of" or "according to". It seems to me that these philosophers mean that perceptual beliefs are justified *in virtue of* experience, But even if these philosophers

---

[2]See for example [54] and [60].

are not guilty of mixing up different readings of "a is justified by b", these philosophers would avoid some murkiness by keeping the belief-belief case separate from the belief-experience case, by reserving "is justified by" for the belief-belief case and using "is justified in virtue of" for the belief-experience case. And, even if these philosophers are not using "a is justified by b" ambiguously, we can bring them into the fold of vanilla foundationalists by simply stipulating that, in the vanilla foundationalist theses, the terms only refer to things that can be justified.

Second, the vanilla foundationalist theses are intentionally stated in an abstract way for greater generality. What the 'x' can refer to in "Jx" depends on the details of each vanilla foundationalist theory, in particular what kinds of things can be epistemically justified, according to that theory. Some epistemologists talk about beliefs which are justified or not: John's belief that it is raining is justified; Mary's belief that the future will be bleak is not. And there is a great deal of argument about what it is to have a belief and whether there are any at all.[3] Since many philosophers say that believing is a relation between a person and a proposition (or a sentence), it may be better to talk instead about a justified attitude that a person can have to a proposition (or a sentence): Lily is justified in believing that Rome wasn't built in a day; Hubert is not justified in believing that he is ill-paid. However, others think that, in the first instance, what can be justified is not an attitude to a particular proposition (or sentence) but an entire system of belief. Another view is that, in the first instance, what can be justified are changes in an entire system of belief.[4] The argument I will give should work for any vanilla foundationalist no matter her view about what can be justified. I must admit, however, that this claim deserves some support which I am as yet unable to give. In this chapter, for convenience, I will speak as if beliefs are things which are justified or not. I do not intend this to prejudice any issues.

Now that the three vanilla foundationalist theses have been clarified, the regress argument is easy to state. The vanilla foundationalist uses 3 and the anti-skeptical

---

[3]See [44] for some helpful discussion.
[4]As in [33].

premise that for some belief x, Jx, to argue that either 1 or $\neg 2$. Thus, assuming 2, 1 follows.[5]

2.2.3 The source of justification

Given 1, 2 and 3, it follows that, for any x, if Jx, then either x is *basic*,

Jx and there is no y such that $x < y$,

or *derived*:

For some $y_1...y_n$, $x < y_1 < ... < y_n$, where y_n is basic.

In more vivid terms, according to the vanilla foundationalist, if some x is derived, then x inherits its justificatory status from basic things; Jx *because* of one or more basic things.

But if x is basic, what explains how or why Jx?

In the next section, I will explain why this question appears to require an answer. After that, I will argue that attempting to provide an answer– attempting to explain how or why basic things are justified– brings trouble.[6]

2.3 Some explaining to do

2.3.1 What justifies that?

The conclusion of the regress argument is that some things are basic. However, the regress argument does not reveal what kinds of things are basic, or how it can be that basic things are justified. Being a good philosopher, the vanilla foundationalist (call her "the theorist") goes beyond the regress argument. She does not merely claim that

---

[5]Described this way, the regress argument is closely connected to the Agrippan argument of Chapter 1: 3 is Premise III; 2 is equivalent to Premise IIbb; 1 is the negation of Premise Ib.

[6]Foundationalists may wish to accept that basic things can be justified, simultaneously, in two different ways, a "basic way" and a "non-basic way". (This was suggested to me by R. Feldman.[p.c.]) Call this *redundant foundationalism*. A redundant foundationalist could, for example, think that basic belief b is justified in the basic way in virtue of its relation to experience; but that b could, at the same time, also be justified in the non-basic way by other beliefs. As stated, the vanilla foundationalist theses rule out redundant foundationalism. Despite this, redundant foundationalism is also subject to the coming argument. To see this, readers should keep in mind the relevant question for this form of foundationalism: "But if x is basic, what explains how or why Jx in the basic way?"

there are some basic things. She attempts to explain what makes it the case that the basic things are justified– what makes it the case that Jx when there is no y such that $x < y$. To fail to provide an explanation would be to fail to fully articulate the foundationalist theory. To fail to provide an explanation would leave the nature of the basic things mysterious to us; we would be left with the bare claim that there are some.

The following story shows the trouble facing the theorist who fails to provide an explanation. Suppose that the theorist believes that one of Bill's beliefs is justified but is such that nothing justifies it. If a doubter asks "What justifies that?", the theorist responds, "As I said, nothing justifies that." But the doubter can persist: "No, I didn't mean to ask what justifies Bill's belief. I meant: what justifies your belief that Bill's belief is justified, but is such that nothing justifies it?" This is a different question. The theorist could try to answer by arguing that there are basic things, without attempting to explain how or why Jx when x is basic. Unfortunately, that response leaves the doubter's question unanswered. Why should the theorist believe that *that* belief of Bill's is justified? Why should the theorist believe that *that* belief of Bill's is basic? And, more generally, what makes it the case that the basic things are justified? Stopping at the bare claim leaves the nature of the basic things mysterious to us.

Even more, stopping at the bare claim leaves the theorist vulnerable to skeptical worries. Derived beliefs are supposed to inherit their justificatory status from basic things. If the theorist says nothing about what makes it the case that Jx when x is basic, then we might well doubt that there are any basic things, and so doubt whether there are any justified beliefs.

Of course, vanilla foundationalism may be correct, although the theorist provides no explanation of what makes it the case that basic things are justified. For suppose that the theorist does not or cannot provide an explanation of what makes it that case that basic things are justified. Even so, that need not show that there are no basic things. That Bill's belief is justified need not depend on the availability to the theorist of an explanation. In fact, it would be rather surprising to discover that the

theorist's theoretical skill determines whether Bill's beliefs are justified. (Set aside odd cases like the beliefs that Bill gains after reading the theorist's latest philosophical tract. The theorist's theoretical skill may partly determine whether those beliefs are justified.)

Foundationalists indeed attempt to dispel the mystery, accepting the task of explaining how it is that basic things are justified, that is, what makes it the case that Jx, if x is basic. Here are three examples:

1. if X has a belief b that was formed by an unconditionally reliable process, then Jb.[28]

2. if X has a belief b, then, *prima facie*, Jb.[33]

3. if X has an experience with content p, then, *prima facie*,
$J(X$'s believing $p)$.[54]

In each of the three examples, the foundationalist is trying to explain what makes it the case that a basic thing is justified. For instance, according to 1. being formed by an unconditionally reliable process makes it the case that a basic belief is justified. It should be emphasized that the foundationalist has wide latitude in formulating her explanation. For example, the explanation may be externalist (as in 1) or not. Or the basic things may be thought to be indefeasibly justified (as in 1) or not.

2.3.2   Why give an explanation?

I just offered several sketchy reasons why the foundationalist should explain what makes it the case that Jx, when x is basic. First, absent such an explanation, the foundationalist theory is not fully spelled out. Second, the bare claim that there are some basic things is rather mysterious. The theorist ought to dispel the mystery. Third, the lack of an explanation might lead one to doubt that there are any basic things, and so doubt whether there are any justified beliefs. Fourth, to a claim that b is justified, the doubter's question "What justifies that?" has force that is not fully answered by "Nothing justifies b."

These reasons are interrelated; perhaps they are not even distinct. Despite their sketchiness, together they suggest that something more general can be said here: a

theorist who lacks an explanation is being unreasonable somehow.

As a first stab, I propose to locate the unreasonableness in the following way:

(R0) A person with a reasonable belief b can offer an adequate consideration in favor of b.

If (R0) is correct, then we can diagnose what goes wrong when the theorist fails to give an explanation. Suppose that the theorist believes that a is a justified belief, but nothing justifies it. If the theorist has an explanation of what makes it the case that a is justified, then the theorist has an adequate consideration in favor of her belief that a is justified. She can say: if circumstances are such-and-such a way, then a is justified; and circumstances are that way.

But if the theorist has no explanation, what consideration can she offer in favor of her belief that a is justified? The theorist, in answer to the question "What justifies that?", could simply say "Nothing justifies that; it just is justified." But this is hardly an adequate answer. "Nothing justifies that; it just is justified" is not, by itself, an adequate consideration in favor of a belief that something is justified. And without an explanation, the theorist apparently has no other answer to offer.

So, if (R0) is correct, we can diagnose what goes wrong when the theorist fails to give an explanation. Moreover, something like (R0) indeed seems correct. Imagine that a belief that it is raining in Santa Fe pops into Suzy's head. But suppose that Suzy cannot offer *any* consideration in favor of that belief. When asked why she believes that it is raining in Sante Fe, she replies "I believe it for no reason." Suzy seems unreasonable. She should at least be able to offer *some* consideration. She could say, "I had some reason but I can't remember what it is." Or "I trust the things that pop into my head." But if Suzy can say nothing beyond "I believe it for no reason," her belief seems unreasonable.[7]

(R0) says nothing about justification. This is by design. Since (R0) says nothing about justification, (R0) makes no explicit controversial claim about justification. (R0) does not, for example, say that a person with a justified belief can justify that belief. Neither does (R0) say that if a person is justified in believing that p, then she

---

[7]cf. L. Bonjour's example of a reliable clairvoyant in [11]. Bonjour takes his example to support a very different claim than the one I am defending.

is justified in believing that she is justified in believing that *p*. These claims about justification are very controversial, and likely unacceptable to adherents of certain accounts of justified belief. For example, if being justified in a belief is a matter of having used a reliable belief forming process, then it would be surprising if a person with a justified belief can always justify that belief; it would be surprising for a person always to have access to what it is in virtue of which his justified belief has that status.

I intend my argument to apply to any form of vanilla foundationalism, thus I am trying to remain neutral about accounts of justification. The term "reasonable" helps make this neutrality clearer, even for those philosophers who think that a reasonable belief simply is a justified belief.

Some may object that (R0) is too strong. Even so, everyone, I believe, can agree that at least some beliefs, in order to be reasonable, require the presence of supporting considerations. Perhaps this does not hold true for all beliefs. Perhaps small children and animals have reasonable beliefs without being able to offer considerations in favor of them. Perhaps it is even reasonable for us ordinary folks to merely shrug our shoulders in response to a question "What justifies that?" Still, we can agree that a theorist should be in a position to say at least something in favor of her theoretical beliefs. (A theoretical belief is, for example, a belief that some other belief is justified.) For a theorist, at least, embarrassed silence in theoretical matters is not satisfactory. Let us weaken (R0) to make it— hopefully— acceptable to all:

(R) A theorist with a reasonable theoretical belief *b* can offer an adequate consideration in favor of *b*.

(R) is very weak. (R) says that a theorist ought to be in a position to offer an adequate consideration in favor of certain of her beliefs. What counts as an adequate consideration can be something very weak. For example, Suzy's reply that she had some reason, but can't remember it, seems to be a consideration both weak and adequate.

Although (R) is weak, it is not vacuous. An irrelevant remark does not count as an adequate consideration in favor of a belief. If Suzy says "I believe it because lemons are yellow", we would not reckon her belief reasonable. What counts as an

47

adequate consideration in favor of a belief? That is a significant question which I am not sure how to answer. I believe, however, that the cases I will consider here are clear enough to allow reliance on intuitive considerations in distinguishing an adequate consideration from an inadequate consideration.

At this point, the reader may suspect that I have not answered the question with which I began. I was trying to explain why the foundationalist is being unreasonable who goes no further than the bare claim that there are basic things. I then suggested that a theorist who believes that some thing b is basic lacks adequate consideration in favor of her belief, unless she has an explanation of what makes it the case that Jx when x is basic. But there is a difference between believing that a particular thing b is basic, and believing the bare claim that there are basic things. The foundationalist who believes that there are basic things does have a consideration to offer in favor of that belief: the regress argument. Why is the foundationalist who stops at the bare claim being unreasonable?

The answer, I believe, comes in two steps. First, it is unreasonable for the theorist to believe the bare claim that there are basic things without a plausible instance, some thing which is plausibly basic. But, in order to provide a plausible instance, the theorist needs to believe of some thing that it is basic. Second, it is unreasonable to believe of some thing that it is basic without an explanation of what makes it the case that Jx when x is basic.

The second step is explained by condition (R). The first step is not explained by condition (R), for the regress argument is, apparently, an adequate consideration in favor of a belief in the bare claim. I admit that I do not have a satisfactory argument for the first step. Even so, there is intuitive force behind thinking that in some cases it is unreasonable to believe an existential claim without a plausible instance. You wouldn't feel satisfied to believe that there are some leprechauns just because the old stories say so. You'd want to see one. Now I certainly don't mean to commit myself to a general principle that it is unreasonable to believe an existential claim without a plausible instance. Such a general principle seems false: you might reasonably be convinced that there are numbers too big for you to write down in your lifetime,

without having a plausible instance of such a number. In contrast, the bare claim is offered as the conclusion of a very controversial argument, not a mathematical proof. That makes it seem appropriate to demand a plausible instance. I will take it as a working assumption that the first step is indeed correct. Things will need to be rethought if we find reason to think that it is incorrect.

To summarize: the foundationalist needs an explanation for what makes it the case that a basic thing is justified. This need is (in part) explained by condition (R).

2.4 Obscure foundationalism

2.4.1 The argument

Suppose the theorist offers an explanation of what makes it that case that Jx when x is basic.

Suppose also that the theorist, being a good theorist, is reflective about her theoretical beliefs. For example, she not only believes in her explanation, she believes that this belief is justified. Call her explanation 'T', and her belief in T, 'b_T'. The theorist further believes that Jb_T.

One day, the theorist starts thinking about this last belief of hers; she starts wondering whether Jb_T. The theorist has no particular reason to doubt that Jb_T. She is simply wondering whether it really is the case that Jb_T. I believe that this is something a theorist should do; a theorist should be reflective, and a reflective theorist should sometimes wonder about her own theoretical beliefs. Is this correct? Should a theorist be reflective in the way I am suggesting? I will take this question up later.[8] For now, let us see what ensues when we assume that this theorist, believing that Jb_T, starts wondering whether it is the case that Jb_T. Note, to be clear, the theorist is thinking about her own beliefs, not anyone else's.

As discussed in the last section,

1. Only if she relies on b_T can the theorist offer an adequate consideration in favor

---

[8]That I should take up this question suggests that a theorist should indeed sometimes wonder about theoretical beliefs.

of her belief that some things are basic.

For we are assuming that it is unreasonable to believe the bare claim that there are basic things without a plausible instance, some thing which is plausibly basic. But in order to provide a plausible instance, the theorist needs to offer a particular b which is plausibly basic. However, the theorist can't reasonably believe of some particular thing b that it is basic without relying on b_T.

Let me go through this last step carefully. Suppose the theorist believes of some particular b that it is basic. That is, she believes that Jb but there is no y such that $b < y$. If so, then she can offer no y such that $b < y$ as a consideration in favor of her belief that Jb. For she believes that there is no such y. The only consideration she has to offer involves her explanation of what makes it the case that a basic thing is justified. So she needs to rely on b_T to offer an adequate consideration in favor of Jb.[9]

An example will help make this clearer. Suppose Herman is a vanilla foundationalist who believes that perceptual beliefs can be basic, and are justified when experience is appropriately connected to the perceptual belief. For example, suppose, looking in a mirror, Herman comes to believe that his nose is red. Call this belief b_1. According to Herman's theory, although justified, b_1 is not justified by anything. It is Herman's visual experience which makes it the case that b_1 is justified. I am arguing that if Herman has a further belief b_2 that b_1 is basic, then only if he relies on his explanation can Herman offer an adequate consideration in favor of b_2. For if Herman believes that b_1 is justified but not justified by anything, then what consideration can he offer in favor of his belief that b_1 is justified without involving his explanation of how it could be that something is justified without being justified by anything? Herman could say, "Well, my nose just looks red. That's my reason for believing that it is red." While that reason may support or explain why Herman accepts b_1, the question is not what consideration can he offer in support of b_1. The question

---

[9]Strictly speaking, belief in T*, a rival explanation to T, would also give her something to say. So it is not b_T, in particular, that she needs, but a belief in some explanation. Noting this does not affect the argument. We can presume that the theorist did not start out believing in both of two rival explanations, T and T*.

is what consideration can he offer in support of the higher order belief b_2 that b_1 is justified. And given that Herman thinks that b_1 is basic, it is difficult to see how his nose looking red could be for Herman an adequate consideration in favor of b_2. For one thing, Herman cannot think that this reason justifies b_1 because he thinks that b_1 is not justified by anything. Of course if Herman adds that experiences of the way things look can make it the case that a perceptual belief is justified, then Herman does appear to have an adequate consideration in favor of b_2. But this is to bring in his explanation of how it can be that basic things are justified.[10]

Now, given (R),

(R) A theorist with a reasonable belief can offer an adequate consideration in favor of that belief,

it follows from 1 that

2. If the theorist does not rely on b_T, then it is unreasonable for her to believe that some things are basic.

And, more surprisingly, given that the theorist does believe that some things are basic,

3. If the theorist does not rely on b_T, then many– perhaps most– of her beliefs are unreasonable.

To see this, suppose that the theorist does not rely on b_T. Then, for any thing x, if she believes that x is basic, then her belief is unreasonable. As just noted, without b_T, she has nothing adequate to say in favor of her belief that x is basic. And the theorist's belief that there are basic things is also unreasonable (given our assumption that it is unreasonable for the theorist to believe the bare claim that there are basic things without a plausible instance.)

---

[10] A curious artifact of Herman's theory permits Herman to persist along the following lines: "That it looks that way is a consideration in favor of b_1 although it does not justify b_1. Still, that means that I have a non-justifying consideration in favor of belief b_1. That there is such a non-justifying consideration is a reason to believe that b_1 is justified. So I have a reason to believe that b_1 is justified." I find this reasoning suspicious: why is the presence of a non-justifying consideration a reason to believe b_1 is justified? And what is a non-justifying consideration anyway? But even if this reasoning is accepted, the real problem comes because the theorist is wondering whether her explanation is justified. If Herman is wondering whether his belief (that experiences make perceptual beliefs justified) is justified, then it is illegitimate for him to offer his experience as a consideration in favor of believing that a perceptual belief is justified. For he is wondering whether he is justified in believing that experiences can be such considerations! I will explain this further when I reach step 4.

Worse, the unreasonableness affects certain beliefs about her non-basic beliefs. For each belief x of hers, if she believes that x is non-basic, and believes that x is justified, then her belief that x is justified is unreasonable. For the theorist is a vanilla foundationalist. She believes that if some x is a non-basic belief, then x inherits its justificatory status from basic things; Jx *because* of one or more basic things. But she has nothing adequate to say in favor of a belief that some thing is basic. So she has nothing adequate to say in favor of a belief that some non-basic belief is justified.

Even worse, the unreasonableness in these beliefs about her non-basic beliefs may infect further beliefs. Arguably, unreasonableness in a belief about a belief can infect the lower order belief: if b_2 is a belief that b_1 is justified, and b_2 is unreasonable, then b_1 is unreasonable too. In other words, if the theorist's beliefs about the justificatory status of other beliefs are unreasonable, then those other beliefs are arguably unreasonable too. Whether and how further beliefs are affected depends on how (R) is spelled out, and how reflective the theorist is. If (R) applies not only to theoretical beliefs but to other beliefs as well, then those other beliefs will be unreasonable too. And since a belief will not be "poisoned" unless the theorist has an unreasonable higher order belief about its epistemic status, the number of lower order beliefs affected depends on what higher order beliefs are held. An unreflective person may have few beliefs about her beliefs. The theorist, however, is presumably different. The reflective theorist should have many beliefs about the epistemic status of her beliefs. The more reflective the theorist is, the greater the number of her unreasonable beliefs.

Thus, while wondering whether Jb_T, either many of the theorist's beliefs are unreasonable, or she relies on b_T. But she should not rely on b_T while wondering:

4. The theorist should not, at the same time, wonder whether Jb_T, and rely on b_T.

To do so would be to endorse an illegitimate kind of circularity. This circularity can be displayed by example. Suppose Amy believes that God makes it the case that beliefs about God are justified. Call this belief b_1. Suppose Amy also believes that b_1 is justified. Call this belief b_2. Suppose Amy wonders about b_2; she wonders whether

b_1 is indeed justified. She surely would not be satisfied if she thinks to herself "Well, since b_1 is a belief about God, and God makes it the case that beliefs about God are justified, then b_1 is justified." For that would be to use what she believes as a premise in an argument concluding that her belief is justified. To be sure, the unacceptability of Amy's reasoning does not show that her beliefs are false or unjustified. If God does make it the case that beliefs about God are justified, then her belief b_1 is both true and justified. Still, her reasoning should not satisfy her wondering; it should certainly not renew her confidence that b_1 is justified.

Similarly, if the theorist is wondering whether Jb_T she should not be satisfied by reasoning which relies on b_T. Recall the plight of Herman. Herman believes that his nose is red, and he believes that this perceptual belief is basic. And suppose Herman is also wondering whether he is justified in believing that experiences can make it the case that perceptual beliefs are justified. While wondering, Herman should not offer his experience ("Well, it looks red to me") as a consideration in favor of his belief that this perceptual belief is justified. For he is wondering whether he justifiably believes that experiences can do that.[11]

Thus, while wondering whether Jb_T, either many– perhaps most– of the theorist's beliefs are unreasonable, or she illegitimately relies on b_T.

I do not believe that this conclusion shows that vanilla foundationalism is false or incorrect. However, I think the conclusion does show that there is something very unsatisfying about vanilla foundationalism, even if it is true. The vanilla foundationalist should be able to satisfactorily resolve her wondering. She should be able to satisfy herself that her theoretical beliefs are correct. But she cannot offer an ade-

---

[11]The circularity I am describing resembles what W. Alston calls *epistemic circularity*. A person is justified in believing the premises of an epistemically circular argument only if the conclusion is true. (See [2] and [6]. Also, see [60] for a different understanding of epistemic circularity.) It has come to my attention that in recent unpublished work, M. Bergmann[9] has argued that foundationalism brings commitment to epistemic circularity, and that since foundationalism is a better theory than its competitors, epistemic circularity is sometimes benign. Like Alston, Bergmann is concerned to explain how we know (or are justified in believing) that sense perception is reliable, and to do so without recourse to a malignant epistemic circularity. Inspired by Reid, Bergmann claims that it is a first principle that sense perception is reliable, and we have a faculty of common sense that enables us to know first principles non-inferentially. Whatever the merits of Bergmann's claim, it won't help here; it is doubtful that we have a rational faculty that "sees" the correctness of a theoretical explanation of what makes basic things justified.

quate consideration in favor of her belief that b_T is justified, without illegitimately relying on b_T.

Let me try to get clearer about what the problem is.

2.4.2 What went wrong: dramatic version

Here is a brief, although misleading, version of the argument just presented:

> The question "What justifies that?" can be asked of any belief. Does that question always require an answer which leads to a further question? The foundationalist says "No". The foundationalist argues that at some point the questions stop. When we reach a basic thing, the foundationalist response is: "Nothing justifies that; it just is justified". The trail of questions ends at that basic thing; there are no more things left to question.

> If "What justifies that?", when asked of any belief, requires an answer, then beliefs about beliefs can be questioned too: for example, a belief that some thing is basic– justified but not justified by anything. So the foundationalist is driven to posit T, an explanation how it can be that basic things are justified.

> If "What justifies that?", when asked of any belief, deserves an answer, then the belief that a belief in T is justified can be questioned too. That begins another line of questioning. But if a belief in T cannot be brought in, that line of questioning reaches a point where the theorist has no answer.

Although dramatic, this version of the argument is misleading. In the original version, the theorist wonders about one of her theoretical beliefs. She is not defending her belief in answer to a challenge. She is not trying to convince someone else to change his mind. She is not trying to satisfy a doubt about her belief. The issue is not whether she has what it takes to answer a challenge, or to convince someone else, or to satisfy a doubt. Instead, the issue is whether, in wondering about her beliefs, she can feel satisfied that all is well. I claim that the answer is no. She eventually reaches a point where she needs to rely on b_T, but she should not do so because she is wondering whether b_T is justified.

Here is an analogy to explain what I am mean about wondering, and how it differs from doubting or answering a challenge. Suppose, before going to bed, I look into my

child's room to see if he is ok. In the darkness, I see a vague form curled up in the corner of the crib, and no sign that anything is amiss. I go to bed satisfied. I had no doubt or suspicion that something was wrong. Nor was I am trying to get evidence to answer a person who believes that something is wrong. Otherwise, I might turn the lights on or tickle him to see if he responds. Instead, I am satisfied by finding just a little evidence that all is well. I am looking for reassurance, but of a very weak sort.

The problem is that the theorist can't even get that weak sort of reassurance. To be satisfied when wondering whether p, in the sense I am discussing, is to be satisfied that there is at least a weak consideration in favor of continuing to believe that p. Unfortunately for the theorist, she cannot get that satisfaction without illegitimately bringing in her belief b_T.

2.4.3 What the argument does not show

This is not a skeptical argument. The conclusion is not: the theorist's beliefs are not justified. This is not a skeptical argument, even if wondering temporarily defeats the justification of some of her beliefs. For the theorist can easily and reasonably stop wondering. For example, she might, quite reasonably, think to herself that she is wasting her time wondering whether Jb_T, and stop wondering whether Jb_T. And when she stops wondering, the temporarily defeated justification will presumably come back.

Neither is the conclusion that the theorist's beliefs are not reasonable. Wondering temporarily makes many– or most– of her beliefs unreasonable. But she can easily and reasonably stop wondering. And when she stops wondering her beliefs will presumably become reasonable again.

2.4.4 What the argument shows

I believe that the argument shows that vanilla foundationalism, when coupled with an explanation T, does not sustain reflection. A vanilla foundationalist who critically

scrutinizes her theoretical beliefs cannot satisfy herself that they are correct. She cannot find even a weak consideration to satisfy herself that her belief in her explanation T is justified, without illegitimately bringing in her explanation T.

Someone might rush to deny that this is a problem:

> If a person starts wondering about his entire epistemological theory, many– or most– of his beliefs may be called into question. And he may not be able to satisfactorily answer all questions which arise, without appealing to his epistemological theory. But surely this does not show that his epistemological theory is flawed. Similarly, if a vanilla foundationalist starts wondering whether a belief in T is justified, many– or most– of her beliefs may be called into question. And she may not be able to satisfactorily answer all questions which arise, without appealing to T. But surely this does not show that vanilla foundationalism is flawed.

> We should not conclude from the argument that vanilla foundationalism is flawed. Instead, we should conclude that wondering about T is like wondering about an entire epistemological theory.

In response: even if we accept the objector's claim that nothing is amiss when a theorist wondering about his entire epistemological theory finds himself unable to answer all questions which arise, it is clear that the present case is different. In the present case, the vanilla foundationalist is not wondering about vanilla foundationalism. She accepts vanilla foundationalism on the basis of the regress argument. She is not questioning her commitment to vanilla foundationalism. Rather, she is considering a further question that arises once vanilla foundationalism is accepted. That question— what makes a basic thing justified— does not even arise for non-foundationalist theories. Vanilla foundationalists can, and do, argue about the merits of rival answers to that question independently of acceptance of vanilla foundationalism. The problem is that the vanilla foundationalist cannot be satisfied by her answer to that further question.

2.4.5   Four solutions

I see four possible ways out for the vanilla foundationalist.

The first way out is to offer an explanation T, but refrain from wondering whether or not a belief in T is justified. If the theorist refrains from wondering, then she will

never find herself in the unsatisfactory position I have described. This way out seems rather closed-minded. Surely, it is part of the theoretical project to reflect on one's theoretical beliefs. But even if there is nothing unreasonable or closed-minded about refraining from wondering whether or not her belief in T is justified, this response leaves the theorist no better off. For if the theorist were to wonder, and thought hard enough, she would recognize her unsatisfactory position: not being able to offer an adequate consideration in favor of her belief that her belief in T is justified. It does not help to advise her to refrain from wondering about her theoretical beliefs. That's like advising: "You should accept foundationalism— just don't think about it too hard."

The second way out is to hold no beliefs about the epistemic status of other beliefs. In that case, there would be nothing for the theorist to wonder about. I don't think there is anything unreasonable about this way out. There are probably many reasonable people who hold no beliefs about the epistemic status of other beliefs. Not everyone is a philosopher. However, no one who takes this way out can believe in foundationalism, because believing in foundationalism requires holding beliefs about the epistemic status of other beliefs. To take this way out is to say something like: "Vanilla foundationalism is true but you shouldn't believe in it."

The third way out is to embrace the circularity I have identified. The theorist then says that, while wondering whether Jb_T, she may legitimately rely on b_T. This way out is unacceptable to the theorist. It undermines the reasons to be a vanilla foundationalist. The theorist rejects coherentism; she argues for her view partly on the grounds that circular chains of justification do not suffice to confer justification. But to take this way out is, if not to endorse coherentism outright, at least to hold a view difficult to distinguish from coherentism. If the theorist relies on b_T in satisfying herself that Jb_T, then she uses one of her beliefs as a consideration in favor of her belief that that belief is justified. But if a belief b is a consideration in favor of a belief that b is justified, then b is a consideration in favor of b. (For a consideration in favor of a belief that b is justified, is a consideration in favor of a belief that b is true or likely to be true.) If b can be a consideration in favor of b, then I have trouble

seeing why the theorist would deny that b can justify itself.

The fourth and final way out is more promising. The theorist could maintain that, even while wondering about her belief in T, she can offer an adequate consideration in favor of a belief that a belief is basic: a belief that a belief is basic is itself basic. What adequate consideration has the theorist in favor of *that*? A belief that a belief that a belief is basic is basic is basic. And so on. The theorist thus does not violate (R) while wondering about her belief in T.

However, if the theorist chooses this last way out, she loses her motivation for proposing a T in the first place. I suggested that a T is needed in order to satisfy (R) when the theorist believes of some b that it is basic. If (R) can be satisfied by bringing in a basic higher order belief than a T is not needed. To take this way out is to be satisfied to be an obscure foundationalist– a vanilla foundationalist who refrains from offering an explanation why Jx when x is basic.

I must admit that obscure foundationalism is strange, and rather difficult to endorse. I expect that you agree. But, short of abjuring vanilla foundationalism, I see no better way out. If we wish to be vanilla foundationalists, we should be obscure foundationalists. So to the extent that obscure foundationalism is unsatisfactory, vanilla foundationalism is too.

Part II

The practical problem of
skepticism

Chapter 3

You can't get there from here

Abstract

Rational acts are guided by beliefs. This observation supports a response to skepticism: some skeptical arguments can be rationally avoided.

Introduction: the problem of skepticism

Skeptical arguments are tantalizing. Skeptical arguments display apparently compelling reasoning to implausible conclusions such as

No one can know anything about other minds,

or

No one can have a justified belief about the external world.

A few philosophers accept skepticism. The rest face a question about each skeptical argument: What is wrong with that argument? Suggested answers abound: that argument begs the question, exhibits a misleading form of context sensitivity, has false premises, appeals to implausible epistemic principles, fails to transmit warrant from premises to conclusion, and so on. Some philosophers just hope that each argument has a defect. Whatever the approach, non-skeptics agree that all skeptical arguments are flawed.

Just two options seem available: accept skepticism in some form, or agree that every skeptical argument is flawed. One must either become a skeptic or face the

61

problem of satisfactorily explaining what is wrong with each skeptical argument.

There is a third path. One need not choose between becoming a skeptic or accepting the problem of locating flaws in each skeptical argument. I will argue that it is rational to avoid certain skeptical arguments.

I am not alone on the third path. Hume, for example, can be seen as a non-skeptic who does not look for flaws in each skeptical argument. Instead, Hume claims that humans are psychologically immune to skepticism. However, Hume's views are difficult to interpret. It is not clear whether he offers a stable non-skeptical position.

The difficulties facing a Hume-like response to skepticism are vividly seen by considering a proposal by P. F. Strawson.[61] Inspired by Hume, Strawson claims that, by nature, we cannot become skeptics. Thus, thinks Strawson, skepticism is idle. It is idle to seriously entertain skeptical arguments, and it is idle to rationally confront skeptical arguments: "... skeptical arguments and counter-arguments [are] equally idle— not senseless, but idle— since what we have here are original, natural, inescapable commitments which we neither choose nor could give up."[61, p 27-8] Strawson is concerned with skepticism about external bodies, other minds, the past, and the use of induction; he is claiming that our commitments to the use of induction, the existence of external bodies and so on, are commitments that we, by nature, cannot give up.

Interpreted rather bluntly, Strawson's central claim is that nothing can bring us to give up our "natural, inescapable commitments". Since skeptical arguments aim to do just that, skeptical arguments cannot succeed. If skeptical arguments cannot succeed, then it is idle to seriously entertain a skeptical argument, and it is idle to try to refute a skeptical argument.

This approach to skepticism faces a difficulty. Suppose I can't help believing that the FBI is out to get me. When faced with an apparently compelling argument to the contrary, I dismiss the argument as "idle". I seem unreasonable. Similarly, suppose I can't help believing that some skeptical conclusion is false. When faced with an apparently compelling skeptical argument to the contrary, I dismiss the argument as "idle". Again, I seem unreasonable. The problem here is that it seems unreasonable

to dismiss an argument simply because one cannot help believing that the conclusion is false.

In later writing, Strawson appears, sensibly, to reject the blunt interpretation of his view:

> It is not merely a matter of dismissing the demand for justification of one's belief in a proposition on the ground that one can't help believing it. That would be weak indeed. The position is, rather, that the demand for justification is really senseless. [30, p 370]

To explain the point, Strawson adopts an image from Wittgenstein's *On Certainty.*[72, §96-§99] According to Strawson, our "natural inescapable commitments" are like the banks of a river; they constitute "the boundary conditions of the exercise of our critical and rational competence..." [30, p 371] The notion of a boundary condition is suggestive, but obscure. What is a boundary condition of rational competence? Further, why think that, for example, a commitment to the existence of external bodies is such a boundary condition? Admittedly, Strawson's suggestion is briefly sketched– but one would prefer a more transparent proposal.

Perhaps ambitiously, I will attempt a better response to skepticism. A Humean suggestion is that some beliefs cannot be given up. However, the mere fact that a belief cannot be given up seems insufficient reason to ignore an argument to the contrary. Strawson adds Wittgensteinian complexity: some commitments can neither be given up nor intelligibly questioned. But this thought is obscure. I will defend a simpler claim: some beliefs cannot be given up, while acting rationally. From this claim, I will argue that, ordinarily, no one, while acting rationally, can be convinced by certain skeptical arguments. And I will conclude that, to avoid acting irrationally, it is rational to avoid certain skeptical arguments.

3.1   The plan

This chapter is structured as follows. In section 3.2, I clarify the concept of lasting belief. In section 3.3, I claim that rational acts require lasting beliefs. Then, in section

63

3.4, I use this claim to argue that, ordinarily, no one, while acting rationally, can be convinced by a skeptical argument. I proceed by first describing a certain extreme kind of skeptical argument, and arguing that no one, while acting rationally, can be convinced by an extreme skeptical argument. Then I argue that this result extends, in a weaker form, to other skeptical arguments. In section 3.5, I distinguish the proposal from superficially similar anti-skeptical proposals. After that, in section 3.6, I present an argument for the key claim of this approach to skepticism– that rational acts require lasting beliefs. In section 3.7, I reply to some objections. Finally, I draw some conclusions about the problem of skepticism, and list some questions for future research.

3.2 Lasting belief

Some beliefs last longer than others. If Wallace and Gromit both believe that their tenant is honest, but Gromit will soon change his mind, then Wallace's belief will last longer than Gromit's. If the rock that is about to land on W. E. Coyote's head will cause him to lose a few dozen beliefs, then those beliefs will not last long.

These are claims about the length of time that a belief will last. In claiming that Smith's belief will last three days (or three months or until the next full moon) I am making a claim about the length of time that Smith will continue to hold that belief. But, it should be emphasized, I am *not* making a claim about the robustness of Smith's belief, or the tenacity with which Smith holds that belief, or the belief revision policy Smith adopts, or Smith's beliefs about what he will believe. To avoid misunderstandings, I will explain why.

1. *Robust belief*

A person's belief is robust just in case he would not easily give up that belief.[1] But even if Nora robustly believes that Nestor is loyal, her belief might not last long. For she might be about to give up her belief; she might be about to face incontrovertible evidence of Nestor's treachery. Or, Nora may have beliefs that are very long lasting but not at all robust. Even if, with the slightest prodding, she would stop believing that her life has meaning, if she will never be so prodded, then that belief will last as long as any other belief she has.

2. *Stubborn belief*

Eeyore may decide to stubbornly hold on to a belief. That may make his belief last longer. But it need not. For he may soon give up his stubborn attitude, and then his belief. Or, despite his stubborn attitude, he may soon find himself giving up the belief anyway.

3. *Belief about what one will believe*

Herman may believe that Longfellow was a literary lion, and Herman may believe that next year he will continue to believe that Longfellow was a literary lion. Still, these beliefs may not last long. For Herman may read *Evangeline* tomorrow and change his mind, both about his future beliefs and about Longfellow. Moreover, Herman may have a long-lasting belief that Longfellow was a literary lion, but not have any belief at all about what he will believe about Longfellow next year.[2]

4. *Conservatism*

According to a conservative belief revision policy, a person should continue believing something as long as she finds no reason to doubt it.[3] Perhaps someone who follows a conservative belief revision policy has longer lasting beliefs than someone who does not. But perhaps not. Even if Emma continues to believe things that she finds no reason to doubt, she may regularly find reasons to doubt her beliefs; many of her beliefs might not last long.

---

[1]This formulation comes from [68].

[2]See [27] for recent discussion about beliefs about future beliefs.

[3]This formulation comes from [32, p 23].

3.3 Rational acts require lasting beliefs

You begin typing an e-mail. As you reach the end of the message, you look back and notice that some words at the beginning have vanished. "A computer glitch," you mumble to yourself, while making corrections. Then you notice that some other words have disappeared. Frustrated, you type them again. Then you notice that some more words are gone...

Could you compose a coherent message? Living a life with short-lived beliefs would be like trying to compose a coherent e-mail with this computer. Could you behave reasonably if beliefs disappear just when you need them?

To make this thought more explicit, I propose the following constraint on rational action:

Principle (R) *A person's act is rational only if she has a belief when that act is initiated which lasts until that act is finished.*

Note that (R) is a rather weak claim. (R) does not single out which belief must last long enough. (R) just says that some belief must last long enough.

Here is an example to suggest that we should agree to (R).

> 1. Believing that if P then Q, you come to believe that P. You then start to put these beliefs together. In the midst of your reasoning, you suddenly stop believing that P. Still, you continue on, and conclude your reasoning by forming a belief that Q.

This, I believe, is irrational. It is irrational for you to form the belief that Q on the basis of this deliberation. The problem is clear: You no longer believe that P at the end of your reasoning. At the last step of your reasoning you come to believe that Q, but you need to believe that P at this moment for it to be rational for you to come to believe that Q in this way. Your deliberation is not rational because your belief that P fails to last long enough. If a deliberation is a kind of act (a mental act?), then that act is irrational because your belief that P fails to last until that act is finished.

Some may quibble. For example, does rationality demand that, during your deliberation, you continuously believe that P, or can you waver between believing that

P and not? Perhaps you can waver, as long as you believe that P at the end point of the deliberation. Although this and other complaints may bring us to tinker with (R) let me set this and other worries aside until sections 3.6 and 3.7, where I will scrutinize (R) more carefully. For now, I will suppose that this example shows (R)– or something very much like it– to be a reasonable hypothesis.

Other examples support (R). For instance:

> 2. One evening I develop a mild headache. Unfortunately, there is no aspirin in my home. I wonder if the pharmacy is open. It is not a very strong headache, and it is a long way to the pharmacy, so it wouldn't be worth going out only to find the pharmacy closed. A glance at my watch tells me that it's 8:45. I conclude that the pharmacy is open, and decide to go there. On the way, I notice that the street looks empty, and I wonder if the pharmacy is still open. I glance at my watch. It still says 8:45. I realize that my watch has stopped long ago. I decide that I was mistaken to believe that the pharmacy was open, and I start believing that the pharmacy is closed. Then I continue going there.

This behavior is irrational. My decision to go to the pharmacy required the belief that it was open. When this belief disappears, it is irrational to continue to the pharmacy on the basis of this earlier decision. Given that it is irrational to continue to the store, we can tentatively make a further claim: my act of going to the store is irrational because the belief that the store is open failed to last until I got there.

It is not clear that we can justifiably move to the further claim. Perhaps, in some sense, my act of going to the store is rational because my initial decision to go to the store is rational, given my beliefs when I decided to go. Or, perhaps my continuing to the store is irrational not because I lose the belief that the store is open, but because I gain the belief that it is closed. For these reasons, the second example is not as clear as the first example.

Still, the examples bring out an obvious but important point: acts take place over time. More controversially, the examples suggest that the rationality of an act is sensitive to what happens over time; in particular, rational acts require a certain stability in belief over time. For, in the examples, a particular belief fails to last for the duration of an act, and so, arguably, the act is irrational. In other words, the act

is irrational because a particular belief does not last until that act is finished. The examples thus support (R) because (R) says a person's act is rational only if she has some belief or other which lasts throughout the performance of that act:

Principle (R) *A person's act is rational only if she has a belief when that act is initiated which lasts until that act is finished.*

In fact, the examples support a claim more specific than (R). They suggest that the required belief is not just any old belief or other. The required belief is a particular belief that is closely connected to the act in question. The required belief is apparently a belief that supports– or even guides– the act. But two examples are not enough to understand how to make (R) more specific.

Indeed, two examples are not enough to show that we should accept (R), let alone a more specific version of (R). But the examples do provide some reason to accept (R). That is enough for now. More needs to be said, and, in sections 3.6 and 3.7, I will do so. There I will argue directly for (R), and defend it against some objections. But first, I will use (R) in responding to skepticism. Seeing it in action will help clarify it. Seeing it in action will also permit us to evaluate it in the context of the anti-skeptical role that it is supposed to play.

3.4 A response to skepticism

(R) is a claim about what is required for an act to be rational. So it may be puzzling how this claim can be used to respond to skeptical arguments that lead to conclusions like: no one can know anything about the external world, or no one can have justified beliefs about other minds. Skeptical conclusions concern our epistemic situation, while (R) concerns action.

To apply (R), I will characterize skepticism in an unfamiliar way. Rather than focusing attention on particular skeptical arguments, I will instead discuss the changes in belief state that occur when a person is convinced by a skeptical argument. My suggestion is that rational acts cannot span certain changes in belief state. More precisely, there are pairs of belief states, s_1 and s_2, such that no one, during a rational

act, can be first in s_1, then in s_2. Even so, it may be rational to be in either belief state s_1 or s_2. And it may be possible to move from s_1 to s_2 by rational belief changes. Yet, I claim, it is not possible to move from s_1 to s_2 while acting rationally.

In order to develop this proposal I need to explain what changes in belief state occur when someone becomes convinced by a skeptical argument, and why those changes in belief state cannot occur during a rational act. For heuristic reasons, I will begin with an extreme case. Then I will extend the proposal to more commonly discussed cases.

3.4.1 Principle (R) applied to extreme skepticism

Next I present an illustrative example: I will argue that principle (R) helps provide a response to an extreme form of skepticism. Consider an argument that is intended to bring a person to give up all of his beliefs. (This may have been the point of ancient skeptical arguments; some say that the Pyrrhonian skeptics presented arguments that were supposed to bring their dogmatist opponents to give up all of their beliefs.) A person who becomes convinced by such an argument enters an extreme skeptical state:

Definition 1 *S is in an **extreme skeptical state** if and only if: for every proposition w she neither believes w, nor believes -w; and she is able to form beliefs.*

The last clause excludes things like rocks and trucks, which, though lacking beliefs, are not in an extreme skeptical state.

Now (R) supports the following anti-skeptical claim:

Anti-Skeptical Claim *No one can enter an extreme skeptical state while acting rationally.*

For, suppose a person will perform some act. And suppose that she will enter an extreme skeptical state while performing that act. Then, no belief of hers will last until that act is finished. By (R), that act is not rational. Therefore, no one can enter an extreme skeptical state while acting rationally.

The Anti-Skeptical Claim helps in responding to arguments that are supposed to bring a person to give up all of his beliefs. Since studying an argument is an

act, if the studying is rational, then, by the Anti-Skeptical Claim, a person cannot enter an extreme skeptical state while studying. No one, while rationally studying an argument, can enter an extreme skeptical state.

Someone may hurry to object that, even if no one can enter an extreme skeptical state while rationally studying a skeptical argument, a person might, acting rationally, be convinced by an extreme skeptical argument. For someone might be convinced *by* an extreme skeptical argument, although he is not convinced *while* studying it. Gerry, we may suppose, quite rationally studies an argument, and fails to enter an extreme skeptical state while studying. But then, three days later, she finds herself thinking about the argument again and manages to reason herself into an extreme skeptical state. Hasn't Gerry, while acting rationally, been convinced by an extreme skeptical argument?

No. If, when Gerry turns her thoughts back to the argument, she is continuing her act of studying the argument, then her act of studying is irrational. If, while thinking about the argument, she is performing some other act– like walking to the store– then that act is irrational. Either way she acts irrationally. And if Gerry is not performing any act at all, then it may be that she has, quite rationally, managed to reach an extreme skeptical state. But that is not the same as being convinced by an extreme skeptical argument, while acting rationally.

I am suggesting that no rational act can be performed while certain belief changes are happening. This is so, even if, considered on their own, those belief changes are rational. More precisely, even if it is rational to move from belief state s to an extreme skeptical state e, a transition from s to e could not occur during a rational act. Thus, someone could perhaps reach an extreme skeptical state by rational belief changes; however, he could not do so while performing a rational act.

I will respond to further objections in section 3.7.

The Anti-Skeptical Claim helps a rational person avoid skepticism. According to the Anti-Skeptical Claim, no one can enter an extreme skeptical state while acting rationally. But a rational person should not act irrationally. So a rational person

should not let himself enter an extreme skeptical state while he is performing some act; a rational person should not let himself be convinced by an extreme skeptical argument while he is performing some act. For example, suppose that rational Rhonda decides to spend an evening analyzing the arguments reported by Sextus Empiricus. Rhonda should not let herself be convinced, and so enter an extreme skeptical state. For to do so would mean that her evening study is irrational. Thus, assuming that we are rational, we should not let ourselves be convinced by an extreme skeptical argument while performing some act.

I will now turn to more commonly discussed forms of skepticism.

3.4.2   The proposal extended to other forms of skepticism

Readers are probably used to thinking that to be convinced by a skeptical argument is to accept a skeptical conclusion like: no one is justified in believing anything about the external world; or, no one knows anything about other minds. I instead offer an unorthodox characterization: when convinced by a skeptical argument, a person gives up some portion of his beliefs. I will motivate this unorthodox characterization in the next section. But first, granting this characterization, let us see if the approach will work.

Consider a person who gives up some portion of his beliefs. Perhaps, for example, he gives up all of his beliefs about the external world, or about the past, or about other minds, or even about the Red Sox.

Inspired by the last section, one might try to argue that no one, while acting rationally, can give up all of his beliefs about the external world. Or, no one, while acting rationally, can give up all of his beliefs about the past, or other minds, or the Red Sox. We would just need a version of (R) that says that every rational act requires several beliefs: one about the external world, one about the past, one about other minds, and another about the Red Sox.

This approach looks hopeless. It would be difficult to argue that no one, while acting rationally, can give up all of his beliefs about the Red Sox. For, obviously, I can take a rational walk to the supermarket without requiring any of my Red Sox

beliefs. Moreover, it seems clear that I can take a rational walk to the supermarket without requiring any of my beliefs about other minds. And I can perhaps do the same without relying on any of my beliefs about the past or about the external world. (Maybe I could act based on beliefs about my present perceptual experiences.)

Despite this dead end, there is a tenable, interesting claim in the neighborhood. Note that ordinarily when we act, we do rely on beliefs about the external world. In the earlier pharmacy example, my act relied on a belief that the pharmacy was open; my act relied on a belief about the external world. And, in acting we ordinarily do rely on beliefs about the past (should I trust her again?) and about other minds (how would he feel if I did that?).

If, ordinarily, our rational acts require beliefs about the external world, then, ordinarily, a person, acting rationally, cannot give up all of his beliefs about the external world. And if, ordinarily, our acts require beliefs about the past or other minds, then, ordinarily, a person, acting rationally, cannot give up all of his beliefs about the past or other minds.

Of course ordinarily is not always. This approach concedes the possibility, while acting rationally, of giving up all beliefs about the external world, the past, or other minds. Thus this approach is weaker than the approach in the last section, in which it was claimed that no one, while acting rationally, could give up all of his beliefs.

Recall, however, that in the last section, I left it open that a person, while performing no act, might rationally enter a state without any beliefs. Here, I widen that opening: while performing certain rational acts, a person might give up certain portions of his beliefs.

The size of that opening depends on the portion of beliefs being discussed. I suspect that most of our acts rely on beliefs about the external world. However, I doubt whether Red Sox beliefs are as central to our actions. This approach is more effective for some classes of belief (like external world beliefs) than other classes of beliefs (e.g. Red Sox beliefs).

This approach can handle interesting cases. Suppose you decide to spend the evening sitting by your fireplace, carefully reading a book about skepticism. If this

is an ordinary case, where the rationality of your act depends on your having lasting beliefs about the external world (that you are reading a book, or that you are sitting by your fireplace), then if your act of sitting and reading is rational, you will not be convinced by the book to give up all of your beliefs about the external world.

Note that this approach requires changes to (R). In the external world case, for example:

Principle (R+) *Ordinarily, a person's act is rational only if she has a belief about the external world when that act is initiated which lasts until that act is finished.*

When arguing for (R), I will also argue for principles like (R+).

3.4.3   Accepting a skeptical conclusion

In the foregoing response to skepticism, I described being convinced by a skeptical argument as giving up some or all beliefs. But it is not obvious that becoming convinced by a skeptical argument is best characterized as giving up some or all beliefs. One would have thought that to become convinced by a skeptical argument is to come to believe a skeptical conclusion. A person who comes to believe a skeptical conclusion does not give up beliefs; he gains a new belief– belief in the skeptical conclusion.

These two ways of understanding what happens when a person is convinced by a skeptical argument not only differ, they appear incompatible. In this section, I will draw them together. I can't consider every skeptical conclusion here, so I will consider two representative examples: no one can have justified beliefs about the external world, and no one can have any knowledge about the external world.

Think about Rex, a rational person with an ordinary array of beliefs about the external world. One day, an argument in philosophy class convinces Rex that no one can have justified beliefs about the external world. At this point, can Rex rationally maintain his beliefs about the external world? Can Rex, rationally, believe that no one can have any justified beliefs about the external world, at the same time as Rex maintains beliefs about the external world? If not, then if Rex remains in a rational belief state when he comes to believe that no one can have a justified belief about

the external world, then he gives up all of his beliefs about the external world.

Linking principles

In more general terms, the following principles link the two ways of understanding what happens when a person is convinced by a skeptical argument:

(I) a person cannot, rationally, at the same time hold beliefs about the external world and believe that no one can have justified beliefs about the external world.

(II) a person cannot, rationally, at the same time hold beliefs about the external world and believe that no belief about the external world can amount to knowledge.

Given (I), for example, a person who holds some beliefs about the external world, and who then becomes convinced that no beliefs about the external world can be justified, must, on pain of irrationality, either give up his beliefs about the external world, or give up his belief that no beliefs about the external world can be justified.

(I) and (II) are not outlandish. In a rational belief state, beliefs about beliefs "fit together" with other beliefs. For instance, a person in a rational belief state does not both have a belief and believe that that belief is irrational. So, it is reasonable to think that there are true principles which, like (I) and (II), require that beliefs about the epistemic status of beliefs "fit together" with other beliefs.

I will next briefly explain why (I) and (II) are defensible.

(I) is plausible

There are many ways of understanding the skeptical conclusion that no belief about the external world can be justified, because there are many ways to understand what the status of being justified amounts to. On some of these ways, (I) is plausible. Here are two examples.

A. According to one view, a belief is justified just in case it is permissible to hold that belief. So suppose the skeptical conclusion is that it is never permissible to hold a belief about the external world. A person who accepts this conclusion is a short step from having some belief B, while believing that it is not permissible to have belief B.

74

That seems unreasonable- to believe that you have a belief that you shouldn't have.[4]

B. According to another view, your belief that p is justified just in case a belief that p better fits your evidence than a belief that $\neg p$.[5] So suppose the skeptical conclusion is that, for any proposition p about the external world, no belief that p can ever fit a person's evidence better or worse than a belief that $\neg p$. A person who accepts this conclusion (and who has external world beliefs) is a short step from having a belief that p, while believing that her belief that p can never fit her evidence any better or worse than a belief that $\neg p$. This, arguably, is irrational even if that person believes that her belief that p is morally or prudentially justified. For, arguably, you should not believe that p while believing that you hold that belief solely for moral or prudential reasons.[6]

(II) may be plausible

On some ways of thinking about knowledge, (II) is plausible too. On others, (II) is rather implausible.[7] Here is one example of each.

A. Suppose that the skeptic argues that no belief can amount to knowledge because no belief can be certain, and knowledge requires certainty. It seems unlikely that coming to believe that no belief about the external world can be certain requires you to give up your beliefs about the external world. For it does not seem unreasonable to believe that p, although you are not certain whether p, and you believe that you can never be certain whether p. For instance, it seems reasonable to believe that no brontosaurus was kelly green even though, given the limited evidence, you believe you will never be certain about that.

---

[4]While this view of justification supports (I), this view may render my anti-skeptical approach unnecessary. If a person who ought to believe p, can believe p, then, on this view, if you can't help believing p, then that belief is justified. If so, then if I cannot give up some belief about the external world, then that belief is justified. However, it is very unclear whether a rational person who ought to believe p can believe p. So this view of justification is, I think, worth mentioning. (R. Stalnaker offered this comment to me.)

[5]See [25]. Note that Conee and Feldman discuss justified attitudes (like believing or suspending judgment) rather than justified beliefs. For simplicity I ignore this important distinction.

[6]Adler[1] argues for a thesis similar to (I): it is impossible to believe that p, while, in full awareness, believing that you lack adequate evidence that p.

[7]It does not matter if (II) is false. If the argument of Chapter 4 is correct, then even if (II) is false, skepticism about knowledge of the external world is harmless.

B. According to one view, if a person believes that p, then even if she does not know that p, she reasons as if she knows that p.[8] If you falsely believe that you left the headlights on, although you do not know that you left the headlights on, you will reason as if you do and decide to go back to turn them off. Now suppose you come to believe that no belief about the external world can amount to knowledge. Arguably, it is unreasonable to believe that your belief that p cannot amount to knowledge, and then reason as if you know that p. That would be like treating the naked emperor as if he is wearing clothes.

A conditional claim

Discussing principles like (I) and (II) is a way to discuss a broader issue: how serious or threatening to our world view would it be to accept a skeptical conclusion? Would it hold only theoretical interest? Or would it turn our view upside down?

Perhaps, in accepting a skeptical conclusion, we merely learn something theoretically curious about our epistemic situation. We thought that some of our beliefs were justified, or that we had knowledge, but that was mistaken. Beyond revising a few rarefied theoretical beliefs, we may, rationally, go on believing what we believed before.

Or, perhaps accepting the conclusion would be deeply affecting. Perhaps accepting a skeptical conclusion rationally requires us to revise ordinary, non-theoretical beliefs. If so, accepting the skeptical conclusion would deeply affect how we conceive the world.

Absent a full defense of principles like (I) and (II), I can at least make a conditional claim about the anti-skeptical approach I am proposing.

If there are true principles like (I) and (II), then accepting some skeptical conclusions would be deeply affecting. If there are true principles like (I) and (II), then accepting some skeptical conclusions would rationally require us to give up some or all of our ordinary, non-theoretical beliefs. If there are true principles like (I) and (II), then my unorthodox characterization (of becoming convinced by a skeptical argument as giving up of some or all beliefs) is apposite, and my anti-skeptical approach

---

[8]Williamson[71, pp 46-8] suggests that "believing p is, roughly, treating p as if one knew p."

is appropriate.

If, on the other hand, there are no true principles like (I) and (II), then accepting a skeptical conclusion would only affect beliefs about beliefs. If accepting a skeptical conclusion would only affect beliefs about beliefs, then my anti-skeptical approach is probably not relevant. (It is not relevant unless accepting a skeptical conclusion involves giving up some beliefs about beliefs needed for rational action.) But if accepting a skeptical conclusion would only affect beliefs about beliefs, then one might wonder whether skepticism deserves a response. If only beliefs about beliefs are affected, then, I am tempted to say, skepticism does not seem so threatening after all.

More needs to be said. For present purposes, I only claim this: the plausibility of some interpretations of principles (I) and (II) justifies interest in the present proposal.

3.5 Misinterpretations of the proposal

Next I clarify the present anti-skeptical proposal by comparing it to other approaches.

3.5.1 Skepticism is bad for you?

You probably have not yet heard the funniest joke in the world. That's the one that's so funny that anyone who hears and understands it will quickly laugh himself to death. If you hope to stay alive, you ought remain far from the funniest joke in the world. If you hope to stay alive, it would be irrational to listen to the funniest joke in the world.[9] Someone might think that the present proposal is similar. Someone might think that the present proposal is this: a skeptical state is bad for you, therefore it is irrational to enter a skeptical state.

This is not the present proposal. Even if a skeptical state is good for you, according to the present proposal, no one, ordinarily, while acting rationally, can enter a skeptical state. This is so because rational acts require beliefs that last. The proposal does not rest on any claim that a skeptical state is bad for you, or irrational,

---

[9]An episode of *Monty Python's Flying Circus* showed me the dangers of this joke.

or otherwise undesirable.

It may be instructive to recall that the "goal" of the Pyrrhonian skeptic is *ataraxia*–"tranquility" or "freedom from disturbance". The Pyrrhonian leads his dogmatist rivals to *ataraxia* in the following way. First, the Pyrrhonian uses the dogmatist's own beliefs and methods to bring the dogmatist to suspend judgment on some (or all) propositions. After that, *ataraxia* follows, as if by serendipity. In *Outlines of Pyrrhonism*, Sextus Empiricus tells the following story to display the transition to a state of *ataraxia*.[PH I 28] The painter Apelles is trying to paint the foam on a horse's mouth. After many unsuccessful attempts he gives up in frustration and throws his sponge at the canvas, producing the effect perfectly. The moral is unclear, but we might see this story as suggesting that *ataraxia* is not something one reaches by trying to get there. When it comes, it comes as if by accident.

It may be a stretch, but I think Sextus Empiricus' tale displays an insight connected to the present proposal: setting a goal to achieve *ataraxia* will not get one there. One could get there by accident or mistake, but no one can get there by design. In present terms, we might explain this as follows: no belief can survive a trip to *ataraxia*; so, no rational act will suffice to achieve *ataraxia*.[10]

3.5.2 A skeptic can't live a life?

In ancient times, critics of Pyrrhonian skepticism complained that no one can live a life as a skeptic. It is impossible to live a life without belief, the critics said.[16] In the present framework, this suggestion is that no one without beliefs can act rationally. This suggestion is related to (and part of the inspiration) of the proposal I am making, but it is not the same proposal. In its extreme form, my proposal is that a complete loss of belief cannot occur during a rational act; a transition from a state with beliefs to a state without beliefs cannot occur during a rational act. This is not the same as saying that a person without beliefs cannot perform a rational act. The point of

---

[10]G. Striker [62] helpfully discusses the complicated role of *ataraxia* in Pyrrhonian skepticism. She thinks that the Pyrrhonian reasoning does not require– and would be better off without– talk of *ataraxia*.

my proposal is to try to show that a person with beliefs, who always acts rationally, and who is always performing some act or other, will never reach a skeptical state. Whether or not a person without beliefs can act rationally— or can even act at all— is a different question.

3.5.3 Reasoning your way to skepticism

Someone might understand the proposal this way: no one can rationally accept skepticism about the external world because the act of reasoning your way to skepticism about the external world is irrational. This is an interesting claim that I wish I could make. However, it is not the proposal I am making. It just does not seem plausible that acts of reasoning always require beliefs about the external world. So it is not plausible that focusing on acts of reasoning will help answer skepticism about the external world. It *is* plausible that the acts we perform every day (cooking dinner, reading a book and so forth) require beliefs about the external world. That is why I am claiming that no one, while rationally performing such ordinary acts, can be convinced about skepticism about the external world.

3.6 Arguing for (R)

So far, we have seen a response to extreme skepticism based on principle (R):

Principle (R) *A person's act is rational only if she has a belief when that act is initiated which lasts until that act is finished.*

We have also seen that variations of (R) help with other more interesting forms of skepticism. Here is an example of such an (R) variation:

Principle (R+) *Ordinarily, a person's act is rational only if she has a belief about the external world when that act is initiated which lasts until that act is finished.*

In this section I will defend (R) and (R+). I will not discuss other variations of (R); I will assume that (R+) is a representative example.

It should be emphasized that while (R) gives a necessary condition for any act to be rational, (R+) does not. By (R), every rational act requires a lasting belief.

As far as (R+) goes, some rational act may fail to require a lasting belief about the external world.

I will defend (R) by reflection about the nature of action. In arguing that every rational act requires some lasting belief, I will also argue that the acts that we ordinarily perform require lasting beliefs about the external world. If so, then (R+) is correct.

3.6.1 A failed argument

There is a tempting but flawed way to defend (R). Plausibly, a person cannot rationally change all of her beliefs at once. That is, a rational revision from belief state s to belief state s', requires that at least one belief be preserved from s to s'. If, further, an act is rational only if the belief revisions which occur during that act are rational belief revisions, then we have the makings of a defense of (R).

This defense of (R) is based on two premises: (1) if S's act is rational then every belief revision of S's which occurs during that act is rational; (2) if a revision from belief state s to belief state s' is rational, then at least one belief is preserved from s to s'.

This defense of (R) faces the following counterexample. While performing a rational act, a person makes a series of rational belief revisions, where each revision preserves at least one belief, but the first state in the series, s_i, and the last state, s_n, share no belief in common.

If the transition from s_i to s_n counts as a rational belief revision, then premise (2) is false. If the transition from s_i to s_n is a belief revision, but not a rational one, then premise (1) is false. If the transition from s_i to s_n is not a belief revision (although each intermediate step is) then the example conflicts with (R) because no belief lasts through the rational act.

Beyond this counterexample, the prospects for an argument for (R) based on (1) and (2) are further diminished by noting that premise (1) appears false. Surely a person may make irrational belief revisions that are unrelated to a rational act they are performing. He may rationally walk to the supermarket, while, at the same time,

80

thinking bizarre thoughts about poetry. A person can apparently change belief state irrationally while acting rationally. Thus, it looks hopeless to defend (R) on the basis of (1) and (2).

But perhaps we do not need (R) after all. Premise (2) could ground an anti-skeptical strategy. (2) alone suffices to conclude that no one with beliefs can, by rational belief revisions, reach a state without beliefs. So (2) alone may help with extreme skepticism. And, perhaps, something could be added to help with other forms of skepticism. Perhaps not all beliefs about the external world can rationally be given up because beliefs about the external world are required for rational belief revision to occur. (This is reminiscent of Strawson's suggestion that some commitments underly our "rational competence".) If so, then we don't need (R+) either. For, in that case, no one could rationally give up all of his beliefs about the external world.

I am doubtful of the suggestion that external world beliefs are required for rational belief revision. I also doubt whether a non-question begging defense of premise (2) can be given. Still, perhaps an anti-skeptical strategy could be mustered, based solely on considerations about belief revision. But I will not attempt that here. I turn now to a different argument for (R).

3.6.2 A better argument

(R) says that a lasting belief is required by every rational act. So, to defend (R), an investigation of rational acts is in order.

Rational acts

Performing a rational act, as I conceive it, requires the following two steps.

1. Adopting an appropriate plan. 2. Carrying out that plan.

These two steps are, I claim, necessary, but perhaps not sufficient, to perform a rational act.

Let me fill out this picture a little more.

- A plan specifies how an act will be performed. When adopted, the plan may be incompletely specified, details filled in as the act progresses.

- If a person adopts an appropriate plan, then he adopts a plan in accord with his values, preferences and beliefs.

- When a person acts rationally, the plan plays a guiding role. Think of this role in one of two ways: the plan itself helps to guide the act; or the person acts by following that plan.

- A rational actor follows a plan, but is not slavishly bound to her plan. Sometimes, a rational act is terminated before the plan is carried out to completion.

This is the outline of a framework. Still, enough is on the table to see that the framework is simple and interesting. It is interesting for at least three reasons.

1. The framework encompasses many of the acts in our daily life. We do choose and follow plans for acts: distant and sketchy plans (get a PhD) or immediate and more detailed (cook a tiramisu). And we see this way of acting as reasonable.

2. Adopting an appropriate plan, and following that plan are only claimed to be necessary for an act to be rational. This is compatible with varying ways of spelling out necessary and sufficient condition for an act to be rational.

3. This conception of a rational act is consonant with contemporary work in the theory of action.[11,12]

An argument for (R)

One way to defend (R) would be to fully specify this conception of a rational act, and incorporate it within a plausible theory of action. In the end, I think that is what

---

[11]Action theorists (who agree on little) widely agree that *something* (a 'plan' or an 'intention') plays a role in guiding action. Here are some examples. Myles Brand claims: "the cognitive component of an immediate intention is the guidance and monitoring of ongoing activity."[13, p 173] According to Alfred Mele, "An intention to A incorporates a plan for A-ing ... and one who successfully executes this intention is guided by the plan."[48, p 144] Michael Bratman argues that human agents are planning agents: "we settle in advance on [partial, future-directed] plans of action, fill them in, adjust them, and follow through with them as time goes by."[15, p 1][14]. J. David Velleman speaks of the "direction of guidance" of an intention: an intention causes what it represents.[67, p 25] (Note that these philosophers are discussing intentional action, rather than rational action. However, I believe they would accept that all rational actions are intentional actions.)

[12]Dion Scott-Kakures[59] relies on the claim that intentional actions are guided and monitored in arguing that believing at will is impossible. According to Scott-Kakures, "the reason that I cannot succeed in directly willing to believe that p is that the process which results in the generation of the belief would have to be unmonitored or ungoverned by the content of the intention or the plan... No one can will a belief that p because ... nothing could count as initiating a guided and monitored process which succeeded in producing a belief."[59, p 92] Although Scott-Kakures's argument fails (as ably shown by Dana Radcliffe[55]), I found his attempt instructive. (Thanks to R. Feldman for pointing me toward Scott-Kakures's essay.)

would be necessary to defend (R) in a satisfying way. I will not be able to accomplish that in the confines of an essay. Still, I will outline an argument in favor of (R).

I have just urged that rational acts are planned acts:

Premise 1. *A rational act is guided by an appropriate plan.*

Note, that for present purposes we may remain vague about what it is to be guided by a plan, as long as, in order to play its guiding role, the plan has been adopted when the act begins, and the plan continues to be accepted throughout the performance of the act. I believe that this much is uncontroversial, although some philosophers may prefer to speak of intentions rather than plans.

What needs to be spelled out further, and what may bring some controversy, is what is required to be guided by an *appropriate* plan: what is required to adopt an appropriate plan, and once adopted, what is required to continue to accept an appropriate plan. Here is where I propose to bring in belief.

Premise 2. *When a person adopts an appropriate plan, he relies on some beliefs.*

This premise says that a person relies on beliefs in adopting an appropriate plan. Call these beliefs the *adoption beliefs*. One more premise is needed to reach (R).

Premise 3. *While continuing to accept an appropriate plan, a person maintains at least one of the adoption beliefs of that plan.*

(R) follows from Premise 1, 2 and 3. Premise 1 ensures that a plan is accepted throughout the performance of a rational act. Premise 2 guarantees that there be some adoption beliefs, and Premise 3 requires that at least one of those adoption beliefs last while the plan is accepted. It follows that some belief will last throughout each rational act.[13]

Premise 1 is plausible, if the conception I have given of a rational act is plausible. Premise 2 and 3, however, need support.

---

[13]Some theories of action provide a shorter route to (R). According to some theories of action, intending involves having a certain belief. According to Harman, S intends to A only if S believes that he will A.[31, pp 90-93] According to Velleman, an intention is a kind of belief.[66][67] Other philosophers think that an intention is simply a belief-desire pair. If having an intention involves having a certain belief, and a rational act requires a guiding intention, then a belief lasts through the performance of each rational act.

Premise 2 is plausible

Try to imagine a counterexample to Premise 2. Consider Emma, who adopts a plan P without relying on any beliefs, and then acts rationally by following plan P. She does not rely on a belief that plan P is better than other plans. She does not rely on a belief that following plan P is morally required. She does not rely on a belief that she will carry out plan P. She does not rely on a belief that carrying out plan P is likely to achieve her goals. She does not rely on a belief about the world, about the past or about herself. Not only that, she does not rely on any belief in formulating the content of plan P. If plan P is to eat a sandwich, she does not use any beliefs about sandwiches, or eating, or how sandwiches are eaten in order to formulate the plan.

When I try to imagine a counterexample to Premise 2, I fail. The best I can do is this: Emma chooses plan P because she values following plan P more highly than following any other plan. Even so, I cannot imagine that it would be rational for Emma to follow that plan without relying on a belief that following plan P is highly valuable.

A failure to imagine a counterexample is not conclusive evidence. But it does make Premise 2 look plausible.[14]

Premise 3 is plausible

When should we reconsider our plans?

Clearly, we need not reconsider our plans whenever our beliefs change. Otherwise we would be spending most of our time reconsidering plans, and little time following them. A plan, to some degree, should stick.

On the other hand, we should not be "plan worshippers".[15] Clearly, we should

---

[14]Within the framework of Bayesian decision theory, Premise 2 is not obviously correct. In this framework, a belief state of suspending judgment on all propositions is represented by a probability function that serves in expected utility calculations as well as any other belief state. So, in this framework, a person apparently needs no beliefs to choose a rational act (an act that maximizes expected utility). However, it is not clear to me that the probability model can adequately represent suspension of judgment. Still, a decision theoretic version of the present response to skepticism would be desirable; it is just not clear how to formulate it.

[15]Michael Bratman[15] coins the expression "plan worship" in raising very similar issues.

sometimes reconsider our plans when our beliefs change. If I come to believe that it is not only silly, but dangerous to keep climbing the mountain, I should think about giving that plan up.

When should we reconsider our plans? Here is a partial suggestion. Suppose a person relies on some beliefs in adopting a plan. Then she begins to follow the plan. If, before she completes her act, she gives up all of those beliefs, then she should reconsider her plan. In short, a plan becomes suspect if all beliefs relied on in adopting that plan are lost.

This suggestion is a minimal condition; maybe a plan should be reconsidered for other reasons too. But this minimal condition is enough to support Premise 3, if to reconsider a plan is to stop following that plan, and then decide whether or not to continue following that plan.

Thinking about Premise 3 raises difficult questions. Can a rational person continue to follow a plan while reconsidering it? When is it rationally required or permitted to reconsider a plan? When is it rationally permissible to adhere to a plan? Can a rational person make a binding commitment to a plan, come what may?

Answers to at least some of these questions are needed to fully defend Premise 3.

3.6.3 What about (R+)?

(R+) can be defended by modifying the above argument.

Principle (R+) *Ordinarily, a person's act is rational only if she has a belief about the external world when that act is initiated which lasts until that act is finished.*

Premise 1. *A rational act is guided by an appropriate plan.*

Premise 2A. *Ordinarily, when a person adopts an appropriate plan, he relies on some beliefs about the external world.*

This premise says that, ordinarily, a person relies on beliefs about the external world in adopting an appropriate plan. Call these beliefs the *external world adoption beliefs*. One more premise is needed to reach (R+).

Premise 3A. *While continuing to accept an appropriate plan (that has some external world adoption beliefs), a person maintains at least one of the external world adoption beliefs.*

85

Now I think that, like Premise 2, Premise 2A is appealing. It does seem that we rely on beliefs about the external world in adopting plans for many (or most) of our everyday acts. However, Premise 3A is not obviously true.

We might try to support Premise 3A with Premise 3. Unfortunately, Premise 3 is too weak. It follows from Premise 3 that, while continuing to accept an appropriate plan (that has some external world adoption beliefs), a person maintains at least one of the adoption beliefs of that plan. But, unless every adoption belief is an external world adoption belief, the belief maintained need not be an external world adoption belief. Some other kind of belief could be maintained.

Premise 3A follows from:

Premise 3B *While continuing to accept an appropriate plan, a person maintains every adoption belief of that plan.*

However, Premise 3B is too strong to be plausible. Why can't a person reasonably give up even one adoption belief?

Is there a defensible claim stronger than Premise 3, but weaker than Premise 3B? Perhaps, with better understanding of rational acts, we will find that some, but not all adoption beliefs are *central,* and that the central adoption beliefs are maintained. If, in addition, ordinarily, some central adoption beliefs are beliefs about the external world, then we will have an argument for (R+). As yet, I cannot defend this suggestion.

3.7 Objections to the anti-skeptical proposal

3.7.1 Objection 1

If Gromit acts rationally, then Twin-Gromit, Gromit's internal duplicate, acts rationally too. But, according to (R) and (R+), Gromit may act rationally while Twin-Gromit acts irrationally. For Gromit and Twin-Gromit are in different environments. They might receive contrary news reports tomorrow; Gromit's beliefs may last longer than Twin-Gromit's corresponding beliefs.

Suppose that Gromit has a belief that will last longer than Twin-Gromit's corresponding belief. Therefore, although Gromit and Twin-Gromit are now internal

86

duplicates, they will not be internal duplicates later. Let us also suppose that Gromit now initiates act A, while Twin-Gromit initiates a corresponding act A'. According to the objector, if A is rational, then A' is rational too. The objector is correct if the rationality of an act is fully determined by the internal state of an agent when they initiate the act.

Such a view faces the following counterexample. Aaron is a sailor about to tack. Aaron's present internal state determines that his tacking will be rational. Bart is a sailor about to tack. Bart's internal state is just like Aaron's except that Bart has a disposition to behave rather oddly whenever he sees a dolphin. But not only is it extremely unlikely that Bart will see a dolphin, Bart will not see a dolphin. How can this odd disposition be relevant to the rationality of Bart's tacking? If Aaron's tacking is rational, then Bart's tacking is rational too. Curt is a sailor about to tack. Curt's internal state is just like Bart's. So, on the objector's view, if Bart's tacking is rational, then Curt's tacking is rational too. The unlikely occurs: Curt sees a dolphin, tacks the wrong way, and causes the boat to capsize. Curt's behavior seems irrational.

The objector can perhaps defend his view. But, I think I have said enough to show that this objection should not be accepted without much more argument.

3.7.2 Objection 2

Whether or not an act is rational depends only on what is available to the person when he decides to so act. But, according to (R) and (R+), an act may be irrational because of something that is not available to the person when he decides to act.

As in Objection 1, this objector suggests that the rationality of an act is already determined just as an act is initiated. However, Objection 2 contains a further thought. Stated vaguely, the thought is that rational acts are seen by the agent, from his own point of view, as rational. The point of view of some outside observer is irrelevant. But (R) and (R+) do not fit this picture. What will happen in the future is outside an agent's point of view, and (R) and (R+) have it that what will happen in the future matters to the rationality of an act that a person is about the perform.

For, according to (R) and (R+), the length of time a belief will last matters to the rationality of an act that a person is about to perform.

In response: Surely it needs to be spelled out what it is for an act to be seen as rational from an agent's own point of view. But even without more detail, we can see that (R) and (R+) are not incompatible with that vague thought. The agent has a point of view at each moment during an act. (R) and (R+) say that rational acts require beliefs that last long enough. The loss of a belief is presumably available to the agent, from his own point of view. So, while performing an act, any failure to meet the conditions set by (R) and (R+) is presumably available to him, from his own point of view.

3.7.3   Objection 3

For all the argument shows, the lasting beliefs required for each rational act may be given up just as the act finishes. So, all of these beliefs may be lost at the moment a rational act finishes; at the moment a rational act finishes, a person can enter a skeptical state. Thus, while acting rationally, a person can enter a skeptical state.

The objector raises a difficult question about timing. Let us agree that if an act continues after its last adoption belief is lost, then, by the argument for (R), that act is irrational. Let us also agree that if an act finishes before its last adoption belief is lost, then, as far as the argument goes, that act may be rational. But what if an act finishes at the very moment its last adoption belief is lost? According to the objector, as far as the argument goes, that act may be rational, so a person may enter a skeptical state while acting rationally.

Consider the skeptical state of having given up all external world beliefs. Even if the objector is correct, to enter that skeptical state while acting rationally, a person must rationally finish *all* of her ongoing acts which require external world beliefs at the moment she loses her last adoption belief. Otherwise, she performs at least one irrational act. This situation is odd. Ordinarily we are performing many different acts which require external world beliefs. It would be odd to rationally finish all of these ongoing acts at just the same moment. As far as I can tell, it would cause no harm to allow that, in this odd case, a person, while acting rationally, may enter

88

that skeptical state. However, this objection reveals that the present response to skepticism holds more interest, if, in ordinary life, people indeed are often engaged in multiple acts requiring external world beliefs.

3.7.4   Objection 4

According to the anti-skeptical proposal, some skeptical argument is flawless. But if you believe that some skeptical argument is flawless, you should accept the conclusion; if you believe that some skeptical argument is flawless, you should become a skeptic. Thus, even if (R) and (R+) are correct, the proposal does not help with skepticism.

The anti-skeptical proposal is that, ordinarily, no one, while acting rationally, can be convinced by a skeptical argument. At the beginning of this essay, I said that I would propose an approach to skepticism distinct either from accepting skepticism (in some form or other), or facing the problem of satisfactorily locating defects in every skeptical argument.

Note, however, that the anti-skeptical proposal is compatible with the claim that no skeptical argument is flawless. I do not advocate such an approach to skepticism because I doubt whether it is any more powerful than simply accepting the problem of satisfactorily explaining what is wrong with every skeptical argument.

The anti-skeptical proposal is also compatible with the claim that there is a flawless skeptical argument. I think this approach to skepticism is more interesting. The objection is that this approach falls into skepticism.

This objection is mistaken. Believing that there is a flawless skeptical argument is compatible with failing to believe, of any particular skeptical argument, that it is flawless. But the latter is what is needed for the objection. The objection is apparently based on a principle like this: if you believe that argument A is flawless, then you should accept the conclusion of argument A. This principle does not apply when you merely believe that there is a flawless skeptical argument out there. (If every skeptical argument has the same conclusion, and you believe that there is a flawless skeptical argument out there, then, perhaps, you should accept the conclusion.

89

However, not every skeptical argument has the same conclusion.)[16,17]

3.8 Conclusions

To recap: I first argued that no one, while acting rationally, can be convinced by an extreme skeptical argument. I then argued that, ordinarily, no one, while acting rationally, can be convinced by other skeptical arguments. I have three speculative conclusions to offer.

1. Many philosophers think it reasonable to ignore skepticism. In a sense, they are right. If paying attention to a skeptical argument puts one at risk of acting irrationally, then the risk is avoided by staying away from skeptical arguments.

2. The results of this essay can give comfort to those who have felt the disquieting temptation of skeptical arguments. One may acknowledge the temptation to be convinced by accepting that some skeptical arguments are flawless. Yet, there is reason to resist the temptation to be convinced by a skeptical argument: avoiding an irrational act.

3. It is a weak objection to a theory that it "leads to skepticism". That an apparently valid skeptical argument can be constructed from premises taken from a theory is not sufficient to show that the theory is mistaken. For, if this essay is correct, a non-skeptic may accept that there is a flawless skeptical argument. So, a non-skeptic may accept that there is a flawless skeptical argument constructed from premises taken from a theory he accepts. However, a non-skeptic may not accept, of some particular skeptical argument, that it is flawless. For (as suggested in Objection 4) if a person believes of an argument that it is flawless, he should accept the conclusion.

---

[16]N. M. L. Nathan[49] is one non-skeptic who explicitly claims that some skeptical arguments are sound.

[17]Here, briefly, are two further objections. 1. Objection: The following can be rational: you decide to forget about the world by drinking a bottle of whiskey. When drunk, you lose all your beliefs about the external world. Then you get them back when you sober up. So, beliefs don't need to last through a rational act. Reply: Maybe so. But if so, (R) and (R+) can be weakened to require only that the belief be present at the beginning and the end of the act, without requiring that it persist through the act. Even weakened, (R) and (R+) are enough for an anti-skeptical strategy. 2. Objection: (R) is false because states other than believing (like hoping or accepting) are sufficient for rational action. Reply: Hopes by themselves are not sufficient; beliefs, arguably, are also needed. As for accepting, I have not yet seen a theory that successfully distinguishes accepting from believing.

So, if a non-skeptic comes face to face with an apparently flawless skeptical argument constructed from his theory, he must, on pain of irrationality in action or belief, reject his theory, or find some flaw in the argument. If the non-skeptic stays away from skeptical arguments, he can avoid the need to look for flaws in a skeptical argument constructed from his theory.

3.9 More questions than answers

This chapter has many inadequacies, raising many questions. Without satisfactory answers, I can at least pose questions for future work.

3.9.1 Action and belief

The anti-skeptical proposal is sensitive to details about the way belief is connected to action.

When, exactly, can a belief that supports and guides a rational act be lost?

This is a question about timing. The proposal requires a belief's presence both at the beginning and at the conclusion of the act. If the belief can be lost *just before* the act is completed, then the proposal fails. If not, but the belief can be lost *just as* the act is completed, then the proposal is weakened, as noted in the reply to Objection 3. If the belief can only be lost *after* the act is completed then the proposal may succeed. When, exactly, can a belief that supports and guides a rational act be lost? The tools I use in this chapter are too blunt to tell us.

There are other questions about the way belief is connected to rational action.

What kinds of beliefs are needed in acting rationally? Beliefs about the external world? Beliefs about the future? Beliefs about yourself?

Can beliefs at first needed in acting rationally be replaced during the act? Can every belief needed when the act is initiated be replaced by the time the act is completed? Can every external world belief needed when the act is initiated be replaced by the time the act is completed?

When should a rational act be reconsidered? When should a rational act be continued even though your belief state changes?

91

Again, the tools of this chapter are too blunt to answer these questions. But answers to these questions matter for the success of the anti-skeptical proposal.

Is belief indeed needed to support and guide rational action? Or can some non-belief state— like acceptance— do the job?

Although in this chapter I claim that belief is needed, I don't think this claim is adequately supported. Note, however, that if some non-belief state can do the job, the proposal in this essay may extend. If it is possible to epistemically evaluate the non-belief state (such that there are skeptical arguments that it is not rational to be in that state) then there is a version of the proposal relying on a claim that one of those non-beliefs states must last through a rational act. But if it is not possible to epistemically evaluate the non-belief state, then arguably, for that very reason, rational action is not supported by states of that kind alone. Making progress involves: (1) stating precisely what the non-belief state is supposed to be; (2) examining how that state may be subjected to skeptical reasoning; (3) examining how and whether that state can support action.

The above questions, which explore the connection between belief and action, are interesting apart from any proposal about skepticism. I believe that tackling such questions would be a fruitful way to approach broader questions like: What is a belief? What is an act? When is an act rational?

3.9.2   A formal invitation

Perhaps some of the foregoing questions about the connection between belief and rational action can be settled by adopting a formal framework. The one framework I know of connecting belief to rational action appears insufficient, on its own, to offer answers. According to decision theory, the rational act to choose is the one that maximizes expected utility. One question about this framework, which I mentioned earlier, is whether it adequately represents suspension of judgment. More relevant to the questions I have just asked, an act, in this theory, is a single unit to be chosen or not. As I have been discussing, seeing acts as they take place over time gives rise to questions like: When should an act be reconsidered? When should an act be

continued even though your belief state changes? I suppose the decision theorist would say that you should reconsider when reconsidering has a higher expected utility then not reconsidering, and you should continue when continuing has a higher expected utility then any other possibility. But there are cases which render this suggestion puzzling. This suggestion can leave you ineffectually oscillating between two acts (you rationally choose act A; while performing act A your belief state unexpectedly changes so that it is better to reconsider and switch to act B; while performing act B your belief state unexpectedly changes so that it is better to reconsider and switch to act A.) Surely there should be no puzzle here: you just make a relatively firm commitment to the acts you choose. Perhaps the decision theoretic framework can do better, if a notion of commitment can be incorporated.

3.9.3   Is this a response to skepticism?

Suppose, for the moment, that my claims about rational action are correct, in particular that (R) and (R+) are correct. Even so,

Can claims about action ground an adequate response to skepticism?

Think about the poor philosopher who gets himself tangled up in skeptical doubts. He started reading Sextus, and now finds himself almost— but not quite— convinced that he has no better reason to believe any proposition about the external world than its negation. But he can see nothing wrong with the argument leading to this conclusion. What is the poor philosopher to think?

Even in its strongest form, the proposal I am offering looks rather unhelpful to the poor philosopher: He is told that he will act irrationally, if he becomes convinced, and remains in a rational belief state. Imagine his reply: "I'm doomed to irrationality! I can't just run away. I can't leave this skeptical argument unanswered. I have an intellectual obligation to seek out truth and avoid falsehood. I either have to find a flaw in the argument or accept its conclusion. But I can see no flaw, so I must accept its conclusion, thereby (if you are right) acting irrationally."

Surely we could help the poor philosopher if we could give him an argument that his beliefs are justified. But even the usual project of finding flaws in skeptical

93

arguments would be helpful, if we could find a flaw in the argument he is worried about. Maybe even a Moorean thought could help: that he is more confident of his beliefs about the external world than the premises of some nasty old philosophical argument. How can the proposal I offer help?

Think from the third person point of view about the poor philosopher. If the proposal I offer is correct, then, convinced by the argument, he will either act irrationally or end up in an irrational belief state. Does that mean he is doomed to irrationality? Not if he has other options. And he very likely does. Perhaps he might have dinner, watch a reality show on television, and then go to bed. The next morning, he might get up and start working on philosophy of language. He can very likely do all this acting rationally and remaining in a rational belief state. If not, there is probably something else he can do, acting rationally, and remaining in a rational belief state. Being rational, the poor philosopher should act rationally. So he should choose an option like this. And he can take the third person point of view and think similar thoughts about himself. He can think to himself "I do have another option. I can do something else and remain rational. So I am not doomed to irrationality after all."

But what about the poor philosopher's intellectual obligation to seek the truth? As a philosopher interested in epistemology, he wants to find out whether his external world beliefs are justified. Perhaps a non-philosopher could just as well watch television. But shouldn't the philosopher pay close attention to a serious argument that could help him find what he is seeking? He shouldn't just walk away, should he?

He should. According to the proposal, the rational thing to do is to walk away.

The point here (and this should be developed further than I can at present) is that, in a particular way, what it is rational to do takes precedence over what it is rational to believe. The background thought is that activities of truth seeking and falsehood avoiding have their place, and have their point, within a human life. This is very vague. But, to be clear, I am not saying that you should come to believe something for practical reasons, despite the evidence. I am not saying that you should come to believe something just because it is useful. I am not saying that you should, despite strong contrary evidence, come to believe that you can jump the crevasse because

this will give you a better chance to make it across.[36] Instead, I am saying that, sometimes, when believing something, practical reasons demand that you not give up that belief; sometimes practical reasons demand that you avoid or ignore counter evidence. When believing that your beliefs about the external world are justified, practical reasons may demand that you not give up that belief; practical reasons may demand that you avoid or ignore evidence to the contrary.

This may sound like a plea for intransigence. It is, but only in a limited way. A vexing question, which I cannot answer, is

When is such "intransigence" rational and when is it not?

For, on the one hand, surely we should not change our belief state in response to every bit of information that appears. We have good practical reasons not to. We wouldn't be doing much else; we are constantly bombarded by information.

On the other hand, surely we should not be as intransigent as Woodrow Wilson, of whom an admiring(!) friend said: "Whenever a question is presented he keeps an absolutely open mind and welcomes all suggestions and advice which will lead to a correct decision. But he is receptive only during the period that he is weighing the question and preparing to make his decision. Once the decision is made it is final and there is an absolute end to all advice and suggestion. There is no moving him after that."[45, p 5]

One final question:

How is the present proposal related to the anti-skepticism of Hume and Strawson?

At the beginning of this chapter, I said that the Humean response to skepticism is based on a claim that there are some beliefs which we, by nature, cannot give up. The problem is that we want to know why it is reasonable to hold on to these beliefs. Just because we cannot give up a belief looks like a weak reason to set aside an argument to the contrary. Strawson's proposal is similar, although he brings in Wittgensteinian complexity: skepticism is idle because some commitments can neither be given up nor intelligibly questioned.

In a nutshell, the present proposal is that holding on to some beliefs is reasonable

not for theoretical reasons, but practical reasons. Although I depart from the suggestion that we cannot give up some beliefs, I think the idea is similar in spirit: we cannot give up some beliefs *while acting rationally*. Moreover, like Hume and Strawson, I do not try answer skeptical arguments by looking for flaws, or directly arguing that the skeptical conclusion is false. Instead, I have tried to show that skepticism can be rationally avoided.

Chapter 4

The value of knowledge

Abstract

I argue that knowledge is no more valuable than stable true belief. This surprising claim severely constrains acceptable accounts of knowledge. It also helps accommodate skepticism about knowledge.

4.1 The value of knowledge

Why is knowledge valuable? This question is too rarely asked. I will try to show that examining the value of knowledge sheds light on two central epistemological questions: What is knowledge? Is it possible to attain knowledge?

The main claim of this chapter is that **knowledge is no more valuable than stable true belief**. This claim is surprising. Doesn't knowledge have a unique and special value? If the main claim is correct and if, as it seems, knowledge is not stable true belief, then knowledge does not have a unique value: stable true belief is just as valuable.

I draw three conclusions. First, the main claim does not devalue knowledge or undermine our knowledge gathering practices. Second, skepticism about knowledge is harmless. Even if one cannot have knowledge, one can have something just as valuable. Third, any attempt to analyze the concept of knowledge faces a severe constraint.

But before drawing any conclusions from the main claim, I must first explain and defend it. I turn to that task next.

4.1.1 The main claim motivated and clarified

The main claim motivated

In the *Meno*, Plato raises a question about the value of knowledge.[53, 97A-98A] If, wonders Plato, a man knows which road leads to Larissa, he can surely get there. And if a man does not know, but has a true belief about which road leads to Larissa, he can get there all the same. Why then is it better for him to *know* which road leads to Larissa? Why is knowledge more valuable than mere true belief? Plato's answer is that knowledge is more valuable than mere true belief because mere true belief can be easily lost; knowledge, however, is "fastened" by an explanation and thus more difficult to lose.

The Platonic thought is that knowledge is more valuable than mere true belief because knowledge is more *stable.* Here is an example to make this thought clearer:

> Restaurant reviewers Joe and Renata are dining in a fancy new trattoria. Joe knows that the unusual mushroom on Renata's plate is poisonous. Renata does not know that the unusual mushroom on her plate is poisonous, although she has a true belief that the unusual mushroom is poisonous. Renata concluded this from her false belief that the chef, bitter at her negative reviews, is trying to poison her. The chef, in fact, holds no grudge; he has simply made a mistake. The chef comes out and clearly does not recognize Renata. So she gives up her belief that the mushroom is poisonous, and starts to eat it. Renata is in danger because her true belief was easily dislodged by new evidence. Luckily for Renata, Joe's knowledge is not so fragile; Joe, who knows, is able to stop her in time.[1]

---

[1]T. Williamson gives a similar example.[71, 87]

Examples like the case of Joe and Renata help make Plato's two theses look very plausible:

(1) Knowledge is more valuable than mere true belief.

(2) Knowledge is more valuable than mere true belief because knowledge is more stable.

Indeed, (2) suggests a bolder thesis:

(3) Knowledge is more valuable than mere true belief *only* because knowledge is more stable.

The bolder (3) is a short step from the main claim of this essay. For if (3) is correct, then, plausibly, a true belief which is as stable as knowledge is no less valuable than knowledge. That is my main claim.

So the main claim is motivated by the Platonic thought. But none of these theses and claims are particularly clear. What is stability? What is *mere* true belief? Why does stability help make knowledge (or true belief) valuable? What, exactly, is supposed to be valuable? Is it information, the state of knowing, the concept of knowledge, or what?[2]

The main claim clarified

Here are two steps to clarify the main claim. First, compare the value of states, as opposed to the value of concepts or the value of information. Second, understand stability as the length of time that a state will last. One state is more stable than another just in case the former state will last longer than the latter.

The main claim is thus: **the state of knowing that p is no more valuable than the state of having a true belief that p which will last at least as long.**

---

[2]Somewhat surprisingly, in recent literature, thesis (1) is presented as an obviously true premise to be used in philosophical argument. Zagzebski[73] argues that reliabilist accounts of epistemic justification are incompatible with (1), and, since (1) is obviously true, reliabilism should be rejected. Jones[37] argues, more generally, that any account of epistemic justification which sees epistemic justification as valuable only as a means to gaining true beliefs and avoiding false beliefs, is incompatible with (1), and thus objectionable. Riggs[56] replies that Jones's and Zabzebski's arguments only support the weak conclusion that, in order to preserve (1), certain accounts of justification need to be supplemented. Another sort of argument is given by DePaul[21] and Riggs[57]. They conclude that gaining true beliefs and avoiding falsehoods cannot be the only epistemic values, otherwise (1) would be false.

Since this is rather longwinded, in the interests of brevity, I will often abbreviate this by saying that knowledge is no more valuable than stable true belief.

I could have chosen other clarificatory steps. For example, instead of discussing the value of the state of knowing, I might have discussed the value of the concept of knowledge. The value of the concept of knowledge is not my topic. Indeed, the value of the concept of knowledge should be clearly distinguished from the value of the state of knowing. The concept of knowledge may be valuable even if no one knows anything, just as the concept of a perfect vacuum may be valuable even if no perfect vacuums exist. Conversely, the state of knowing where to find his favorite toy may be valuable to a child, even if he has no need for the concept of knowledge. I will focus on the value of the state of knowing, leaving other questions about value for another occasion.

The second clarificatory step may seem odd. I say that a state is more stable just in case it will last longer. But the word "stable" is usually understood to mean something like "unlikely to be changed" or "not easily changed"; stability is more naturally understood as a modal property, like fragility. Here I must ask the reader's patience; it will soon become clear why I have chosen an unusual notion of stability.

Some readers may continue to feel puzzled by the main claim that the state of knowing that p is no more valuable than the state of having a true belief that p which will last at least as long. In what respect are we to compare these two states? Their value for action? Their intrinsic value? Or what? The answer: we are to compare these two states in any respect in which knowing is valuable. In gory detail, then, the main claim is: **in any respect in which the state of knowing that p is valuable, the state of having a true belief that p which will last at least as long, is just as valuable.**

4.2 What is the value of knowledge?

I will now defend the main claim by exhaustively examining proposals to explain why knowledge is valuable. For each proposal, I will argue that if knowledge has that

value, then stable true belief does too.

4.2.1 Knowledge is useful

It is sometimes suggested that knowledge is valuable because knowledge is useful. Proposals along these lines include: knowledge enhances your chance of survival; knowledge helps you make good decisions; knowledge helps you satisfy your desires; knowledge helps you to act morally; knowledge gives you power over others.

These proposals are easily dealt with. In each case, it is easy to see that a stable true belief is just as valuable.

According to these proposals, knowledge is useful for action— useful either in deciding which action to choose or in carrying out an action. Plato has already gone down this road. In trying to explain why knowledge is more valuable than mere true belief, Plato considers why knowledge is more useful to a person performing the action of traveling to Larissa.

Unlike Plato, I am comparing the value of knowledge to the value of stable true belief. Is knowledge more useful for action than stable true belief? Consider a case where knowledge enhances your chances of survival. If I know there is an angry, hungry tiger behind the door, then, as long as I am trying to stay alive, I won't open it. My knowledge keeps me from life-threating danger. If I have a stable true belief that there is an angry, hungry tiger behind the door, then, as long as I am trying to stay alive, I won't open it either. Knowledge seems to have no advantage here.

The example suggests a general point. When deciding what to do, it is useful to have an accurate picture of the world. And when acting, and it is useful to *maintain* an accurate picture of the world. But my picture of the world, with respect to the proposition p, is no more accurate when I know that p, than when I truly believe that p. And when maintaining an accurate picture of the world, knowing that p, and stably, truly believing that p, are on a par.

I hear grumbles about this line of thought. One objector claims that false beliefs are sometimes more useful than knowledge. A rough and ready picture of the world can be better than an accurate one. For one thing, gaining true beliefs or knowledge

sometimes carries too high a cost. In response, although false beliefs may sometimes be more useful than knowledge, I am comparing the value of knowing that p to the value of having a stable true belief that p. Whether or not it is sometimes better to have a false belief is beside the point.

A second objector notices that, in some situations, knowing that p is clearly more valuable than having a mere stable true belief that p. Consider an evil demon who will hinder my actions as long as I fail to know that p. Then, I'd be better off knowing that p than having a mere stable true belief that p. Such cases are to be expected. The presence of an instrumental value can require that certain background conditions hold. An example: a refrigerator, normally useful to keep the milk from going sour, doesn't help much during a long blackout. Although a stable true belief may normally be useful for action, in this evil demon case, it is not. Although the second objector's point is correct, the claim here should anyway be limited to normal conditions: Knowledge is, in normal conditions, no more useful for action than stable true belief. That said, for brevity, I'll now drop the reference to normal conditions.

I conclude that if the value of knowledge lies only in its usefulness for action, then knowledge is no more valuable than stable true belief. But there are other proposals for the value of knowledge.

4.2.2 Valuable for others

Your knowledge is not only valuable for you, it is valuable for others as well. You can inform others. You can give them a stock tip, warn them of imminent danger, or just satisfy their intellectual curiosity. Could your knowing that p be more valuable for others than your having a true belief that p? This looks like a non-starter. When you know that p you can inform your friend that p. Just as easily, when you have a stable true belief that p you can inform your friend that p.

There is more to say. As E. J. Craig[20] notices, being a good informant involves more than having a true belief. A good informant is recognizable as a good informant; someone trying to find out whether p needs to be able to pick a good informant out of a crowd. According to Craig, a good informant not only has a true belief that p,

she has a detectable property X which correlates well with being right about p.

Craig's broader project is curious. He suggests that thinking about our practical need to find good informants helps us to explain why our concept of knowledge is the way it is. Thus, he tries to understand the concept of knowledge by means of the concept of being a good informant. But Craig does not make the straightforward claim that to know that p is to be a good informant whether p. Instead, according to him, the concept of knowledge is a stretched and twisted version of the concept of being a good informant.

I simplify Craig's wide-ranging discussion, and extract the following thought: knowing that p is valuable to others because someone who knows that p often has a detectable property which correlates well with having a true belief that p. So let us look to see if knowing does have such a detectable property, and, if it does, whether that makes knowledge more valuable than stable true belief.

Here is one possibility: A person who knows that p can often give reasons in support of his view about p. If so— if a person who knows that p can often defend his view when asked— then knowing looks to have an advantage over mere stable true believing. When you are looking for a good informant, you can ask your target what reasons she has. If she knows that p, she can often give you some. So you can often get some evidence that she knows that p.

Stable true belief differs from knowledge in this respect. If a person has a stable true belief that p, she just has that belief. She may or may not have reasons. If knowing that p often comes with reasons, then knowing that p is, apparently, more valuable to others than having a stable true belief that p. This conflicts with the main claim.

I see two satisfactory ways to respond. The first way is to retreat to an extended version of the main claim. Rather than comparing knowing that p to having a stable true belief that p, compare knowing that p to having a stable group of true beliefs. Include in the stable group of true beliefs not only a true belief that p, but some reasons in favor of p that the person can give. A person who has this stable group of beliefs can give evidence that she has a true belief that p. Although a retreat

from the main claim, this response does not depart from its spirit— that there is nothing distinctively valuable about knowledge not found in stable true belief. If indeed knowing that p is often detectable because knowing that p often comes with reasons, then this group of stable true beliefs is equally detectable.

A second way to respond maintains the main claim in the face of the objection. Consider, for the moment, a person who is rational, and who will continue to be rational, in the following sense: She will respond to evidence by modifying her belief state appropriately. Thus, she will not, for example, hold on to some belief out of stubbornness despite the evidence. Now a person rational in this sense at least arguably often has reasons for her beliefs. For, without reasons for many of her beliefs, how can a person always respond appropriately to evidence? But if a rational person often has reasons for her beliefs, then a rational person with a stable true belief that p often has reasons in favor of her belief that p. So a rational person with a stable true belief is often able to defend her view. Thus, if we restrict the discussion to rational people, there is no advantage here for knowing that p.

And there is reason to restrict the discussion to rational people. Consider one example. Poor Julia knows that her car is parked on floor 3 row R. Unfortunately, she changes her belief state willy-nilly whenever she sees a red car. On the way to her car, she suddenly sees a red car, and she gives up her earlier belief, and comes to believe that her car is parked on floor 2 row B. A moment later, she comes to believe that her car is parked on floor 12 row D. It is difficult to make sense of the thought that Julia's knowing where her car was had any value for her. That state had no reasonable role in her cognitive life. Julia does not respond appropriately to evidence, and it is difficult to understand why knowing has value for her. This is just one case, but it supports restricting the discussion to rational people. More should be said, but I think that this is enough to show that there is at least some plausibility in responding to the objection by restricting the discussion to rational people.

I prefer the second of these responses, because it holds on to the main claim. But either response, I believe, satisfactorily dispenses with the objection that knowing that p is more valuable because a person who knows that p often has the detectable

property of having reasons to believe p. Even so, knowing that p may sometimes be accompanied by some other detectable property.

I can think of only one other plausible candidate detectable property.[3] When you know that p, you often know related things. When searching for a good informant as to p, I can ask you about related topics of which I am already informed. If you know about these topics, I have evidence that you know whether p. Arguably, then, knowing that p is more valuable than a stable true belief that p because a stable true belief that p may stand alone, unaccompanied by true beliefs about related topics. Again, this conclusion conflicts with the main claim.

As before, one may retreat to an extended version of the main claim. Talk not of a stable true belief that p, but a group of stable true beliefs, including beliefs about related topics. Knowing that p, even if accompanied by knowledge of related topics, is not more valuable that this group of stable true beliefs.

Although this reply is adequate, we can do better, and preserve the main claim. There is no reason to think that a person with a stable true belief that p is less likely be exposed to information about related topics than a person who knows that p. A person with a stable true belief surely has the same opportunity to learn about related topics. So there is no reason to think that a person with a stable true belief is any less informed about related topics. So knowing that p is not more valuable even if often accompanied by knowledge of related topics.

I can see no further plausible reason to think that knowledge is more valuable to others than stable true belief. Let's move on.

4.2.3 Knowledge as an achievement

One somewhat vague suggestion is that knowledge is valuable because knowledge is an achievement. Riggs[56][57] tries to cache out this vague suggestion. Riggs' thought is that knowing that p is valuable (in part) because a person who knows that p deserves

---

[3]Well-known analyses of knowledge fail to provide a plausible candidate. A true belief reliably formed[28], or the output of properly functioning cognitive equipment[52], or tracking the truth in nearby possible worlds[50], is no more detectable than a generic stable true belief.

some epistemic credit for having reached that state:

> Being in the state of "knowing that p" entails of a person that she have a true belief for which she deserves a certain degree of epistemic credit. Believing something true by accident entails no credit of any sort to the person. This is so despite the fact that the belief is no more valuable in the former case then the latter, nor need we assume that the believers in question differ in their respective epistemic qualities. The difference that makes a *value* difference here is the variation in the degree to which a person's abilities, powers, and skills are causally responsible for the outcome, believing truly that p.[57, 94]

Riggs seems to think that he has isolated a value distinctive of knowing, as opposed to mere true believing. A Gettier-style example shows that this is not correct; mere true believers can deserve just as much epistemic credit as knowers. Suppose that Ned comes to know that someone at the lab owns a Toyota, having often seen Jones drive around in Jones' own Toyota. Jed, who works in the same lab, comes to truly believe that someone at the lab owns a Toyota, having often seen Smith drive around in Jones' Toyota. Since Smith owns no car, although Jed's belief is true, Jed does not know. Ned and Jed, we can assume, are causally responsible to just the same degree for their true beliefs. So, by Riggs' lights, they deserve the same epistemic credit for their respective true beliefs.

Perhaps Riggs' point is instead that knowers always deserve some epistemic credit, but mere true believers only sometimes do. This may be correct. But that does not mean that knowers always deserve at least as much epistemic credit as mere true believers do. In fact, knowers sometimes deserve *less* epistemic credit than mere true believers: Joe the policeman happens to walk by at the moment Lefty the gardener pets the dog. Joe comes to know that one of the gardeners petted the dog. Smitty the detective, after an exhaustive and careful investigation, comes to falsely believe that Righty the gardener petted the dog. Smitty concludes that one of the gardeners petted the dog. Surely Smitty is more causally responsible for his true belief than Joe. So, apparently, by Riggs' lights, Smitty deserves more epistemic credit. That means Joe, who knows, deserves less epistemic credit than Smitty, who has a mere true belief; someone who knows can deserve less epistemic credit than someone with a

mere true belief. Thus Riggs' proposal does not support the conclusion that knowing is more valuable than stable true belief.

Step back from Riggs proposal, and return to the thought that knowledge is valuable because it is an achievement. Sometimes knowledge is an admirable achievement. But not always; sometimes knowledge comes easy. And sometimes stable true belief is an admirable achievement. But not always; sometimes stable true belief comes easy. Either state can be an achievement or not. So we have not yet found something distinctive of knowing which makes knowing that p more valuable than having a stable true belief that p.

4.2.4   A close connection to the world

Another thought is that knowing involves a close connection to world; knowing gives you a firm grip on the way things are.

More specificity is needed to evaluate this thought. Robert Nozick's tracking theory of knowledge is one way to make this close connection idea more precise.[50] Bells and whistles aside, Nozick claims that to know that p is to have a true belief that p which tracks the truth, where your belief that p tracks the truth just in case: if p were false you wouldn't believe that p, and if p were true you would believe that p. Possible worlds talk helps make the notion of truth-tracking clearer. Your belief tracks the truth just in case in nearby possible worlds where p is true, you believe it, and in nearby possible worlds where p is false, you don't. A person with a truth-tracking belief has a grip on the way things are— small changes in the way things are would not disturb his grip on the truth.

Even if (as compelling examples suggest) Nozick's conditions are neither necessary nor jointly sufficient for knowing that p, his proposal has an appeal which suggests that he may be on the right track. So it is worth asking whether a truth-tracking belief is more valuable than a stable true belief.

Initially, it seems implausible to think that a truth-tracking belief is more valuable than a stable true belief. Why would you would care whether or not your beliefs track the truth in some nearby possible world? You are not in some nearby possible world;

you are in the actual world. If in some nearby world you get hit by a car, come down with the measles or get a terrible toothache, too bad for you in that possible world. Luckily, you are in this world, where these things didn't occur. Similarly, in some nearby world you may have some extra false beliefs, or fail to have some true ones. So what. You are in the actual world where this didn't occur.

Yet, there is reason to care whether your beliefs track the truth in nearby worlds: you don't know which world you are in. Moreover, since some of your beliefs are almost certainly false, you are almost certainly in a world that you consider to be merely possible. Thus, since you care whether or not your beliefs are true in the actual world, you should care whether or not your beliefs are true in worlds that you consider to be merely possible.

This reasoning does not show that truth-tracking beliefs are more valuable than mere stable true beliefs. According to this reasoning you should care whether or not your beliefs are true in worlds that you consider to be merely possible because you care whether or not your beliefs are true in the actual world. But if your belief is stable and true, then it is stable and true in the actual world; if your belief is stable and true then you have what you want already. Having true beliefs (and avoiding false beliefs) in some other world does not satisfy any further want.

(A similar line of thought explains why the notion of stability in the main claim is persistance in the actual world. Ask a parallel question: why should you care whether a true belief is resistant— whether a belief would be held truly in some other possible situation? The answer that comes to mind is that since you don't know what situation you are in, you want to hedge your bets. So you ensure that the belief is truly held in a range of situations. But if you have a persisting true belief, you have what you want already– that the belief continue to be held and be true in the actual world. Having a resistant true belief gives you nothing further that you want.)

Let's take stock. I turned to Nozick's tracking account of knowledge in developing the idea that knowing involves a close connection to the world. Then I tried and failed to find a reason to think that a truth-tracking belief is more valuable than a stable true belief. But perhaps truth-tracking was not the right place to start. Let's

try again.

One thought is that stable true belief can be lucky in a way that knowing cannot, and in a way that weakens your connection to the world. Big Al thinks the government owes him $200. Big Al is right, but only by luck. His accountant made several errors on his tax return which cancelled each other out. Luckily for Al, his belief is stable; he won't notice any of the errors. Al's grip on fact that the government owes him $200 seems tenuous. Al's connection to the world seems much weaker than that of Slim, who comes to know that the government owes him $200 because his accountant, who made no mistake, tells him so.

There is a good reason to disvalue luck: as any gambler knows (or should know), you can't count on luck. Luck comes and goes. But why think that luck makes Al's state more tenuous (and thus less valuable) than Slim's? One reason is that Al is in a state that could easily be lost if his luck turns bad. However, we should not disvalue Al's state because his good luck might run out. For his luck won't run out; his state won't be lost— it is stable.

Another reason to think that luck makes Al's state less valuable than Slim's is that Al was very lucky to have a true belief. He might easily have had a false one. His accountant could have easily made one less error. Again, however, we should not disvalue Al's state because he might easily have had a false belief. Knowing can depend on luck in just the same way: I should have mentioned earlier that Slim usually uses Al's mediocre accountant, but this year he was lucky that his cousin the excellent accountant was visiting at tax time.

This is enough, I think, to see that this is another false lead. Knowing can occur by luck, be maintained by luck, or not. Stable true believing can occur by luck, or be maintained by luck, or not. So stable true belief is not dependent on luck in a way that makes it less valuable than knowledge.

Here is one last reason to think that knowing brings a closer connection to the world than stable true belief. When you know that p you have a thorough understanding which may be lacking when you have a stable true belief that p. For example, Andrew Wiles knows that Fermat's Last Theorem is true because he proved it. He

understands how FLT connects to other theorems; he understands why it is true. Silly Billy, on the other hand, believes that Fermat's Last Theorem is true on the basis of a mistaken proof he concocted in the eighth grade. Luckily for Billy, his belief is stable. Unlike Wiles, however, Billy has no deep understanding.

The contrast between Billy and Wiles neglects an important point: Surely, you can know that Fermat's Last Theorem is true without understanding why it is true. You can know because the respected mathematician tells you so. Knowing that p does not always bring a thorough understanding. But if knowing that p does not always bring a thorough understanding, then knowing has no advantage over stable true believing.

Maybe this answer is too quick. Perhaps someone who knows is more likely to have a thorough understanding— and thus a closer connection to the world— than someone who has a stable true belief. Now, as far as I can tell, having a thorough understanding is just to know more things. For example, in addition to knowing that FLT is true, Wiles knows how it follows from other theorems. If this is correct, then one response is, as before, to retreat to the extended main claim, and maintain, that knowing is no more valuable than a certain group of stable true beliefs. But, as before, we can do better and maintain the main claim. One can insist that a stable true believer is no less likely to acquire a thorough understanding than a knower. A stable true believer is no less likely to be exposed to situations in which surrounding stable true beliefs can be acquired. Although the retreat to the extended claim will do, I prefer the latter response, in order to maintain the main claim.[4]

4.2.5   Valuable as a means

I may have missed something when arguing that truth-tracking belief is no more valuable than stable true belief. Nozick himself (in a section called "What's so special

---

[4]S. Haslanger[34] has a interesting suggestion worthy of mention. She suggests that knowledge is valuable because knowers have and respond to reasons, and being responsive to reasons is part of what it is to be autonomous— which is a constitutive part of a flourishing life which is valuable for creatures like us. But if comparison between knowing and having a stable true belief is restricted to people who are rational in the sense of being responsive to evidence, then the value that Haslanger points to is plausibly present for a rational person, whether or not he has knowledge.

about knowledge?") posits a value for truth-tracking belief. Nozick says that evolutionary processes can't directly put true beliefs in our heads— they can only produce a capability to form true beliefs in a changing world. But the kind of capability which evolution can produce is a capability to form beliefs which track the truth.[50, p 283ff]

It is difficult to evaluate evolutionary considerations at this level of detail. But Nozick's thought suggests an interesting proposal. How can you get a stable true belief? Perhaps an effective way is to get— or at least try to get— knowledge.[5]

If this is right, then knowing that p may have a value that a stable true belief that p lacks. Knowing that p is useful as a means to have a stable true belief that p. Or, at least, trying to know that p is useful as a means to have a stable true belief that p.

I am unsure how to deny that knowing has this extra value— that knowing is valuable as a means to stable true belief. But then the main claim— that knowledge is no more valuable than stable true belief— is false.

However, admitting that the only extra value had by knowing is that knowing is valuable as a means to stable true belief, is not to accept that knowing has a significant value over and above stable true belief. For it is to accept that knowing is valuable because stable true belief is. Stable true belief is the valuable end, not knowledge. If stable true belief is not worth much, then a means solely to reach stable true belief is not worth much either.

Even so, some readers may wish to add some fine print to the main claim: knowledge is no more valuable than stable true belief, *except insofar as knowledge is valuable as a means to stable true belief.*

4.2.6 Knowledge is worthless

An unusual proposal is that knowledge has no value. Jonathan Kvanvig[42] makes this surprising claim, arguing that we value other things which we confuse with knowledge. Suffice it to say that if knowledge has no value, the main claim is correct. If knowledge has no value, knowledge surely has no more value than stable true belief.

---

[5]A. Byrne suggested something like this to me.

4.2.7 Knowledge is valuable in itself

I have discussed several different kinds of value which knowledge might have. The first proposal was that knowledge is useful for action. The second proposal was that knowledge is valuable for others. Both of these two proposals see knowledge as valuable for the sake of something else. The third proposal was that knowledge is valuable because it is an achievement. The fourth proposal was that knowledge is valuable because knowing brings a close connection to the world. I have suggested that these proposals see knowledge as valuable because something other than knowledge is valuable. The fifth proposal was that knowledge is useful in getting stable true beliefs. Again, this proposal sees knowledge as valuable for the sake of something else.

Here are two distinctions in value.[6] The first distinction is between being *valued for its own sake (as an end)* and being *valued for the sake of something else.* For example, if knowledge is valued because it is useful for action, then knowledge is valued for the sake of something else— successful action. The second distinction is between being *valuable in itself* and being *valuable in virtue of something else.* For example, if knowledge is valuable because it is an achievement, then knowledge is valuable because something else— achievement— is valuable. The first distinction is a distinction in the way we value things, as ends or for the sake of something else. The second is a distinction in the way things have value, in themselves (intrinsically) or in virtue of something else (extrinsically).

So far I have discussed every proposal I can find according to which knowledge is valued for the sake of something else. And I have discussed every proposal I can find according to which knowledge is valuable in virtue of something else. (Some of the proposals I have discussed fall into both of these categories.) I have argued that, if knowledge has that value, then stable true belief has that value too. (Setting aside the one exception discussed in the last section— knowledge as valuable for the sake

---

[6]C. Korsgaard[40] draws similar distinctions and argues that should not be collapsed together. I am indebted to Langton[43], who amends and criticizes Korsgaard's discussion. However, I ignore many of the complexities which Langton uncovers.

of stable true belief.) It seems fair to conclude that if knowledge is valuable for the sake of or in virtue of something else, then stable true belief has that value too. But that leaves two open questions: is knowledge valued as an end? is knowledge valuable in itself?

There is evidence that it is. People talk about the value of learning for its own sake. Research projects get funded on the grounds that they might bring us new knowledge. Scholars spend years trying to answer obscure questions. But we need stronger evidence than this, if we are to conclude that knowledge is valuable in itself, or valued as an end.

A test for intrinsic value comes to us from G. E. Moore. To apply the *isolation test* to see whether X has intrinsic value, you imagine a world in which nothing exists except X, and ask whether X has value. Knowing presumably does not exist without a knower, so imagine a world in which the only existing thing is a person P_1 who knows something. Since nothing exists except P_1, one thing he might know about is himself. So let's suppose that he knows that he exists. Is his knowing valuable? Maybe. I am not sure. At least, I don't see any contrast with stable true belief. If I imagine a different world in which nothing exists except P_2, who has a stable true belief that he exists, and ask whether his having a stable true belief is valuable, I have the same reaction. As far as I can tell, this test only shows that if knowledge has intrinsic value, then stable true belief does too.

Here is a second test. To apply the *choice test* to see whether X is more valuable than Y, ask yourself whether you would rather choose a life in which you have X or one in which you have Y.[7] This is not directly a test whether knowing is valuable in itself, or as an end. However, given the conclusion that knowing is no more valuable than stable true belief in virtue of, or for the sake of something else, this test will indirectly help decide whether knowing is valuable in itself or as an end.

So here goes. Would you rather choose a life in which you know that the Federated States of Micronesia gained independence in 1989, or one in which you merely have a stable true belief that the FSM gained independence in 1989? Again, this is not

---

[7]N. Hall suggested this test.

clear.

Here's a thought that may decide the question. In the life in which you merely have the stable true belief, you may have the stable true belief because you concluded it from a false belief, for example the patently false belief that all U. S. territories gained independence in 1989. But in the life in which you know, you couldn't have gained this knowledge by drawing a conclusion from this false belief. Otherwise you would not have knowledge. Choosing the life with the mere stable true belief risks having such patently false beliefs. So, since false beliefs are bad, you should choose the life in which you know so as to avoid false belief.

This is not a good reason to choose the life in which you know. For you might still have this false belief– indeed you might have many, many false beliefs— in the life in which you know. You just can't gain your knowledge by drawing a conclusion from such a falsehood. And in the life in which you merely have the stable true belief, you might, in addition have a multitude of other stable true beliefs, comprising a complete and comprehensive picture of the world. Thus, avoidance of false belief is not a good reason to choose the life in which you know.

So let's try again. Would you choose the life in which you know over the life in which you merely have a stable true belief? I find myself without a clear answer. (I suppose someone might say, "Of course, I'll choose the life in which I know because knowing is intrinsically valuable." But that won't help here, when trying to use the choice as a test to decide whether or not knowing is intrinsically valuable.) This test has not helped decide whether knowledge is valuable in itself.

Although more could be said, I believe that it is now time to give up the search for some value of knowledge over stable true belief, and tentatively conclude that the main claim is correct: knowledge is no more valuable than stable true belief (except insofar as knowledge is valuable as a means to stable true belief). I will now see what conclusions follow from this claim about the value of knowledge.

4.3   Conclusions

4.3.1   Skepticism not considered dangerous

Skepticism about knowledge is an affront to common sense. No one can know anything about the past? Ridiculous. No one can know anything about the external world? Absurd. No one can know anything at all? Crazy. According to common sense, we know many things; according to common sense, such skeptical claims are obviously false.

A difficulty arises. There are some appealing arguments leading to skeptical conclusions about knowledge. Many philosophers have faced this difficulty, strenuously laboring to find flaws in skeptical arguments about knowledge.[8] But we should pause to ask whether such effort is necessary.

After all, there is an easier way out of the difficulty. We might instead— to consider the extreme case— accept that no one can know anything. If the main claim of this essay is correct then we need not hesitate to accept that no one can know anything; skepticism about knowledge is not intellectually threatening. If knowledge is not stable true belief, but knowledge is no more valuable than stable true belief, then even if no one can have knowledge, one can have something just as valuable: stable true beliefs.

Is this an acceptable response to skepticism about knowledge?

I will now argue that it is.

Reasons for

The first reason that this response to skepticism should be accepted is that it fits well with the way people actually respond to skepticism about knowledge. In my experience the typical neophyte philosophy student gets convinced by skeptical arguments about knowledge; the typical student agrees that, strictly speaking she doesn't know anything. But that doesn't bother her, she says, as long as that does not mean that her ordinary beliefs about the world are unreasonable, or should be given up.

---

[8][39] and [54] contain recent attempts.

The typical neophyte philosophy student is often puzzled by philosophers like Barry Stroud who get themselves worked up about skepticism conclusions about knowledge:

> "The consequences of accepting Descartes's conclusion as it is meant to be understood are truly disastrous. There is no easy way of accommodating oneself to its profound negative implications." [63, 38]

I don't mean to rely too heavily on the evidence of my own personal experience, but it does seem to me that my experience is not unusual. (Perhaps one day an experimental scientist will do philosophers a favor and get some good data about how non-philosophers respond to skepticism.)

A second reason why this response to skepticism should be accepted is that it is compatible with a straightforward explanation why some skeptical arguments are compelling: some skeptical arguments about knowledge are sound. Those who try to dismantle a skeptical argument have the burden both of explaining why that argument is unsound, and why it appears sound. They try to explain both why we are almost taken in, and why we should not be taken in. The approach I am suggesting is more straightforward.

A third reason why this response to skepticism should be accepted is that other approaches are inadequate. I won't argue this point here, except to note that no approach to skepticism about knowledge is widely accepted (if not to simply ignore skepticism and think about something else).

Further reasons make this approach to skepticism about knowledge appealing. It is simple. It is general, applying to any argument that concludes that no one can (or does) know anything about a certain subject matter. Nonetheless, it is perfectly compatible with sometimes taking the usual approach for a skeptical argument: trying to find a flaw in that argument.

Reasons against

There are objections to this approach to skepticism. But these objections do not show that this approach is inadequate.

116

One objection is that skepticism about knowledge is an affront to common sense. According to common sense, skepticism about knowledge is obviously false. So it is a mistake to accept skepticism about knowledge. In reply, if this objection is not to be simple intransigence, one needs to see why common sense should carry the day. That a belief is a common sense belief is certainly not, by itself, sufficient reason to reject every argument to the contrary. So-called common sense has failed us in the past. It was once common sense that women are less rational than men. It was once common sense that slavery is morally permissible. Those common sense beliefs have, thankfully, been left in the past. This is not to deny that common sense has weight; it is where we start after all. But rejecting this approach to skepticism merely by claiming that skepticism about knowledge conflicts with common sense is very weak.

A second objection is that it is not clear that we have stable true beliefs. Why is it reasonable to believe that we have stable true beliefs? This objection looks serious. If we cannot have stable true beliefs, then my main claim leads to an uninteresting conclusion: even if we cannot have knowledge there is something else, just as valuable, that we cannot have either. In reply: we have very good evidence that at least some of our beliefs are both true and stable. Here is just one example. I believe that I have blue eyes. Every time I look in the mirror I get more evidence that this belief is true; my memory testifies that I have had this belief for a long time.

A third objection is that skeptical arguments about knowledge extend to skepticism about stable true belief. A good skeptical argument about knowledge (if there are any) might not simply threaten our knowledge alone. It may threaten to show that none of our beliefs are stable or true. If so, then saying that even if one cannot have knowledge one can have something of equal value– stable true beliefs– is a Pyrrhic victory. I grant that the approach I have suggested for skepticism about knowledge would be feeble indeed if accepting that no one can have knowledge about some subject matter required one to accept that no one can have stable true beliefs about that subject matter. Whether one form of skepticism leads to the other is not clear. But if it does– if a skeptical argument about knowledge does lead to skepticism about stable true belief, then, admittedly, a different response to skepticism is

needed. I recommend the response in Chapter 3, where I argued that certain skeptical arguments can be rationally ignored.

Setting skepticism aside, let me now turn to other lessons to be drawn from the main claim.

4.3.2 Knowledge devalued?

The main claim that knowledge is no more valuable than stable true belief is initially surprising. One might have thought that knowledge has a unique and special value, that knowledge is the worthy goal of laborious effort, that knowledge is to be cherished when attained. If the main claim is correct, then there is something just as valuable as knowledge, which is not knowledge. So there is no special value unique to knowledge alone.

However, it certainly does not follow from the main claim that knowledge is not the worthy goal of laborious effort, or that knowledge is not to be cherished when attained. Rather, if knowledge is the worthy goal of laborious effort, then stable true belief is too. And if knowledge is to be cherished when attained, then stable true belief is to be too. The main claim does not devalue knowledge; it reveals that there is something else which is just as valuable.

Still, to the extent that our educational practices and methods of scholarship are aimed at achieving knowledge rather than mere stable true belief, our practices and methods may need revision. Maybe it is easier to get stable true beliefs than knowledge; if so, we should give up our attempts to gain knowledge, and turn to stable true belief activities.

Even so, I doubt that our practices need much revision. First, knowledge is a close cousin of stable true belief, even if they differ. Assuming that knowledge requires true belief, both require the hard work of getting true belief. That is no mean matter. Second, as pointed out earlier, maybe a good way to get stable true beliefs is to get knowledge. Or maybe a good way to get stable true beliefs is to seek knowledge. If so, knowledge seeking practices should remain in place, as long as we are trying to get stable true beliefs.

4.3.3 Analyzing the concept of knowledge

The main claim sharply constrains any feasible analysis of the concept of knowledge into true belief plus something more. If knowing that p requires a true belief that p, then knowing that p requires an equally stable true belief that p. So, if knowing that p requires a true belief that p plus something more, then knowing that p requires an equally stable true belief that p plus something more. But what value does that something more add to the value of knowing that p? Call that something more 'X'. If the main claim is correct, then knowing that p is no more valuable than an equally stable true belief that p. So if the main claim is correct, and if knowing that p requires a true belief that p plus X, then, while the presence of X might help bring stability, beyond that X adds no value. (The presence of X before a person comes to have a true belief might be valuable in helping to get a true belief. But once the true belief is held X no longer has this value. Also X might be valuable on its own— it might, for example, be a stable true belief. But that does not add value to the true belief that p.) If X is justification, then justification adds no value to the true belief beyond helping make the true belief stable. If X is justification plus something to solve the Gettier problem, then justification plus something to solve the Gettier problem adds no value beyond helping to make the true belief stable.

This is surprising. But any attempt to produce an X is constrained by this fact: If X adds any value to true belief, it adds stability.

4.3.4 Epistemic concepts

The foregoing considerations inspire a bolder claim: belief, stability and truth are the central concepts for epistemology, not knowledge. If knowledge is true belief plus X, then knowledge is a species of stable true belief. If, in addition, the main claim is correct, then knowing is just one way to be in the valuable state of believing stably and truly. Arguably, the central theoretical concepts are the ones needed to talk about the valuable state: belief, stability, and truth. The concept of knowledge may be interesting, but it is not the central concept.

119

I offer the bolder claim only as a speculation. It would take several essays defend it. However, it is interesting to note, as Timothy Williamson does, that "in recent decades, questions of knowledge seem to have been marginalized by questions of justification. ... Once Gettier showed in 1963 that justified true belief is insufficient for knowledge, and therefore that knowledge is unnecessary for justified true belief, it became natural to ask: if you can have justified true beliefs, why bother with knowledge?" [71, p 184] (Williamson, by the way, disapproves. He thinks knowledge is the unanalyzable central concept of epistemology.)

Failure to analyze the concept of knowledge may have led some philosophers to give up on the concept of knowledge. But if the bolder claim is correct, even successfully analyzing the concept of knowledge would not vindicate the concept of knowledge.

4.4 Final remarks

The main claim I have defended in this essay is that knowledge is no more valuable than stable true belief. (With one possible exception: knowledge may be valuable as a means to get stable true belief.) I drew three conclusions from the main claim.

1. Skepticism about knowledge is harmless.

2. Knowledge is not devalued. It is doubtful that we should give up our knowledge gathering practices and methods.

3. Any attempt to analyze knowledge into true belief plus X is severely constrained by the following fact: if X adds any value to the true belief, it only adds stability.

Finally, I added one speculative thought: knowledge is not an important concept for epistemology; belief, stability and truth are.

Bibliography

[1] Jonathan E. Adler. *Belief's Own Ethics*. MIT Press, Cambridge, MA, 2002.

[2] William Alston. Epistemic circularity. *Philosophy and Phenomenological Research*, 47(1):1–30, 1986.

[3] William Alston. *Epistemic Justification*. Cornell University Press, Ithaca, New York, 1989.

[4] William Alston. Level confusions in epistemology. In *Epistemic Justification* [3], pages 153–171.

[5] William Alston. What's wrong with immediate knowledge? In *Epistemic Justification* [3], pages 57–78.

[6] William Alston. *The Reliability of Sense Perception*. Cornell University Press, Ithaca, New York, 1993.

[7] Robert Audi. *The Structure of Justification*. Cambridge University Press, Cambridge, UK, 1993.

[8] Jonathan Barnes. *The Toils of Scepticism*. Cambridge University Press, Cambridge, MA, 1990.

[9] Michael Bergmann. Epistemic circularity: Malignant and benign. Ms., 2003.

[10] Laurence BonJour. Can empirical knowledge have a foundation? *American Philosophical Quarterly*, 15:1–13, 1978.

[11] Laurence BonJour. Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5:53–73, 1980.

[12] Laurence BonJour. *The Structure of Empirical Knowledge*. Harvard University Press, Cambridge, Massachusetts, 1985.

[13] Myles Brand. *Intending and Acting*. MIT Press, Cambridge, Massachusetts, 1984.

[14] Michael Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, Massachusetts, 1987.

[15] Michael Bratman. *Faces of Intention*. Cambridge University Press, Cambridge, 1999.

[16] Myles Burnyeat. Can the sceptic live his scepticism? In Schofield et al. [58], pages 20–53.

[17] Stanley Cavell. *The Claim of Reason*. Oxford University Press, Oxford, 1979.

[18] Stanley Cavell. *A Pitch of Philosophy*. Harvard University Press, Cambridge, MA, 1994.

[19] Roderick Chisholm. *Theory of Knowledge*. Prentice-Hall, Englewood Cliffs, New Jersey, 3rd edition, 1988.

[20] Edward Craig. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press, Oxford, 1990.

[21] Michael R. DePaul. Value monism in epistemology. In Matthias Steup, editor, *Knowledge, Truth, and Duty*, pages 170–183. Oxford University Press, Oxford, 2001.

[22] David Embick. *Voice and the Interfaces of Syntax*. PhD thesis, University of Pennsylvania, 1997.

[23] David Embick. Participial structures and participial asymmetries. MIT Handout, 2000.

[24] Richard Feldman. *Epistemology*. Prentice Hall, Upper Saddle River, NJ, 2002.

[25] Richard Feldman and Earl Conee. Evidentialism. *Philosophical Studies*, 48:15–34, 1985.

[26] Robert J. Fogelin. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford University Press, Oxford, 1994.

[27] Richard Foley. *Intellectual Trust in Oneself and Others*. Cambridge University Press, Cambridge, 2001.

[28] Alvin Goldman. What is justified belief? In George S. Pappas, editor, *Justification and Knowledge*, pages 1–23. D. Reidel, 1979.

[29] Alvin Goldman. *Epistemology and Cognition*. Harvard University Press, Cambridge, Massachusetts, 1986.

[30] L. E. Hahn, editor. *The Philosophy of P. F. Strawson*. Open Court, Chicago, 1998.

[31] Gilbert Harman. *Change in View*. MIT Press, Cambridge, MA, 1986.

[32] Gilbert Harman. *Reasoning, Meaning and Mind*. Oxford University Press, Oxford, 1999.

[33] Gilbert Harman. Skepticism and foundations. Ms., 2000.

[34] Sally Haslanger. What is knowledge and what ought it to be. *Philosophical Perspectives*, 13:459–480, 1999.

[35] Christopher Hookway. *Scepticism*. Routledge, London, 1990.

[36] William James. *The Will to Believe and Other Essays in Popular Philosophy*. Dover, New York, 1956.

[37] Ward E. Jones. Why do we value knowledge? *American Philosophical Quarterly*, 34(4):423–437, 1997.

[38] Peter Klein. Human knowledge and the infinite regress of reasons. *Philosophical Perspectives*, 13:297–325, 1999.

[39] Peter Klein. Skepticism. *http://plato.stanford.edu/archives/sum2002/entries/skepticism/*, 2001.

[40] Christine Korsgaard. Two distinctions in goodness. *Philosophical Review*, 92:169–195, 1983.

[41] Angelika Kratzer. Building statives. Ms., 2000.

[42] Jonathan Kvanvig. Why should inquiring minds want to know? *The Monist*, 81(3):426–451, 1998.

[43] Rae Langton. Values, conditioned and conferred. Ms., 2003.

[44] David Lewis. Reduction of mind. In Samuel Guttenplan, editor, *A Companion to Philosophy of Mind*. Blackwell, Oxford, 1994.

[45] Margaret MacMillan. *Paris 1919: six months that changed the world*. Random House, New York, 2002.

[46] Alec Marantz. Reconstructing the lexical domain with a single generative engine. MIT Handout, 2000.

[47] Tatjana Marvin. Past participles in reduced relatives. Ms., December 2000.

[48] Alfred Mele. *Springs of Action*. Oxford University Press, Oxford, 1992.

[49] N. M. L. Nathan. *The Price of Doubt*. Routledge, London, 2001.

[50] Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, Massachusetts, 1981.

[51] Terence Parsons. *Events in the Semantics of English*. MIT Press, Cambridge, MA, 1990.

[52] Alvin Plantinga. *Warrant and Proper Function*. Oxford University Press, New York, 1993.

[53] Plato. *Meno*. W. R. M. Lamb, translator, Loeb Classical Library, volume 165, Harvard University Press, Cambridge, MA, 1924.

[54] James Pryor. The skeptic and the dogmatist. *Noûs*, 34:517–549, 2000.

[55] Dana Radcliffe. Scott-Kakures on believing at will. *Philosophy and Phenomenological Research*, 57(1):145–151, 1997.

[56] Wayne D. Riggs. Beyond truth and falsehood: The real value of knowing that p. *Philosophical Studies*, 107:87–108, 2002.

[57] Wayne D. Riggs. Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64(1):79–96, 2002.

[58] Malcom Schofield, Myles Burnyeat, and Jonathan Barnes, editors. *Doubt and Dogmatism*. Oxford University Press, Oxford, 1980.

[59] Dion Scott-Kakures. Belief and the captivity of the will. *Philosophy and Phenomenological Research*, 53(4):77–103, 1993.

[60] Matthias Steup. Foundationalism, circularity, and infinitism. Ms., 2003.

[61] P. F. Strawson. *Skepticism and Naturalism: Some Varieties*. Columbia University Press, New York, 1985.

[62] Gisela Striker. Ataraxia: Happiness as tranquility. In *Essays on Hellenistic Epistemology and Ethics*, pages 183–195. Cambridge University Press, Cambridge, 1996.

[63] Barry Stroud. *The Significance of Philosophical Scepticism*. Oxford University Press, Oxford, 1984.

[64] Peter Unger. *Ignorance: A Case for Skepticism.* Oxford University Press, Oxford, 1975.

[65] Peter Unger. *Philosophical Relativity.* University of Minnesota Press, Minneapolis, 1984.

[66] J. David Velleman. *Practical Reflection.* Princeton University Press, Princeton, 1989.

[67] J. David Velleman. *The Possibility of Practical Reason.* Oxford University Press, Oxford, 2000.

[68] Brian Weatherson and John Hawthorne. Beliefs old and new. *ms: http://web.syr.edu/weathers/belief.pdf,* 2001.

[69] B. A. O. Williams. *Descartes: The Project of Pure Enquiry.* Penguin, London, 1978.

[70] Michael Williams. *Unnatural Doubts.* Princeton University Press, Princeton, 1996.

[71] Timothy Williamson. *Knowledge and its Limits.* Oxford University Press, Oxford, 2000.

[72] Ludwig Wittgenstein. *On Certainty.* Basil Blackwell, Oxford, 1969.

[73] Linda T. Zagzebski. *Virtues of the Mind: An Inquiry into the Nature of Virtue and The Ethical Foundations of Knowledge.* Cambridge University Press, Cambridge, 1996.