

52. Mortimer Adler draws this distinction very clearly, in addition to making the case that the confusion between these two things has had a great influence on modern philosophy. It is the first mistake discussed in his *Ten Philosophical Mistakes*. Cf. Kelley, 36, 67, where the point is made that one does not perceive one's means of awareness.

53. Kelley, 66-69, makes this point; Rand, 80-82, appears to be making a similar one.

54. Locke, *Essay*, I.i, section 8 (altering punctuation and capitalization to conform to modern style; italics Locke's).

55. Berkeley, *Principles*, section 1; italics Berkeley's.

56. Berkeley, *Dialogues*, 113.

57. Berkeley, *Dialogues*, 114.

58. Kant, *Critique*, B40-45.

59. See Barker; Strawson, 249-56.

60. Kant, *Critique*, B298. Cf. B303, where he says, "that which is not appearance cannot be an object of experience" and Bxvii, where there appears the phrase, "the objects [of our intuition], or what is the same thing, . . . the *experience* in which alone, as given objects, they can be known" (emphasis Kant's).

61. Kant, *Prolegomena*, Remark II, 36.

62. Kant, *Critique*, B45.

V. A Version of Foundationalism

1. What Is Perceptual Knowledge?

Perception is direct awareness of the external world. *Perceptual knowledge*, however, is not the same thing as perception. Perception is awareness in the form of perceptual experiences (that is, awareness wherein the type of apprehension involved is a perceptual experience). Knowledge is awareness in the form of beliefs (awareness wherein the type of apprehension is a belief). A perceptual experience is not a belief, so perception itself is not a kind of knowledge.

To see this, consider a hypothetical case. Jack comes home one evening, and as he walks in the door, he sees a large pink rat on the dining room table. He firmly believes that there are no pink rats in reality, that it is highly unlikely that a rat of any kind should be on his table in any case, and that people periodically have hallucinations. Jack concludes that there is not really a pink rat on the table, and that he is merely hallucinating it. It does not matter for the example whether Jack is reasonable in drawing this conclusion; it matters only that he may in fact draw it. In fact, let us suppose, Jack is mistaken; there is nothing wrong with his visual system. There really is a pink rat on his table (never mind how it got there), and it is causing his visual experience in the way middle-sized physical objects normally do, by reflecting light to his eyes and so on. In other words, Jack is perceiving (is perceptually aware of) the pink rat. But does he *know* that there is a pink rat on his table? No, for he does not believe it; in fact, he believes the opposite. This shows that perception is not a kind of knowledge, nor does perceptual experience imply belief, for it is possible to have perception in the absence of knowledge and belief.

Knowledge is traditionally defined, approximately, as justified, true belief. In other words, in order to know that something is the case, one must believe that it is the case, one's belief must be true, and one must be justified (in the

From Huemer, M., 2001
Skepticism + The Veil
of Perception

internalist sense) in that belief. I say knowledge is defined "approximately" this way, because it is widely recognized that at least some fourth condition beyond the three just mentioned is needed, although there is no general agreement on what that condition is.¹ Here is a well-known example: suppose I look at a clock to find out what time it is. The clock reads 3:00 and I both correctly and justifiably accept that the time is approximately 3:00. But suppose that, although I have no reason to suspect this, the clock is actually stopped (even a stopped clock is right twice a day). In this case, I do not genuinely *know* that the time is 3:00, even though I have a justified, true belief that it is.² There is a voluminous (a bit too voluminous, in my opinion) literature on the definition of knowledge, mostly concerned with how one should deal with cases like this (so-called "Gettier cases"). My aim is not to resolve the disputes that have arisen in that literature, so I will only briefly mention how I think the definition of knowledge may be amended to deal with such cases: one may add the condition that there be no *defeaters* for the justification of one's belief. A defeater for your justification for *P* (where *P* is something you justifiably believe) is defined to be a true proposition that, if it were added to your beliefs, would render you no longer justified in believing that *P*. Thus, in the clock example, the proposition "the clock is stopped" is a defeater, and this explains why I do not have knowledge.³

Notice how this traditional definition of knowledge mirrors my account of awareness. In awareness, one must have an apprehension, its content must be at least roughly satisfied, and it must be nonaccidental that its content is satisfied. Knowledge is simply one form of awareness: the form in which the apprehension is a belief, the content of the belief is satisfied (since the belief is true), and it is nonaccidental that the belief is true (since the believer has an adequate and undefeated justification for it). Seen in this way, my proposed definition of knowledge as *awareness in the form of beliefs* is not incompatible with the traditional *justified true belief* definition; the latter is merely a more precise version of the former, in which it is specified more precisely what it takes for a belief to be *nonaccidentally* true.

Perceptual knowledge is a particular kind of knowledge, namely, the kind that involves perceptual beliefs. What is a perceptual belief? It is a belief that is directly based on a perceptual experience. I say "directly" because there are cases in which perceptual experiences lead to certain beliefs only through a chain of inferences—as in the case of scientific theories—and I would not call the beliefs at the end of the chain "perceptual beliefs." In general, an apprehension, *B*, is directly based on another apprehension, *A*, when *B* is based on *A*, and there is no intermediary apprehension, *C*, such that *A* causes *B* through causing *C*. Inference is a particular species of the basing relation; it is the process by which one belief comes to be based on another belief. When I say

that I inferred *Q* from *P*, I am saying that I believed *P*, I came to believe *Q*, and my belief that *Q* was based on my belief that *P*.

From what I have said here and in chapter IV (section 2), it follows that perceptual knowledge, if it exists, would be a form of indirect awareness, but noninferential knowledge. (All inferential knowledge is indirect awareness, but not all indirect awareness is inferential knowledge.)

2. Do We Have Perceptual Knowledge?

There are perceptual beliefs—even an external world skeptic can agree with that. When we have perceptual experiences, we normally are prompted, without any intermediary thoughts or other apprehensions, to accept propositions about the external world. For instance, when I look at a tomato on the counter, the visual experience I have will typically cause me to believe, at the least, that there is a red, round thing there. (It is due to various background beliefs of mine that I accept the further proposition that the thing in question is a tomato. The latter belief, then, is not purely perceptual because it is based on my perceptual experience *together with* background knowledge.) Furthermore, these beliefs are logically related to the experiences that cause them in the way appropriate for the "based on" relation: the content of my perceptual experience entails the content of my belief. As we indicated earlier (section IV.4.2), the content of a perceptual experience will normally not be the same as that of the perceptual belief it causes; rather, it will be far more specific and detailed. But we can be assured that it will entail the belief because there is a determinable-determinate relation between the two.⁵

It is easiest to explain the determinable-determinate relation with examples. Consider the relation between *shape* and *squareness*. Shape is a determinable property; squareness is a more determinate property, relative to shape. Similarly, *red* is a determinate, relative to the determinable *color*. A particular shade of red, say crimson, would be a determinate relative to *red* (and also relative to *color*). In these examples, we see that one property is more specific than, and falls under, the other. Crimson is a specific kind of red, so for a thing to be crimson implies that it is red, but not vice versa. Squareness is a particular kind of shape, so for a thing to be square implies that it is shaped, but not vice versa. Since our perceptual experiences are more fine-grained in their content than our concepts, and since the beliefs we form are constructed from our concepts, the properties represented in perceptual experience will typically be determinates, relative to those represented in beliefs. Thus, my perceptual experience might represent that there is a thing with a certain very specific color (one of the two million colors the visual system can distinguish), while

my belief might represent simply that there is a red thing. The color represented by my visual experience will be a specific *kind* of red. This is why the contents of my perceptual experiences will entail the contents of my perceptual beliefs, but not vice versa.

The question of whether we have perceptual knowledge comes down to the question of whether our perceptual beliefs are true and justified, in the absence of defeaters. Although the truth condition is theoretically distinct from the justification condition (one can have true but unjustified beliefs or justified but false beliefs), we cannot really address them separately. An account of why our perceptual beliefs are justified just is, ipso facto, an account of why we should take them to be true (this is a tautology). If I could successfully argue that our perceptual beliefs are *true*, then, upon reflecting on that argument, we would also see that we are *justified* (by that very argument, if nothing else) in holding those beliefs. Conversely, if I could show that we are justified in our perceptual beliefs, then we would thereupon see that we should accept those beliefs (or rather, we should reaffirm them, since we have already accepted them prior to philosophical reflection), which is equivalent to accepting them as true. Although in general the proposition, "*S* is justified in believing *P*" does not imply "*P*," if I believe "I am now justified in believing *P*," I must believe "*P*" as well (unless I am irrational).

The skeptical challenges to our perceptual knowledge with which we are concerned are all arguments against the *justification* of our perceptual beliefs, not against their truth. The skeptic does not say that our perceptual beliefs are all false; he only says they may or may not be true. But the skeptic does not say they may or may not be *justified*; he says they *are not* justified (recall chapter II). Nor, by the way, does the skeptic say there are defeaters for our perceptual beliefs (again, there may or may not be defeaters). Thus, in answering the question of how we can have perceptual knowledge, the main issue to address is how our perceptual beliefs can be justified.

Some philosophers would seek to avoid skepticism by denying that justification is part of the meaning of "knowledge" in ordinary English.⁶ They would maintain that we can have "knowledge" even if our beliefs are unjustified. I think these philosophers are mistaken, but that isn't really important. Notwithstanding the tradition of analytic philosophy, the important question here is not semantic. The important question is how we might justify our own beliefs about the external world, to ourselves; this question remains of interest even if I am mistaken about what "know" means in ordinary English. If you think I am thus mistaken, then you may regard this chapter as an inquiry after a special, philosophers' sense of "knowledge."

My position, then, is that our perceptual beliefs are justified *by* the perceptual experiences on which they are based. A perceptual experience can

justify a belief in much the same way that a belief can justify another belief. If I believe *P*, and *P* entails *Q*, and I come to believe *Q* by correctly inferring *Q* from *P*, then my belief that *Q* is justified by my belief that *P*. The essential features of this situation are mirrored in the case of perceptual experiences and perceptual beliefs; namely, that there is an entailment relation between the two apprehensions and that the one causes the other by virtue of the logical relation between them.

But, it will be objected, one belief can only justify another belief if the first belief is, itself, justified. Similarly, therefore, shouldn't we say that a perceptual experience can only justify a perceptual belief if the experience, itself, is justified?

But the latter condition makes no sense. It does not make sense—it is a category error—to say that an experience is justified or unjustified. (A category error is the mistake of attempting to apply a predicate to a kind of thing to which it logically cannot apply. For instance, to say "The number 7 is red" would be a category error, since numbers cannot be colored; only physical objects and light can be colored.) You may be justified or unjustified in accepting a belief, but you cannot be justified or unjustified in having experiences; you simply have them. If I go near a fire, I just *will* feel a sensation of warmth; what would it mean to say I was "justified" in feeling warm?

To say that a person is unjustified in doing *x* is to say that, in some sense or for some reason, he shouldn't do *x*. Thus, you may say, "The American invasion of Vietnam was unjustified"; this would be to say that the Americans should not have invaded Vietnam. Interestingly, the predicate can also be applied to emotions. Suppose I say, "Your resentment toward Fred is unjustified; he hasn't done anything wrong." This would be the same as to say, "You shouldn't resent Fred; he hasn't done anything wrong." Likewise, if I say, "Belief in God is unjustified," this means, "One shouldn't believe in God." So the reason why an *experience* cannot be justified or unjustified is that it cannot be the case that you *should* or *shouldn't* have an experience.

There is a deeper explanation, for why it does not make sense to say a person should or shouldn't have a perceptual experience, and that is that perceptual experiences are automatic responses to external stimuli. Although I can choose to open or close my eyes, I cannot control what I will see if my eyes are open. My evaluations of my experiences will have no effect on those experiences. For instance, if I judge my present visual experience to be a mirage, and thus a misrepresentation of reality, this judgement will have no effect on the experience itself; I cannot stop "seeing" the mirage. Matters are otherwise when it comes to actions, beliefs, and emotions. If I come to judge that action *A* is morally obligatory, this will normally (except in cases of weakness of will) result in my doing *A*. If I come to judge that *P* is justified,

this will normally result in my believing that *P*. If I come to judge that Fred did nothing wrong, this will normally result in my ceasing to feel resentment toward Fred, if I felt such resentment initially. Of course, it is possible that these results not occur—it is possible for a person to act, think, or feel irrationally. That is the point of the “should” statements, after all; they would have no point if our actions, thoughts, and feelings were automatic.

There will be some who will say that their feelings *are* automatic, in just the way I have said perceptual experiences are. I think this is false, because I think that an emotion is a kind of evaluation, or comes about through one's evaluation, of the events towards which one has the emotion. Rational-emotive therapy has demonstrated the possibility of controlling one's emotions through these evaluations.⁷ Those who say they cannot control their emotions, I believe, are those who simply choose not to. But let us not argue this tangential point further; the point to make is that *if* emotions were purely automatic responses, then it would make no sense to speak of an emotion as justified or unjustified. And similarly, if (as is in fact the case) perceptual experiences are automatic responses, then it makes no sense to speak of them as justified or unjustified.⁸

But it is unclear where this leaves us, *vis-à-vis* the justification of perceptual beliefs. We have agreed that when a belief, *B*, is based on and logically supported by another belief, *A*, if *A* is justified *B* will be justified, and if *A* is unjustified *B* will be unjustified. Given that perceptual experiences are neither justified nor unjustified, it is difficult to know how to extend the analogy. We might say that, since perceptual experiences cannot be unjustified, perceptual beliefs based on them cannot be unjustified. But it would make just as much sense to argue that since perceptual experiences cannot be justified, perceptual beliefs cannot be justified.

Clearly, to explain why perceptual beliefs are justified, it is not enough to say that they are based on perceptual experiences that are neither justified nor unjustified. Consider another kind of state that is neither justified nor unjustified: imagination. Suppose I *imagine* that *P*, and this induces me to believe that *Q*, where *Q* is something that logically follows from *P*. I am neither justified nor unjustified in imagining that *P*. But obviously I am *unjustified* in believing that *Q* on the basis of it. Perceptual experiences likewise have no justificatory status, but I claim that a belief based on one is *justified*. Why?

3. A Principle of Foundational Justification

My theory is a version of foundationalism. Foundationalism says that there are certain beliefs, the so-called “foundational beliefs,” which we are justified in holding and which do not depend on any other beliefs for their justification.

Perceptual beliefs, in my view, are foundational. Notice that the definition of foundational beliefs does not say that they do not depend on *anything else* for their justification; it says they do not depend on *other beliefs* for their justification. In my view, perceptual beliefs certainly do depend on something else for their justification: namely, perceptual experiences.

As we mentioned earlier (section II.1), foundationalists must explain what differentiates foundational beliefs from merely arbitrary beliefs. They must state a condition or set of conditions under which a person has foundational (noninferential) justification for believing a proposition. A “principle of foundational justification” is a principle identifying such a condition. A foundationalist might put forward any number of principles of foundational justification, to account for different kinds of foundational beliefs. I believe, however, that a single principle of foundational justification can account for all foundational beliefs.⁹ The principle is what I call the rule of Phenomenal Conservatism (PC) (this is using “phenomenal” in the sense of “pertaining to appearances,” not in the sense of “fantastic”):

(PC) If it seems to *S* as if *P*, then *S* thereby has at least *prima facie* justification for believing that *P*.

In the remainder of this section, I explain what that means.

First, note that its seeming to *S* as if *P* is a distinct state from *S*'s believing that *P*.¹⁰ This is important, since otherwise PC would be granting foundational justification, automatically, to all beliefs, and this is not what we want; we want to identify a special class of foundational beliefs, to be distinguished from merely arbitrary beliefs.

One kind of seeming-as-if state is perceptual, and the argument showing this state to be different from belief was that it is possible for it to seem to you as if *P*, even when you do not believe that *P*; in fact, even when you know that $\sim P$.¹¹ There are also other, nonperceptual seemings, that is, cases in which it seems to you as if *P* is not an aspect of your perceptual experience. There are memory-related seemings:¹² for example, I seem to remember that Saturn is the fifth planet from the sun. The argument showing this to be different from belief parallels the argument in the perceptual case: we can easily imagine a case in which it seems to me as if *P* in just this way but in which I do not trust my memory and so do not believe that *P*. (The argument would be at least as strong for experiential memory as it is for factual memory.) Likewise, there are intellectual seemings: when I think abstractly about the proposition, “The shortest path between any two points is a straight line,” it seems to me to be true. Philosophers commonly call these intellectual seemings “intuitions.” Intuitions can be distinguished from beliefs by means of the same kind of

argument we have used above: it is possible for a person to have the intuitive sense that P but not to trust his own intuitions, and so to not believe that P . As a case in point, consider the intuitively obvious proposition, "In a plane containing a line and a point not on the line, there exists exactly one line parallel to the given line passing through the given point." There are some physicists who say that this proposition, as obvious as it seems, is false. Nor do these scientists lack the intuitions the rest of us (and the rest of the mathematicians and scientists in history) have. They admit that the proposition *seems* obviously true, but they say that our intuitions on this matter are not to be trusted and that there is powerful empirical evidence against it. (I refer, of course, to the non-Euclidean geometry used in Einstein's general theory of relativity.) The point here, as in the earlier examples, is not whether the scientists are reasonable to make this judgement, but simply that it is possible to make such a judgement. This suffices to show that its seeming to one, intellectually, that P is different from one's believing that P .

Next, we consider the notion of prima facie justification.¹³ When a belief is said to be prima facie justified, two things are meant: first, that the belief does not depend on other beliefs for its justification, and second, that the belief's justification can be defeated by countervailing evidence. In other words, prima facie justification is both *foundational* and *defeasible*. Now, many people feel this combination of properties to be paradoxical, but it really is not. Compare the legal concept of the presumption of innocence: the defendant is presumed innocent, until proven guilty. This means that the defense need not present evidence proving the defendant's innocence (though they *may* do so, of course). If no proof is brought forward during the trial, neither of the defendant's innocence nor of his guilt, then the defendant is acquitted. But of course, this does not mean that a defendant can never be convicted. If the prosecution brings forward enough positive evidence that the defendant is guilty, then the presumption of innocence is overcome and the defendant is convicted.

Prima facie justification in epistemology works similarly. According to phenomenal conservatism, the epistemological default position is to accept things as they appear. The appearances are presumed true, until proven false. That means that when it seems as if P and no evidence emerges contravening P , it is reasonable to accept P . (Since, as I claim, phenomenal conservatism is the sole principle of foundational justification, "evidence against P " would consist of other things that seem to be the case and that, directly or indirectly, either contradict or render it improbable that P .) But that does not mean that P can never be refuted, of course; it is possible that sufficient evidence will emerge against P to overcome its presumption. Prima facie justification only provides us with an initial starting point; it does not guarantee that we have to

end up where we started.

For a concrete example, take the Müller-Lyer illusion (figure 5.1).




Figure 5.1. The Müller-Lyer illusion.

The top line appears to be longer than the bottom line, so other things being equal, you would be reasonable in thinking that the top line is longer. However, you can get out a ruler and measure these lines. If you do this, you will find them to be of the same length. At that point, it would be reasonable for you to revise the initial belief and conclude that the lines are really of the same length. This is because you know that measurement is generally a more reliable way of determining lengths than just eyeballing things. Of course, your new belief, that the lines are of equal length, is also based on how things seem, including what you seemed to see when you were measuring the lines; there is no (rational) way around that. It is just that you rightly take the latter seemings to be more reliable. As to why we take measurement to be more reliable than unaided visual length-estimates, I think this is partly a matter of past experience (beliefs based on measurements tend to exhibit fewer and less serious conflicts with each other and with other beliefs than unaided visual estimates) and partly due to the fact that measurement relies on situations in which what seems to be the case is particularly clear and unambiguous—which, of course, is the point of measurement procedures. This example shows that we normally treat perceptual beliefs as having just the sort of presumption of truth that phenomenal conservatism attributes to them.

Now why do I say, in my statement of PC, that S has "at least" prima facie justification? Well, I don't want to rule out the possibility of cases in which S would have incorrigible justification for believing that P . That is, I want to allow that there may be (as in fact I think there are) cases in which S has justification for believing P not dependent on other beliefs, and in which this justification *cannot* be defeated by further evidence. An example would be my present belief that for every number x , $(x + 1)$ is greater than x . I think I am not

only justified in believing that, but furthermore, no future evidence could disconfirm it. Similarly, suppose I have a splitting headache. At this time, it seems to me that I am in pain, and I believe I am in pain. Again, I think my belief that I am in pain is not only justified, but justified in a way that is immune from countervailing considerations.

Those examples were not examples of *perceptual* beliefs, of course. I do not believe there are any examples of perceptual beliefs that have incorrigible justification. But I intend phenomenal conservatism to be a general principle of foundational justification. It explains not only why perceptual beliefs are noninferentially justified, but also why any other belief that is noninferentially justified is such. And this is important, for it means that my account of perceptual knowledge does not depend upon ad hoc principles of justification contrived specially to let in knowledge of the external world; I propose to account for perceptual knowledge by the same general principle I apply to all other kinds of knowledge. This should make my account more persuasive for those who are skeptical about perceptual knowledge but who allow knowledge of other kinds (that is, external world skeptics).

Notice that phenomenal conservatism does not hold that from the proposition, "It seems to me as if *P*," one may infer *P*. That would be a principle of inferential justification. What makes you justified in believing *P*, under PC, is not that you believe or know that it seems to you as if *P*, and you thereupon get to conclude that *P*; rather, it is the mere fact that it does seem to you as if *P* that makes you *prima facie* justified in believing *P*. Thus, to satisfy the antecedent of PC one does not need to have any beliefs—one does not even need to have the concept of an appearance, nor the concept of justification. In the case of perceptual beliefs, it is our perceptual experiences themselves—and not our assessments of, or beliefs about them—that justify our initial beliefs about the external world.

Even so, however, note that I am not saying a perceptual belief, or any other kind of foundational belief, *cannot* be supported by other beliefs. I am saying it *need not* be supported by other beliefs. These are two quite different things. Foundational beliefs are defined to be beliefs that do not depend upon other beliefs for their justification; they are not defined as beliefs that are not supported by other beliefs.¹⁴ To return to the legal analogy: the presumption of innocence in a criminal trial means that the defense is *not required* to give positive evidence of the defendant's innocence. It does not mean the defense is *prohibited* from giving positive evidence of the defendant's innocence. If they have such evidence to give, all the better for them. So there could arise a case in which the defendant's acquittal was overdetermined: suppose the prosecution failed to prove the defendant's guilt, and, in addition, the defense was able to prove his innocence. There is certainly no contradiction in that, and

in that case, there would simply be two conditions, each sufficient for acquitting the defendant. Likewise, there can be cases in which a belief has two kinds of justification, both a foundational one and an inferential one.

Incidentally, this is more than a merely theoretical possibility. There are many cases in which a belief has both kinds of justification. Suppose, for instance, that I see a red ball on the table. I am foundationally justified in believing that there is a red sphere there, due to the character of my perceptual experience. You happen to be in the same room with me, and I ask you, "Do you see a red sphere there?" whereupon you answer, "Yes." This gives me a second, inferential justification for thinking there is a red sphere there.¹⁵ Obviously, the existence of this second justification does not remove or change the character of the original, noninferential one.

Lastly, notice that phenomenal conservatism does not assert, as a general rule, "If it seems to someone as if *P*, then *P*." If it did, phenomenal conservatism would be false, for there are many cases in which it seems as if *P* but *P* is not the case. Nor does phenomenal conservatism assert or imply, "Most of the time, when it seems as if *P*, *P* is true." The latter is true, but it just is not what PC is about. Phenomenal conservatism is epistemological, not metaphysical. It says that when it seems as if *P* and there is no evidence to the contrary, it is *reasonable to believe P*. Compare the legal presumption of innocence again: the principle is not that defendants in general are innocent; it is not even that most defendants are innocent (in fact, quite the reverse is the case). Rather, it is the normative principle that if the prosecution fails to prove the defendant's guilt, he should be acquitted.

A corollary is that phenomenal conservatism is a necessary truth, not a contingent one. There is no possible world in which phenomenal conservatism is false, although there are possible worlds in which most of the things that appear to be so are not so. Recall the brain-in-a-vat scenario. If you were a brain in a vat, then most of the things that seemed to you to be the case would in fact not be the case. Nevertheless, it would still be reasonable for you to believe those things, given your situation, and given that you would have no evidence of your being a brain in a vat. The diabolical scientists stimulating your brain make it so you have a bunch of false beliefs, but they do not make it so you have *unjustified* beliefs; in fact, their method of getting you to adopt the false beliefs is precisely to give you justification for them.

4. In Defense of Phenomenal Conservatism

Why should we accept my proposed principle of foundational justification? I claim that PC is self-evident, once it is seen in its proper light. To see it in its

proper light, we have to remind ourselves of two things. The first is that the notion of “justification” we are talking about is internalist, rather than externalist. That is, we are talking about what propositions are justified from the subject’s own point of view. Thus, if *P* were false, but there were no way that *S* could reasonably be expected to guess that, then its falsity would have no bearing on *S*’s justification for believing *P* in the internalist sense. What is “unjustified” in the internalist sense is what one could be blamed for believing.¹⁶

The second thing we must remember is that the notion of justification we are concerned with is epistemic (rather than, e.g., prudential or moral). Epistemic justification is the kind of justification that is assessed from the standpoint of the pursuit of the truth and the aversion to error. The relation between justification and truth is like the relation between expected utility and utility; that is, the acceptance of (only) justified propositions is what constitutes the rational pursuit of the truth.

Now consider an account of rationality offered by Richard Foley: “A decision is rational as long as it apparently does an acceptably good job of satisfying your goals.”¹⁷ This is plausible. Notice the need for the qualifier, “apparently.” In order for a decision to be rational, it is not necessary that it *in fact* will satisfy your goals. A rational person can be mistaken about what will satisfy his goals, in which case he will do the thing that it appears to him will satisfy his goals rather than the thing that actually will. To return to my example from chapter II: it would have been irrational for the fisherman to just throw the pearl back into the ocean—even though that action would in fact have satisfied his goals better than keeping the pearl.

Now, if my goal is to have true beliefs and avoid having false ones, and if *P* seems to me to be true, while I have no evidence against *P*, then from my own point of view, it would make sense to accept *P*. Obviously, believing *P* in this situation will appear to satisfy my epistemic goals of believing truths and avoiding error better than either denying *P* or suspending judgement.

Another thing that will help to bring the self-evidence of PC into focus is reflection on the way you actually do form beliefs (omitting the cases of self-deception and leaps of faith, which are cases of epistemically irrational belief). When you are conscientiously seeking to know, you weigh a proposition in your mind. What determines whether you accept it or not? Well, in some cases you consider an argument for it. But this will help only if you accept the premises of said argument. How do you decide whether to accept a premise of an argument? You do not, in real life, demand an infinite series of arguments. Rather, you consider the premise on its own: if it is sufficiently obvious (to you) you accept it; otherwise, not. Moreover, it is difficult to imagine what possible alternative a person critical of this procedure could have in mind.

Should you, perhaps, accept the propositions that seem *false* instead? Surely this cannot be rational. Accept all and only the propositions that *are in fact* true? But acting on this advice is the same thing as accepting the propositions that seem true to you. Imagine I have a child who is having difficulty with fruit names. I tell him to put all the fruits that he takes to be apples into a basket. When he’s done, I empty the basket, start over, and tell him this time, put all the fruits that *are* apples into the basket. From my point of view, I might seem to have said something different, but (leaving aside the implication the child may pick up that he did it wrong the first time, which is not part of the literal meaning of what I said) from the child’s perspective I have essentially just repeated the same instruction, and he will obviously just do the same thing.

Or perhaps you should accept nothing at all? But half of your epistemic aim is to gain true beliefs. When *P* seems to you to be true and there are no grounds for doubting it (no defeaters), what more are you looking for? This is as good as it gets. Why should you withhold belief in *P*, keeping in mind the internalist sense of “should” (recall section II.6 above)—because it is still possible that *P* is false? But this seems to be just an extreme and unreasonable bias toward the “avoiding error” part of your epistemic goal, as against the “believing truths” part. You would not let the mere possibility that *P* is *true* suffice for you to *accept* it, so why let the mere possibility that *P* is false suffice for you not to accept it?

Consider, in fact, the argument you have just read. No doubt some philosophers will accept it, while others will not. Which ones will accept it? The ones to whom it seems correct, of course. Even if you do not accept it, you still will be thinking in accordance with the rule of phenomenal conservatism. The difference will merely be that to you, it does not seem correct. There is no (rational) escape from the reliance on how things strike *you*. (Even if you decide to rely on the opinions of a reliable expert, you will still be relying on how things seem to you: what the expert seems to you to be saying, and who seems to you to be an honest and reliable source.)

Because of this fact, any attempt to deny the principle of phenomenal conservatism will be self-defeating, for all thought and reasoning presupposes the principle in a certain sense. Even the reasoning put forward by philosophical skeptics, even as they try to present grounds for doubting all appearances, depends for its force on the way things appear to the skeptic and his audience. To illustrate, reconsider the following argument in favor of philosophical skepticism, the infinite regress argument:

Argument A:

1. In order to know anything, I must have an adequate reason for believing it.

2. In order for a reason to be adequate, I must know it to be true.
3. I cannot have an infinite series of reasons.
4. A series of adequate reasons cannot be circular.
5. Therefore, I cannot know anything.¹⁸

This argument, on the face of it, carries some force. Some work will need to be done to respond to it, to show, in the face of the argument, how we can have knowledge. But now compare the following “argument” for the same conclusion:

Argument B:

1. 3 = 5.
2. Therefore, I cannot know anything.

This “argument” carries no force at all, so that it is not even clear whether it deserves to be considered an argument. *This* form of skepticism does not call for a response; there is no philosophical work to be done in explaining how knowledge is possible in the face of it. Similarly, consider the following anti-skeptical “argument”:

Argument C:

1. There are seventeen inhabited planets in the Andromeda galaxy.
2. If there are seventeen inhabited planets in the Andromeda galaxy, then skepticism is false.
3. Therefore, skepticism is false.

Just as argument B does not call for a response, argument C does not count as a response. Just as we do not need to take argument B seriously, the skeptic does not need to take argument C seriously.

Why? What is the difference between (A) and (B)? The difference is not that (B) has a false premise, for (A) has a false premise too (which one, we will see later). Most of the arguments that have been presented throughout the history of philosophy have contained one or more false premises (for philosophical arguments are almost always logically sound, but their conclusions are usually false). Still, they have not been bad in the way that (B) and (C) are. What we are trying to account for is not why arguments (B) and (C) are *unsound*, but why (B) and (C) are not serious arguments at all; why do they not have a place in the discussion of philosophical skepticism?

Notice that in a way, this issue is prior to the question of whether philosophical skepticism is true, for unless we (at least implicitly) have a way of distinguishing what counts as a serious argument from what doesn't, we

cannot even approach the issue, philosophically. This does not mean the skeptic would win dialectically; it means the practice of dialectic could never get started. We would have no conception of what it would be to motivate skepticism, or anything else.

What is wrong with argument (B) is that, not only does it have a *false* premise (like most philosophical arguments), but its premise *doesn't even appear* true; indeed, it seems obviously false. Hence, it doesn't even give us a *prima facie* reason for doubting that I can have knowledge. Similarly, the premises of (C), while not obviously false, do not particularly seem true either. Hence, we do not take them even *prima facie* as grounds for thinking I can have knowledge. This is why neither (B) nor (C) requires philosophical discussion or response.

The above illustrates the way in which the rule of phenomenal conservatism is presupposed in the practice of dialectic: what we count *prima facie* as motivating a position is a matter of what at first seems true. Clearly the intuitively *plausible* (true-seeming) premises of (A) are taken as *prima facie* justified; else the argument would give no motivation at all to skepticism. (Also clearly, but less importantly for our present purposes, they are taken as *only* *prima facie* justified, or else discussion of the issue would terminate once the validity of the argument was agreed upon.) If phenomenal conservatism is false, then dialectic as we practice it is fundamentally and intrinsically irrational. In calling it “intrinsically” irrational, I mean to imply that the problem would be irremediable; there is nothing that could be recognized as a form of dialectic or reasoning that would not be irrational. As a result, it is impossible coherently to argue against phenomenal conservatism. In engaging in argumentation, one is presupposing that there is a distinction between serious arguments and mere strings of arbitrary statements such as (B) and (C). If phenomenal conservatism is false, we cannot draw this distinction in practice. Thus, the skeptic who would argue against phenomenal conservatism labors at pulling the rug out from under his own feet. And if the skeptic eschews argument altogether, we shall hardly be concerned by our failure to “answer” him.

I think, in fact, that the lesson is even more general than I have just indicated—that is, the rule of phenomenal conservatism underlies not only our practice of argumentation (the attempt to convince others through giving expression to one's reasoning), nor, even, just the practice of reasoning. I think the principle of phenomenal conservatism underlies *judgement* in general. I think reflection will reveal that all judgement, whether inferential or not, is a process in which one accepts a proposition on the basis of how things seem to oneself. If phenomenal conservatism is false, so that the way things seem to oneself is irrelevant to epistemic justification, then all judgement must be irrational. And this is something which no philosopher, not even a skeptic, can

accept.

5. Questions and Objections

(1) *Phenomenal conservatism seems to be overly liberal in doling out justification: whenever it seems to us as if P, that affords us prima facie justification for P. But doesn't it seem to us as if P whenever we believe that P? Didn't you, in fact, just assure us that beliefs are always formed on the basis of how things seem to the believer? Therefore, how can a belief ever be irrational according to your theory of epistemic rationality?*

Reply:

There are three general ways that a belief can be irrational. The first is when the subject has grounds for doubting the proposition believed that ought to, but don't, cause him to revise the belief. Remember that phenomenal conservatism is only a principle of *prima facie* justification.

Here is an illustration of this case: To a casual observer, it looks as if the sun moves around the earth. Therefore, other things being equal, it would be reasonable to assume that the sun does in fact move; the heliocentrists have the burden of proof. Nevertheless, if I, right now, were to believe that the sun orbits the earth, my belief would be irrational. This is because I know that the opposite belief has been established scientifically, and I also know a bit about the sort of experiments that establish it. That is to say, I have evidence against the belief that the sun orbits the earth which, epistemically, should overrule that belief. However, I might be dogmatic and refuse to reconsider the belief in the light of that evidence, in which case my belief is unjustified.

Phenomenal conservatism, then, does not imply that all beliefs are justified, full stop. Nevertheless, the question may arise whether it implies that all beliefs are *prima facie* justified, which in itself would seem at least odd, though not as obviously unacceptable as the claim that all beliefs are justified. The answer to this question, however, is no, which brings us to the second way that a belief may be unjustified.

A belief can fail to be even *prima facie* justified, if the believer adopts it for reasons other than that the proposition believed seems to be the case. I said above that all *judgement* involves accepting a proposition on the basis of how things seem to oneself. However, not all *beliefs* are formed on the basis of how things seem to oneself. (The fact that "judgement" and "belief" appear to most philosophers to be synonymous is perhaps a testament to their rationality.) Judgement is a particular process by which a belief can be formed, but there are

other ways of forming beliefs too, which do not depend on how things seem to the subject.

What ways? Well, there are two that I know of: self-deception and faith. (On reflection, the latter may turn out to be a species of the former.) In self-deception, a person adopts (or retains) a belief, not because it seems to be true, but because they *prefer* to believe it. A special case of this is wishful thinking—adopting a belief because one wishes it were true. For instance, say a man's son has been accused of a serious crime. Despite powerful, and almost conclusive, evidence against the son, the father continues to believe his son is innocent. What sort of factors are causing the father's belief? Not attention to how things honestly seem. It is not as if the father has something wrong with his reasoning capacity, such that he can't perceive the link between evidence and conclusion. For we may suppose—also perfectly realistically—that this same father, if any other boy were accused of a similar crime and similar evidence were brought forward, would conclude that the boy was guilty. The evidential situation could be the same; the only difference would be the emotional situation. If the father is "unable" to see how the evidence points to his son's guilt, it is because he does not *wish* to see that; rather, he chooses not to *accept* it. The general point: it often happens that what appears to be the case is something that it would be emotionally painful for a person to believe. People generally try to avoid pain. As a result, people often refuse to accept those propositions. On the other side, it often happens that what does not appear to be the case is something that it would be emotionally pleasant to believe. People generally seek pleasure. Hence, they often choose to accept, to talk themselves into, such beliefs. And those beliefs are epistemically irrational.

Compare the case of religious faith. Suppose a person has been presented with a certain wafer, which for all the world appears to be an ordinary piece of bread. Nevertheless, he believes it is a part of the body of Christ. Why? Because he has been told very solemnly, by an authority he respects, that this belief is essential to his religion. But why does he accept that religion, when it contradicts the evidence of the senses? I am not speaking here of those Catholics who have sophisticated philosophical arguments to defend their religion, but of the average Catholic, who takes his beliefs on faith.

To take a "leap of faith" is precisely to choose to believe something that does not appear to be the case, or even appears false. If I believe that there is a cup on my desk because I am seeing it, that is not called "faith." If I choose to believe that there is a cup on my desk, even though I can't see anything here, *that* is a leap of faith. Faith is typically associated with authority—for example, someone else tells me that there is a cup here even though I don't see anything, and I decide to take their word for it. (No one would just decide on their own that the communion wafer turns into the body of Christ.) This decision could

be prompted by a feeling of respect for the authority figure, a desire to remain part of the community he represents, and even moral beliefs about the virtue of holding certain beliefs and/or the vice of questioning them.

There is a parallel here (as in so many other places) between epistemology and ethics. Unjustified belief is belief which is epistemically blameworthy, just as immoral action is action that is morally blameworthy. True immorality results when a person is aware of what ought to be done, and chooses to act otherwise. It happens because people have other, nonmoral sorts of motivations, including self-interest and various emotional motivations. If a person is entirely unaware that his action will result in some harmful consequences to others, and there is no reasonable way he could be expected to be aware of that, then the occurrence of such harmful consequences does not make his action blameworthy. Blameworthy actions occur only when other motivations override a person's sense of morality. Nor is a person who lacks a sense of morality immoral; he is amoral, just as animals are amoral. Likewise, irrational beliefs occur when other motivations override a person's epistemic goals, the goals of having true beliefs and avoiding false ones.

But there are really two kinds of blameworthy action—those that result from evil intent¹⁹ and those that result from negligence. If a person is unaware of the harmful consequences of his action because he simply chooses not to investigate the issue, when any reasonable degree of caution would have resulted in awareness of them, then he can be blamed for negligence. And this suggests our third and final way that a belief can be unjustified: a belief might be unjustified, despite its seeming to be true, if the believer fails to exercise due caution in accepting it (fails to investigate the issue, to gather evidence). Exactly what is “due” caution is, of course, a matter for debate, and I will not attempt to answer that in detail. My purpose here is only to explain how unjustified beliefs can come about, in the light of my claim that mere appearances are sufficient to generate prima facie justification. Consider an example.

Suppose that Smith has decided to investigate the issue of creationism versus evolution, which he has heard a little about. He goes to a lecture by a creationist and listens to all the arguments the creationist presents. At the end of this lecture, it appears to Smith that the theory of evolution is unreasonable and counterintuitive. Smith accepts creationism, and that is the end of his investigation. Notice that this case need not involve any self-deception or leaps of faith in Smith's weighing of the arguments; the creationist lecturer might have presented arguments that honestly seemed correct when Smith heard them. Nevertheless, I would consider Smith's belief to be unjustified, or *insufficiently* justified, due to Smith's failure to investigate the other side of the issue. This is epistemic negligence. In order for Smith to be justified in accepting creationism, he would have to also attend a lecture (or read a book, or some

such) by an advocate of the theory of evolution, to see whether they could make evolution seem reasonable.

This is a fairly uncontroversial point—that rational belief requires a reasonable attempt at investigating the issue. But is it consistent with phenomenal conservatism? Phenomenal conservatism would seem to imply that creationism is justified for Smith, since it seems true to him, and—due to failure to listen to what the evolutionists have to say—Smith has no evidence against it.

I think, however, that Smith really does have grounds for doubting creationism. In general, there are at least two different kinds of “grounds for doubt” that can defeat a prima facie justified belief. One is evidence against the proposition believed—that is, if *P* is what is prima facie justified, evidence that supported $\sim P$ would obviously be grounds for doubting *P*. But another kind of defeater would be evidence that the means by which one formed the belief is unreliable.²⁰ (Notice that this isn't the same thing—even if the means by which I came to believe *P* are unreliable, that doesn't tend to justify $\sim P$.) Smith knows that there is a controversy between evolutionists and creationists, and he knows that he has not investigated the other side of the controversy. These facts are not evidence for the theory of evolution, but they are sufficient evidence that Smith formed his belief in creationism in an unreliable manner. Note that the situation would be different if Smith were unaware that there was any controversy about the issue, and if he thought the creationist was presenting the received view.

To sum up: a belief may be unjustified if (a) one forms it for reasons other than how things seem (e.g., self-deception), (b) one has evidence against it that one chooses to ignore, or (c) one has reasons for thinking one's belief-forming method was unreliable (including that one was negligent in the investigation of the issue).

(2)

Richard Foley has presented a compelling argument against a thesis called “epistemic conservatism,” which is very similar to phenomenal conservatism.²¹ Epistemic conservatism holds that, if a person believes that *P*, the mere fact that they believe it produces at least some degree of prima facie justification for *P*.

Foley imagines a case in which *S* has almost, but not quite, enough evidence for *P* to make it rational for him to believe it. Suppose *S*, irrationally, accepts *P* anyway. As soon as he formed this irrational belief, if epistemic conservatism is true, it would immediately become rational—for now, in addition to the evidence *S* originally had for *P*, there is also the fact that he believes it, and this pushes *P* over the threshold for rational belief. This result

is counterintuitive, and so we should reject epistemic conservatism.

*How does phenomenal conservatism relate to epistemic conservatism?
Does Foley's argument apply to phenomenal conservatism as well?*

Reply:

Foley's argument succeeds in refuting epistemic conservatism as he defines it, but it does not apply to phenomenal conservatism.

Matters will be somewhat clearer if we redefine "epistemic conservatism." Let's say that an epistemic conservative is a person who believes there is some sort of representational mental state such that its mere occurrence makes one at least to some degree prima facie justified in believing its content. Then we can distinguish two different forms of epistemic conservatism: A *doxastic* conservative holds that the mere occurrence of a *belief* provides some form of justification for that belief, so all beliefs have at least some prima facie justification automatically.²² ("Doxastic" means pertaining to beliefs.) A phenomenal conservative, on the other hand, holds that the *appearance* that *P* provides some kind of justification for believing that *P*. Foley's argument, I think, only refutes doxastic conservatism.

Suppose we tried to modify Foley's case to make it applicable to phenomenal conservatism. The only way to do this, while keeping the structure of the case the same, is to imagine that at first, *S* has almost enough evidence to make it rational to believe *P*, and then on top of that, there occurs a state of its seeming to him as if *P*—for instance, *S* might have a perceptual experience or an intuition representing that *P*. This appearance state substitutes for the belief in the original version of the story. (But the appearance is neither rational nor irrational, unlike the belief.) Is *S* now justified in believing that *P*?

It seems to me that he is. When the example is changed in this way, I no longer have the intuition that it would be irrational to believe *P*. If you still do, consider the example I discuss under objection (3), which is a similar case.

(3) *Phenomenal conservatism would seem to imply not only, as any foundationalist holds, that it is possible for a belief to be justified when one has no argument for it, but also that a belief can be justified even when one believes it on the basis of a bad argument. For it might happen that a person gives a fallacious argument for P, that they do not notice the fallacy, and that as a result, it seems to them that P. In this case, according to PC, a foundational justification will arise in lieu of the failed inferential justification.*

Reply:

Phenomenal conservatism does indeed imply that a belief based on a

fallacious inference can be justified. This result, however, is not a problem for the theory, because the result is correct and not counterintuitive.

Consider the case of the Unfortunate Mathematician. A certain mathematician, let's say, has just constructed a proof of *P*. His argument is valid and sound, the premises are self-evident, and he exercised all due care in constructing it (he checked it over carefully, etc.), so he both believes and is justified in believing that *P*. Now he goes on to construct an argument for another mathematical proposition, *Q*. This second argument appears to him just as sound as the first. Furthermore, he has no reasons to doubt *Q* and no reasons to suspect his method of forming the belief that *Q* to be unreliable. He has exercised the same due caution in checking over this second argument as the first, and he found no errors in it. Each step appeared perfectly correct as he checked it, just as in the first argument. Unfortunately, however, the argument contains a subtle error in step sixteen, rendering it invalid. Since he doesn't know about this error, he continues to accept *Q*. Is his belief rational?

I think yes. Remember that we are talking about internalist justification, and keep in mind the alternatives open to the mathematician. In this situation, he must either accept *Q*, deny *Q* (which amounts to accepting $\sim Q$), or suspend judgement. But surely it would not be rational for him to deny *Q*. He may have a fallacious argument for *Q*, but he has no argument at all against *Q*. Should he suspend judgement? Well, he already accepted *P*, and we agreed that that was rational. As far as he can tell, the case for *Q* is exactly as strong as the case for *P*; therefore, how could he be expected to refrain from believing *Q*? If he accepts *P* but declines to accept *Q*, it would surely be appropriate for someone to ask him why. And what reason would he be in a position to give? Given his epistemic position, it is as rational to believe *Q* as it is to believe *P*, and since *P* is justified, we must conclude that *Q* is justified as well.

Obviously, if he becomes aware of the error in the "proof," or even comes to have reason to suspect there is an error, then his belief is no longer justified. But this is also true in the case of the valid proof of *P*—if he comes to have reason to suspect there is an error in it, he will also cease to be justified in believing *P*.

(4)

*What if it seems to one person as if P, but it seems to another person as if $\sim P$?
What should we believe then?*

Reply:

Each person should believe what seems to himself to be correct after a careful examination of the issue. If it seems to John that *P* and it seems to Sally that $\sim P$, then John will believe *P*, and Sally will believe $\sim P$. If you are a third

party observer, then you should believe whatever seems to you to be the case.

This brings out the point, obvious in any case, that internalist justification is relative to an individual. One person may be justified in believing that *P*, while another person is not. Since different people have different experiences, interests, cognitive abilities, and so on, it is not realistic to think that what propositions are justified should be the same for everyone.

(5)

But is there any way for John and Sally to resolve their disagreement?

Reply:

There are two ways: either John makes it seem to Sally that *P* is true, or Sally makes it seem to John that *P* is false. No one has ever resolved a disagreement, except by one of those methods. (By "resolving" a disagreement, I mean convincing the other person. Hitting them over the head or shooting them doesn't count.) When you give someone an argument, what you're trying to do is simply to change the way things appear to them. You do this by drawing their attention to some other things that seem true to them, which support your conclusion. That is all you can possibly do. Of course, if they don't find your premises plausible, then your argument won't work; it won't persuade them. Disagreement is often unfortunate, but it is also a fact of human life. It is not the responsibility of a theory of knowledge to produce a technology for inducing universal agreement.

(6)

Does phenomenal conservatism imply that "truth is relative"?

Reply:

No. It implies that *justification* is relative to an individual. Justification is not the same as truth. You can be justified in believing something that, in reality, is false, and you can be unjustified in believing something that is true. This is where it is important to see, again, that phenomenal conservatism is not a metaphysical principle but an epistemological norm. It does not hold that things must be as they appear (which would indeed imply that reality is subjective, with multiple different "truths"), but rather that it is reasonable, other things being equal, to believe things are as they appear.

(7)

If justification is separate from and does not imply truth, then why should we care about justification? Why not just focus on truth?

Reply:

Although believing only what is justified does not guarantee you of getting only true beliefs (nothing within your power will), believing only what is justified for you is *doing the best you can* to get to the truth.

Here is an analogy. Suppose you want to shoot some baskets. The best thing to do is to carefully aim the ball at the basket, and then throw it. Will this guarantee you of making a basket? No, but it's the best you can do. It'll do a lot better (in all probability) than throwing the ball without aiming, or not throwing the ball at all.

When trying this technique, some people will have much better results than others. Some people will make a lot of baskets; others will make relatively few. The latter group of people are still doing the best they can: all they can do is aim the ball and throw it, and they cannot use someone else's aim to do it.²³

Similarly, when you believe the things that are for you justified, this won't guarantee that you'll never be wrong. But you'll be doing the best you can. You'll do a lot better (in all probability) than if you tried believing the things that seem false.

This is what epistemic justification is about: does one believe "responsibly," so to speak, when one accepts *P*? The task for an epistemologist, in my view, is to address this question with respect to the main kinds of beliefs that we hold. If accepting *P* is part of doing the best one can to pursue truth and eschew error, then accepting *P* is justified in the relevant sense; no more could be asked.

Notes

1. This is not to say that all epistemologists accept that the first three conditions I have stated are necessary for knowledge, but all epistemologists agree that they are *not sufficient*. See Gettier.
2. Russell's example, *Human Knowledge*, 154.
3. There are further refinements required to the definition, having to do with distinguishing "misleading defeaters" from "genuine defeaters," but these details do not concern us here. See Klein, "A Proposed Definition of Propositional Knowledge," and Lehrer and Paxson.
4. Admittedly, even this description is oversimplified. It would be hard to identify the background beliefs of mine that lead me to categorize this object as a tomato. "All red, round things are tomatoes" won't do, nor do I have concepts adequate for describing the exact sort of shape and color that tomatoes have (unless you trivially count "shaped like a tomato" and "the color of a tomato"). What I seem to have is a much more diffuse mental state than a belief or collection of beliefs; roughly, I have a *conception* of tomatoes, which makes possible the classification of objects as tomatoes.

(or nontomatoes as the case may be) on the basis of observed features, although the conception also involves features tomatoes are supposed to have that I am not now observing (for example, the tomato should have seeds inside it). So my belief in a tomato will be based on this conception of tomatoes, plus my present visual experience.

5. Evan Fales, ch. 5-6, defends this view of the relation between perceptual experiences and beliefs. Fales also discusses and answers traditional objections to the notion of noninferential, perceptual knowledge. See my review of his book for a summary.

6. See Nozick, 172ff.; Dretske, "The Pragmatic Dimension of Knowledge"; Goldman; and Lewis.

7. Ellis, *Reason and Emotion in Psychotherapy*. See also papers 2 and 3 in Ellis and Grieger for a review of the empirical evidence supporting rational emotive therapy.

8. There is one sort of exception, which occurs in the viewing of ambiguous drawings such as the Necker Cube and several of the illustrations in Jastrow, 283-95. In these cases, an observer can often choose to see the picture in one way or another; for instance, in viewing the Necker cube, one can choose to view a particular square alternately as the front face or the back face of the cube depicted, so that the content of one's visual experience is, in that sense, dependent on the will. This calls into question my proposed explanation in the text for why experiences are neither justified nor unjustified, because in such a case I still do not think it would make sense to characterize one's experience as justified or unjustified. It would not make sense, for instance, to speak of one's being justified in seeing the Necker cube in a certain way.

9. That is, I think the principle accounts, for every foundational belief, for why that belief is justified. Additional principles, or a stronger principle, would be needed to account for the *degree* of justification each such belief has.

10. Chisholm has observed that "appear" words are sometimes used to express tentative judgements, though I do not think that this refutes my contention here. See chapter IV, note 39.

11. See section IV.4.2-3.

12. The concept of a memory appearance is related to, but distinct from, Shoemaker and Parfit's concept of a "quasi-memory." Quasi-memories are states similar to experiential memories (as distinct from factual memory), except that it is conceptually possible to quasi-remember someone else's experiences (but it is not conceptually possible to quasi-remember things that never happened at all) (Shoemaker, "Persons and their Pasts," 271; Parfit, 14-15). "Memory appearances," on the other hand, are states like memories (experiential or factual), except that it is possible to have a memory appearance of something that is not true or that never happened.

13. A number of foundationalists, probably most foundationalists these days, rely on the notion of *prima facie* justification. See for example Audi, *The Structure of Justification*, 307-10; Pollock, *Contemporary Theories of Knowledge*, 177-78. For the argument that there is nothing inconsistent about this, see Alston, "Has Foundationalism Been Refuted?"

14. Alston makes this point, along with several other important clarificatory ones, in "Has Foundationalism Been Refuted?"

15. There are at least two inferential steps here: first, from the fact that you spoke certain words to the fact that you see, or at least think you see, a red sphere there, and second, from the fact that you see or think you see a red sphere there to the conclusion that there probably is a red sphere there. According to Burge (though I think he is wrong), I would actually have foundational justification for believing that you see the red sphere; however, the second inferential step would still be necessary, even on his theory. Moreover, there is no difficulty in generating any number of examples in which I would have a further, inferential justification for thinking there is a red sphere there.

16. See section II.6.

17. Foley, *Working without a Net*, 8. I do not know, however, whether Foley would agree with the use to which I have put his theory.

18. This is an abbreviated version of the argument from section II.1.

19. To be more precise, the distinction should be made between doing evil knowingly, and doing it unknowingly as a result of negligence. Doing something *knowingly* is not exactly the same as doing it *intentionally*. If x is neither my end nor a means to my end, but I know that it will occur if I do y , and I proceed to do y , then I bring about x knowingly but not intentionally. The law of double effect notwithstanding, a person may be blamed for bringing about evils knowingly, whether intentionally or not.

20. Compare Pollock's distinction between undercutting and rebutting defeaters in *Contemporary Theories of Knowledge*, 38-39. The distinction, with different terminology, originates in Pollock, "The Structure of Epistemic Justification."

21. Foley, "Epistemic Conservatism."

22. Chisholm takes this position in *The Foundations of Knowing*, 14.

23. I owe this analogy to Ben Kovitz (in email).