1

The Structure of Justification

Alvin I. Goldman

1.1 THE CONCEPTS AND QUESTIONS OF EPISTEMOLOGY

Epistemology is the study of knowledge and related phenomena such as thought, reasoning, and the pursuit of understanding. It is less a study of customary thinking processes—although they are relevant—than a study of better versus worse ways to think, reason, and form opinions. Moral theory reflects on what is right and wrong in the sphere of action, while epistemology reflects on what is rational or irrational, justified or unjustified, in the sphere of the intellect.

Why are matters of the intellect important? This can be approached from many vantage points of ordinary life. Do I want to make good decisions in life, decisions that promote my own welfare, my family's welfare, my community's welfare? If so, I had better figure out which of the available choices would best promote favorable outcomes. In other words, I need to form correct, or *true*, beliefs about the consequences that would ensue from the performance of different actions. If I form true beliefs about the consequences of each choice, I am more likely to make choices that lead me in useful directions. If I form false beliefs, my choices may be unfortunate however good my intentions. Accurate beliefs tend to guide us down desired pathways, inaccurate beliefs down pathways we don't mean to travel. A student decides to get trained in a given field because employment opportunities are predicted to boom in that field within a few years. Training in this field should bode well for the student if the prediction is true, but a wasted effort if it is false. One hopes to act on true predictions rather than false ones.

The same point emerges when deciding whom to trust in this or that domain. In consulting a physician about an ailment, I want her to have extensive medical knowledge plus the skill of applying that knowledge to new cases. I want her to predict correctly which treatments would cure or alleviate my ailment (and have no serious side effects). The same point extends to

a choice of a financial advisor or auto mechanic. The claim here is not that instrumental or practical value is the only value of true belief or the sole ground of the intellect's importance. People have intellectual interests not rooted in practical affairs. We are curious creatures. We want to know, for example, what caused the dinosaur extinction many millennia ago, even if knowing this has no immediate action implications. This curiosity does not rest on the pursuit of practical ends.

All right, you concede, we have reasons to try to get truths, and to enlist the help of others who are adept at getting truths. But how is truth to be acquired? Can one ascertain the truth by just reaching out and grasping some facts? It isn't clear how one does that. This problem is what epistemology is largely about. A customary way to approach the issue is to focus on another central topic in epistemology: *justification*. If I can get a justified or warranted belief, rather than a randomly or haphazardly chosen one, such a proposition is more likely to be true. How, then, do I go about getting justified beliefs? Some epistemologists link justification to having good evidence. Others link it to following good methods or procedures. These are among the prospects to be explored in what follows. In this chapter and Chapters 2 and 6 our discussion focuses on justification; in Chapters 3, 4, and 5 it focuses on knowledge.

As these introductory remarks make clear, epistemology talks a lot about belief. What is belief (or opinion)? It is a species of psychological attitude toward a proposition, where a proposition (roughly) is a content that purports to express a fact. The statement that the dinosaur extinction was caused by a massive asteroid impact on the earth purports to state a fact. To believe this proposition is to mentally "assent" to it, or think that what it states is so. Belief belongs to a family of psychological attitudes directed at propositions. Other members of this family include *disbelief*, which is the mental attitude of rejecting, or denying, a proposition, and *agnosticism*, or suspension of judgment, which is the mental attitude of neutrality, or indecision, with respect to a proposition's truth. Collectively, these different attitudes (plus more finely graded attitudes such as being 75 percent convinced of something) are called *doxastic attitudes*. The term *doxastic* comes from the Greek word *doxa*, meaning opinions.

It is widely thought that forming a belief that is reasonable, warranted, or justified is the best means available of forming a true belief. Hence, epistemology is particularly interested in the question of how one should go about acquiring a belief that is warranted, justified, or reasonable. *Justification* and *warrant* are examples of terms of epistemic evaluation or appraisal. To call a person's belief "justified" is to commend it, or appraise it positively, along some evaluative or normative dimension, whereas calling a belief "unjustified" or "unwarranted" is to criticize it, appraise it negatively, along the same dimension. Suppose I say that John thinks that Gregory is a terrible fellow, but this belief is totally unwarranted because it is wholly based on Gregory's appearance. I call it "unwarranted" because judging people by their looks is like judging a book by its cover. It is a poor basis for judgment (or belief). Such a normative assessment has parallels with moral discourse, in which actions are called right or wrong. Epistemologists agree, however,

that epistemic terms of appraisal like *justified* and *unjustified*, *warranted* and *unwarranted*, are not terms of moral evaluation; they express evaluations along some intellectual dimension. If you believe in the absence of good evidence, or by reliance on unsound reasons, this is an intellectual failing rather than a moral one.

Although justification may be significantly related to truth, truth and justification are not equivalent concepts. A proposition can be true although nobody believes it, and it can be true although nobody is justified in believing it. Consider a precisely delimited expanse of seaside beach. Some proposition of the following form is true: "The number of grains of sand on this beach is N." But nobody is justified in believing this truth (with the correct value of N filled in) because nobody has determined or ascertained the correct value of N. A major question in epistemology is how justification is related to truth, but they are not equivalent concepts.

Truth, it may be said, is a purely metaphysical concept rather than an epistemological one. Given a proposition, what makes it true or false is simply the state of the world. Its truth-value is not affected by cognitive relations people have toward the relevant state of affairs. But cognitive relations to a proposition are precisely what are crucial in determining justification or warrant. A person's justifiedness with respect to a proposition P is never (or rarely) fixed by P's actual truth-value. Despite a proposition's truth, it is possible for someone to lack any evidence for its truth (as in the grains-of-sand example). Conversely, it is possible to have highly favorable (though misleading) evidence that justifies one in believing a proposition despite its falsity. Thus, truth and justification must be carefully distinguished.

Justified and *warranted* are not the only terms of epistemic evaluation; another familiar term of epistemic evaluation is *rational*. Some epistemologists equate rationality with justification, but we shall keep them distinct. Epistemologists who favor rationality-talk as compared with justification-talk often opt for a more finely delineated range of doxastic attitudes. Instead of the tripartite classification scheme of belief, disbelief, and withholding (or suspension), they prefer degrees of belief arrayed along a zero-to-one interval. One represents the highest possible strength of credence in a proposition, and zero represents the lowest. Our own treatment here will mainly use the tripartite classification scheme. But in Chapter 11, on probabilistic epistemology, finer gradations will often be invoked.

In addition to justification, warrant, rationality, and reasonability, a critical concept in epistemology is *knowledge*. Indeed, the term *epistemology* just means the study of knowledge (*episteme*, in Greek). According to many theories, there are intimate relations between knowledge and some of the other epistemic concepts we have introduced. Here are some points of widespread agreement. First, knowledge implies truth; you cannot know that P unless P is true. (This idea is often conveyed by saying that truth is *factive*.) If you are certain of a proposition P, in your own mind, you will be inclined to *claim* to know it. But if P is not true, as a matter of fact, then you don't really know it. (One cannot know what isn't so.) Second, a person must believe P, or have reasonably high

credence in P, in order to know it. Knowledge is (partly) a psychological state. Third, most theories hold that knowledge requires justification. You do not know that P by simply believing P where P is true. Being justified in believing it is also necessary. Thus, there appears to be a web of relationships between knowledge and other important concepts. Closer examination of these matters occupies Chapters 3, 4, and 5. This book begins with justification, which some regard as the more fundamental epistemological concept.

1.2 THE EPISTEMIC REGRESS PROBLEM

Many justified beliefs attain this status when people infer them from other things they believe. Inferential relations between beliefs are often expressed in conversation or debate with others. Suppose Henry says to Tony, "I hear you believe that New York will win the basketball championship next year. How can you think that? New York was so miserable last season; how can you think it will turn things around and win next year?" Tony responds: "New York landed the top new draft pick in the league this year, and he is rumored to be the most talented player since Michael Jordan. It has also acquired one of the best big men in the league. With these crucial additions, I feel confident it will win."

Tony offers an argument featuring P1 and P2 as premises and C as conclusion:

P1: New York will have the most talented draft pick in years.
P2: New York also has one of the best big men in the league.
C: Therefore, New York will win the championship next year.

Tony defends C by appeal to P1 and P2, implying that C is a reasonable conclusion to draw from P1 and P2 (given various unmentioned assumptions). C does not deductively follow from P1 and P2, that's clear. But it may be a reasonable nondeductive inference from P1 and P2 (given the additional assumptions).

Our story presents Tony as a speaker who verbally defends his belief to a challenger, but a similar story might be told where there is no challenger and Tony makes no verbal defense of C. If he holds his belief in C in utter silence, but it is (properly) based on (justified) beliefs in P1, P2, and the additional assumptions, the same justificational upshot will hold with respect to proposition C. Thus, interpersonal speech is not required for there to be justified inferential belief. Personal justification has no essential tie to a "dialectical" situation, where verbal reasons-giving takes place.

What is the structure of Tony's justification for C? As described so far, it resembles the roots of a tree. In Figure 1.1, the belief in proposition C (labeled Bel(C)) appears as the top node of a tree, with beliefs in premises P1 and P2 (labeled Bel(P1) and Bel(P2)) shown as downward-branching roots. The arrows indicate that Bel(P1) and Bel(P2) lend joint support to Bel(C). Inferential justification is supposed to be transmitted upward from one or more root

nodes to at least one higher-level node. However, if the lower-level nodes have no justification of their own to transmit—in other words, if Tony does not believe them *justifiedly*—then the higher-level node(s) cannot inherit justification from them.

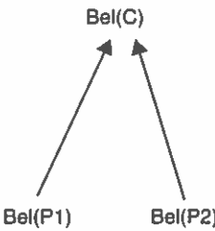

FIGURE 1.1 Belief in C is Justified by Inference from Justified Premises P1 and P2.

Might the root structure of Tony's justification run deeper? If his beliefs in premises P1 and P2 are justified, how was *their* justification acquired? Perhaps Tony had other justified beliefs from which P1 and P2 were properly inferred. His belief in P1, for example, may have been based on his justified beliefs about the new draft pick's previous scoring records and the tall recruit's rebounding prowess. To depict this situation in the diagram, we would expand the root structure by adding another level of branching roots, roots that extend downward from Bel(P1) and from Bel(P2). Of course, each additional root node must also be justified if Tony's beliefs in P1 and P2 are to acquire justification from them.

What emerges here might be called a *regress* of reasons, grounds, or justification. How does this regress of justification continue? This question poses what is called the *regress problem* in epistemology. It was posed in ancient times by Sextus Empiricus, who asked about the chain of "proofs" for a person's assertion. Here we focus on beliefs rather than assertions and reasons rather than proofs, but the core idea—and worry—is much the same. Can the regress of reasons continue indefinitely, with additional beliefs being invoked at each step, with no repetition and no end points? Or must the regress eventually terminate along each root, where a terminus is a justified belief that acquires its justificational status from an epistemic source distinct from inference? What are the possible structures here, and which possibility captures the real structure of inferential justification?

Here are three views about the correct structure of a regress—that is, a structure that enables justified beliefs to be derived from other (justified) beliefs:

1. *Infinitism.* The correct structure is an unending continuation of reasons, without repetition or end. A justification tree for each inferentially justified belief contains roots that never end and never contain repetitive nodes.
2. *Foundationalism.* The correct structure is a tree with roots of *finite* length. Each root has a terminal, or final, node that represents a justified but

noninferred belief. Such beliefs are *foundational*, or *basic*, beliefs. All justified beliefs that are not themselves basic, however, ultimately derive their justification from one or more other beliefs that are basic. So, all justification ends at—or begins with—basically justified beliefs.

3. *Coherentism.* The correct structure is one in which some of the roots circle, or loop, back on themselves. In other words, some nodes reappear earlier in the chain. Thus, the regress has no end points. At the same time, it is not infinitely long. In rejecting end points, coherentism joins infinitism in rejecting the notion that justification depends on there being foundational, or basic, beliefs.

A fourth response to the regress problem belongs in a separate category because it does not purport to "solve" the problem of justification—that is, to show how inferential justification is possible (and feasible):

4. *Skepticism About the Regress.* None of the first three solutions to the regress problem is satisfactory. One cannot get justified belief via any of the structures they describe. Because those other three solutions jointly exhaust the possibilities for a positive solution, there is no way by which inferential beliefs can be justified.

In claiming that none of the three solutions works (for reasons to be examined below), the skeptical position implies that justification can never be derived by inference. This response is *skeptical* because it denies justification for a huge swath of beliefs—all beliefs based on reasoning—that we normally consider justified.

Before examining these rival responses to the regress problem, let us note a shared feature of the positive solutions. All of these theories restrict beliefs that do the justifying to ones held *at the same time* as the target belief. In the case of Tony's basketball championship belief, our question is whether his belief at the time of the conversation is justified. Obviously, Tony's justification for C depends on whether he is justified in believing that the new draft pick is as talented as he claims. This is something he might be justified in believing in, say, December, although he wasn't justified in believing it in October.

The three positive theories under discussion are all "classical" theories of justification in the sense that they center on *synchronic* reasons. A synchronic reason for a belief is a factor occurring at the same time as the target belief. Earlier or later beliefs make no justificational difference in the traditional view. Presumably everybody would agree that Tony's belief that New York will win the championship cannot be justified in the way indicated if Tony's beliefs about the new players either weren't really in place at the time he believed C or weren't yet justified at that time, say, because the draft of the super-talented player had not yet occurred or the player had not yet signed a contract. Surely the supporting premise beliefs must be held and justified at the time of the conclusion belief; so the synchronic view would have it. We can agree that justified premise beliefs that only occur *after* a target belief

cannot help the latter belief (at the earlier time). However, this leaves open the possibility that earlier beliefs and their justificational statuses can influence the justificational status of a later belief (albeit indirectly). That is a prospect we shall encounter in connection with a diachronic theory of justification to be considered in Chapter 2. Diachronic theories as such do not solve the regress problem, certainly not simply by virtue of being diachronic. The point is only to indicate that theories of justification need not be purely synchronic. All classical theories of justification, however, have been synchronic, including infinitism, foundationalism, and coherentism.

1.3 INFINITISM

Infinitism is rarely given serious consideration. Two problems jump out on first encounter. First, if a justification tree goes back infinitely far (on one or more roots), there seems to be no point at which justification originates. But if there are no originating places for justification—places where justification gets "kick-started"—how can there be any justification to transmit to higher nodes? Second, if a tree contains an infinite chain of reasons, doesn't this imply infinitely many beliefs? But, surely, no human person has infinitely many beliefs.

Does infinitism have any good response to these problems? In response to the first problem, it may be argued that justification need not originate anywhere; starting points aren't needed. Just as time and the universe may each be eternal (in both directions), with no real starting points, so justification can lack starting points. In the case of synchronic justification, of course, the issue of temporal origination doesn't arise. What the critic of infinitism denies is that all beliefs on inferential trees are epistemically dependent on others. The critic insists that some beliefs must be independently justified if they are to generate any justificational "juice" to transmit to the others. This is not a proof, but it seems intuitively compelling.

What of the criticism that nobody possesses infinitely many beliefs? Here the infinitist has a bit of maneuvering room (Klein 2005). The principal maneuver appeals to a standard epistemological distinction between two kinds of justification. *Doxastic* justification is a property of existing beliefs; *propositional* justification arises from the state or condition of an epistemic agent that entitles them to believe a proposition (even if they don't opt to believe it). Proposition P may be justified "for" such an agent in the sense that they *would* be justified in believing it in virtue of their condition or situation. If you are shivering with cold on a wintry night, you are justified in believing that you feel cold, even if you fail to notice this feeling and don't assent to the proposition. Once it is conceded that propositions can be justified even without belief in them, infinitism can be interpreted as the view that there are infinitely many propositions (rather than beliefs) that form chains of inferential justification. Even if no finite being like us has infinitely many beliefs, this leaves open the possibility that there are infinitely many propositions we are justified in believing at a given time.

The propositions might be arrayed in a chain of (potential) inferential dependence. Not that we actually execute all of the corresponding inferences, but we would be justified in executing any finite segment of them.

Does a switch from an infinite chain of beliefs to an infinite chain of propositions really work? What the infinitist presumably holds is not only that a proposition can be justified without itself being believed but also that it can receive (inferential) justification from other propositions without the latter being believed. So infinitism supports the viability of a chain of propositions, each of which bears (or possesses) justification but none of which is believed. This is extremely dubious: Unless a proposition is believed, it is doubtful that it can convey inferential justifiedness to other propositions. Don't *some* propositions in the chain have to be believed? Perhaps so, the infinitist might reply, but this only requires one proposition to be believed, not infinitely many of them. Is this correct? If a single believed proposition is inferentially justified, doesn't it have to be inferred from other beliefs earlier in the chain? But then we are off and running on the same sort of infinite regress of beliefs the infinitist sought to avoid.

Of course, we have already encountered an example in which a proposition is deemed (propositionally) justified without receiving that justificational status from a belief. That's the feeling-cold example of two paragraphs back. The problem is that no infinitist can embrace this sort of case as prototypical, because it features a proposition that acquires justification from a nondoxastic source, and hence is an example of basic, or foundational, justification. This supports not infinitism but one of the rival solutions instead (foundationalism).

1.4 COHERENTISM

We proceed to the coherentist solution to the regress problem. Coherentism is a minority view today but historically was an influential theory and doubtless retains some of that luster. In terms of the regress problem, it is commonly conceptualized as tolerating circular inference. But this may not be the best way to introduce it; let us consider another perspective.

Coherentism depicts a body of justified beliefs as a holistic system whose parts mutually support one another. There are no "privileged" beliefs that play the role accorded to foundational beliefs under foundationalism (i.e., beliefs that obtain justification from outside the doxastic system and transmit this justificational juice to the remaining elements of the system). Under coherentism, all justification arises from interrelations among beliefs (or other doxastic states). Moreover, the justificational interrelations are reciprocal or bidirectional rather than unidirectional.

A good metaphor for coherentism is a house of mutually supporting cards leaning against one another. Such a structure stands upright because its components are each supported by all of the others, or at least many of the others. Support runs in many directions. In a four-card house consisting of A, B, C, and D, support might run from A, B, and C to D, from B, C, and D to A,

and so forth. There are no "foundational" cards that support others but are themselves wholly unsupported. The very formulation of the regress problem assumes that justificational support is linear—that is, runs in a single direction. Coherentism challenges this assumption.

Now consider how justificational coherence might be spelled out, metaphors aside. Unfortunately, there is little consensus here. One ingredient that presumably contributes to system coherence is logical consistency. Unless the beliefs in a total doxastic system are mutually consistent (i.e., their conjunction does not imply a contradiction), the system isn't coherent. But this is an extremely weak condition. Thus, consider the following set of propositions as someone's total system:

1. There are cats.
2. There are bananas.
3. Sometimes it rains.
4. There is a sun.
5. There is a moon.

These five propositions constitute a consistent system, because their conjunction entails no contradiction. But is this consistent system a highly coherent one? No; at least not in the rich sense of coherence that coherentists typically advance. The system fails to be coherent because (i) no proposition is deducible from any of the remainder (either singly or in conjunction with one another) and (ii) none of the propositions would have its probability raised by taking the truth of the others as given. In short, the five propositions are mutually independent, both logically and probabilistically. Coherence, however, is supposed to be the opposite of this; a coherent system should display a high degree of mutual interdependence.

Here is an illustration of a system with greater interdependence:

6. Edna loves cappuccinos.
7. On Tuesday afternoon Edna had a yen for a cappuccino.
8. On Tuesday afternoon Edna thought that the closest place to get a good cappuccino was Café Nero.
9. On Tuesday afternoon Edna went to Café Nero for a cappuccino.

Like system (1)–(5), this system of propositions is logically consistent. In addition, however, it also features several probabilistic support relations among various members of the system. The truth of (6) raises the probability of (7); the truth of the conjunction of (7) and (8) considerably raises the probability of (9); and arguably (9) raises the probability of (7). So there is a fair amount of mutual coherence among these propositions, and hence the system as a whole has a fair degree of systemic coherence (especially for a small system).

Now comes the jackpot question: Does a system's level of coherence guarantee or contribute materially toward a positive justificational status for the propositions that constitute the system? Does the fact that a belief system is

highly coherent guarantee, or even make it likely, that each of its members (or most of them) is justified? No.

Notice first that the propositions constituting a consistent system need not be true, and similarly for the members of a coherent system. System (6)–(9) is more coherent than system (1)–(5), but this is no clue as to which system has more true members. Now turn to justification. Would the fact that somebody believes all and only the members of (1)–(5) and another person believes all and only the members of (6)–(9) provide any clue about which believer has a larger proportion of justified beliefs? No. It all depends on how the beliefs were formed, or arrived at, by the two believers, and nothing said thus far speaks to their respective processes or methods of belief formation. The believer in propositions (1)–(5) may have formed each of her beliefs by careful observation, whereas the believer in propositions (6)–(9) may have formed his beliefs by mere wishful thinking. Thus, the justificational quality of the member propositions does not correlate with system coherence. Contrary to coherentism, nothing follows about a belief's justifiedness—certainly not its doxastic justifiedness—from the fact that it belongs to a system that ranks high on a systemic-coherence scale. It may also fail to be propositionally justified if the relation of coherence among the propositions is too complex or obscure for the subject to appreciate it.

This point may be reinforced by considering the following possibility. Although we may believe each member of a coherent system, we might not recognize or (intellectually) appreciate the fact that they cohere. We might form each belief separately from the others, without recognizing their logical and/or probabilistic relations. In this scenario, the fact that our beliefs are mutually coherent contributes nothing to their level of justifiedness. Their justifiedness depends on whatever leads (or causes) us to believe them. Since, by hypothesis, we are not influenced by their coherence relationships, those relationships have no impact on their justifiedness.

We can highlight these points by tweaking an example often used against coherentism. Fiona loves to fantasize. She sits in her armchair weaving complex stories that rival the best detective mysteries for intricacy of plot and depth of detail. Fiona's fantasy is, in effect, itself a well-composed novel with a high degree of coherence (the plot "hangs together"). Moreover, the fantasy is so vivid and realistic that Fiona believes each of the elements in the story. Is she justified in believing them (to be true)? Surely not. Yet each such element is a member of a highly coherent system. The system would rate very high on the coherence dimension, judged, at least, by their internal relations with one another.

Here are two additional problems for coherentism, the first commonly called the *isolation* objection. According to coherentism, a given belief's justificational status is exclusively a function of the believer's other belief states. Psychological states of other kinds, such as perceptions and feelings, are irrelevant. According to coherentism, your feeling cold right now has no bearing on any of your beliefs' justificational statuses, even the status of a belief to the effect that you feel cold or feel hot. How you actually feel makes no difference, because justification is exclusively a function of *inferential*

relations between beliefs, and neither feeling cold nor feeling hot is a belief. Similarly, if we hear what sounds like a motorcycle careening down the street, hearing this sound is justificationally irrelevant to any motorcycle beliefs we may have. This approach seems wrong-headed. Beliefs that are misaligned with concurrent experience (especially vivid experience) are unjustified. After all, perceptual experience is input from the world, and ignoring such experience does not conduce to sound belief. Some epistemologists (e.g., BonJour 1985) try to tinker with coherentism to avoid this counterintuitive feature, but none of them has had much success.

Finally, a seemingly small defect in coherentism shows how dramatically off-track it is. Coherentism says that a belief is justified if and only if it belongs to (is a member of) a highly coherent total system of beliefs. An obvious consequence of this is that all members of a highly coherent system are justified and all members of a weakly coherent system are unjustified. In short, all members of any total system have the same justificational status. This makes little sense: Normal people have a mix of justified and unjustified beliefs, a mix that coherentism has no obvious resources to accommodate.

1.5 FOUNDATIONALISM AND BASIC BELIEFS

Foundationalism's response to the regress problem says that every root in a tree of (successful) inferential justification terminates after finitely many steps. Each stopping point (or starting point, one might say) is a belief that possesses justification it does not get from further (premise) beliefs. Such starting points, or "basic" beliefs, have two crucial properties: (A) they are uninferred and (B) they are justified. In addition, foundationalist theories say that there are inferential relations between basic and nonbasic beliefs such that enough of the latter are justified to avert the threat of global skepticism. Roughly speaking, a wide swath of our commonsense beliefs qualify as justified under foundationalist principles.

To succeed, a developed form of foundationalism must address four questions (we shall frequently use the terms *immediate* and *mediate* justification for basic and nonbasic justifiedness, respectively):

1. What does it take for a belief to attain the status of "justified" in an unmediated or noninferential fashion?
2. Which types of belief qualify as immediately justified? (For example, what types of propositional content lend themselves to immediate justification?)
3. What strength of justification must immediately justified beliefs possess according to foundationalism? Must they exhibit the highest grade of justification (i.e., certainty or infallibility)?
4. Assuming that many immediately justified beliefs are available, does this enable epistemic agents to draw enough reasonable inferences to further beliefs to dispel the specter of skepticism often laid at foundationalism's doorstep?

The foregoing questions pose the constructive tasks confronting the foundationalist program. Can it provide satisfactory answers? Foundationalism's critics have attacked it with energy and zeal, challenging the very possibility of its delivering the goods. We shall give examples of some roadblocks erected by the critics. Do these roadblocks eviscerate all prospects for a successful foundationalism, or do the opponents' criticisms fall apart when examined closely?

1.6 BASIC BELIEFS AS SELF-JUSTIFYING BELIEFS?

Focusing on the first (and central) question, one answer is that immediately justified beliefs are self-justifying beliefs. The term *self-justifying* often crops up in attempts to explain what is distinctive to immediate justifiedness. *Self-*justifying sounds like the right contrast with inference-based justifiedness, since inference is always a relation with other beliefs. So perhaps immediate justifiedness is what accrues when a belief is self-justifying. But what does it mean for a belief to justify itself? Is this even possible? Perhaps a belief justifies itself in case its occurrence logically entails the truth of its content. A belief to the effect that one has some beliefs seems to satisfy this condition. To be sure, few beliefs exemplify this property: Believing it is going to rain tomorrow does not entail it will rain tomorrow. Still, in those few cases in which a belief has this property, maybe it is immediately justified.

Would this answer to question (1) satisfy the goals of foundationalism? It is widely assumed that paradigm cases of immediately justified beliefs concern our own current (nonfactive) mental states—for example, "I believe I am in pain," or "I believe I want an espresso." But the proposed definition will not guarantee this result, first because the desired logical entailment will rarely hold. Does the fact that I believe I am sad logically entail that I am sad? That is highly dubious. Perhaps we are always correct when we classify our own current mental states—but is this logically, or even metaphysically, guaranteed? These are problematic theses.

However, even if it is true in such cases that believing we are in mental state M guarantees that we are in M, why should this imply justifiedness of the belief? For any truth L of logic, my believing L (trivially) entails its truth (because it is necessarily true). But it hardly follows that I am justified in believing it. L might be such a complex proposition that although I manage to grasp its content, I fail to understand how or why it is true. Thus, even when the content of a belief guarantees its truth, a bit of reflection shows that this does not guarantee its justifiedness.

Here is a slightly different way to understand self-justification. If person S believes that she is currently in mental state M and S *is* in M, then S is justified in believing this proposition. Under this second conception, self-justification works a little differently: It isn't the believing that confers justifiedness, but the belief's being true that confers justifiedness. This idea was defended by Roderick Chisholm (1977). Chisholm offered the following formula: "What justifies me in thinking I know that *a* is *F* is simply the fact that *a* is *F*."

Chisholm conceded that the formula does not apply to every proposition that instantiates it:

> Thus, in answer to "What justification do you have for counting it as evident that there can be no life on the moon?" it would be inappropriate—and presumptuous—simply to reiterate "There can be no life on the moon." But we can state our justification for *certain* propositions about our beliefs, and certain propositions about our thoughts, merely by reiterating those propositions. (21, italics added)

Unfortunately, Chisholm offers no explanation of *why* it is appropriate to defend the justification for certain beliefs by reiteration but not others. Why does it work for first-person mental-state propositions but not for third-person mental-state propositions? What is the crucial difference between them?

Chisholm introduced a phrase to distinguish the one class of propositions or states of affairs from the other class. States of affairs are called *self-presenting* when it is appropriate to reiterate the fact in defense of a justification claim. This provides a label, but not the slightest explanation of the alleged epistemic difference. Perhaps such an explanation is intended when Chisholm says that a self-presenting state is one that is "apprehended through itself" (22). But this obscure notion is never explained. If the phrase of immediate justification is not explained more clearly, a defense of immediate justification remains elusive.

1.7 DEBUNKING SOME DEBUNKERS OF IMMEDIATE JUSTIFICATION

Perhaps the case for foundationalism takes on an obscure character from these unhelpful attempts to clarify the nature of immediate justification. In principle, however, this might not be a hopeless task. To get some appreciation for what foundationalism is after, consider first the following analogy.

In European history becoming a monarch was usually a matter of lineage: One became a monarch by standing in the right lineal relationship to someone else who had been a monarch. Not every monarch, however, became one by inheritance; some became monarchs by leading a conquest of new territory. So, in addition to monarchy being passed on by inheritance, it was also sometimes acquired de novo. This parallels the idea of justificational status (J-status) being generated without any justificational "juice" being transmitted from one possessor to another. This makes good sense in the abstract, and there seems no reason to preclude it, at least in principle, in the realm of epistemic justification.

Nonetheless, many epistemologists are pessimistic about the prospects for foundationalism. Several have tried to debunk immediate justifiedness by demonstrating its inherent unrealizability. William Alston (1983) has surveyed some of these attempted debunkings and undertaken to debunk those

debunkings. Let us consider a few of these critical assaults and Alston's rejoinders to them.[1] We begin with the ground rules Alston lays down, more specifically the core notions of mediate justification (MJ) and immediate justification (IJ):

> (MJ) S is *mediately* justified in believing that p if and only if S is justified in believing that p by virtue of some relation this belief has to some other justified belief(s) of S.
>
> (IJ) S is *immediately* justified in believing that p if and only if S is justified in believing that p by virtue of something other than some relation this belief has to some other justified belief(s) of S. (Alston 1989, 58)

Immediate justification is defined here purely negatively: It is justification that does not involve any inferential relation or other relation to other justified beliefs. This skeletal definition does not supply an answer to question (3) posed above; it does not provide positive conditions for immediate justifiedness. Alston goes on to argue that the relevant debunkers he addresses make sneaky attempts to substitute a different conception of immediate justification than the simple conception, (IJ). In effect he accuses debunkers of quietly seeking to impose more stringent conditions on immediate justification than foundationalism needs.

The basic maneuver may be illustrated in terms of an analogy. You are going to a concert, ticket in hand. You approach the entrance to the concert venue and hand your ticket to the ticket-taker standing at the door. The ticket-taker says, "Well, OK. You have a ticket; but that's not enough for entry. You must also produce proof that you personally purchased this ticket from an authorized ticket agency plus the date of your purchase. No such proof is on the ticket itself." You would be outraged, of course. Permission to enter public events of this kind is normally guaranteed by presentation of an applicable ticket. Further proof of how and when the ticket was acquired is not needed. Such a surprise "upping of the ante" would be frowned upon in everyday life. Analogously, Alston's plan is to rebut certain debunkers of immediate justification who (in his view) unfairly and unreasonably "up the ante" for immediate justifiedness.

Alston picks out several debunkers who offer "level ascent" arguments against the possibility of immediate justification. Level ascent is exemplified in our concert ticket example. The first-level requirement for admission—normally the only requirement—is presentation of a legitimate ticket. The ticket-taker goes up a level by requiring not only a legitimate ticket but proof that it is a legitimate ticket. A similar inflation of reasonable requirements for immediate justification is what Alston claims to detect in the work of certain critics.

The first such critic is Wilfrid Sellars (1956). Alston describes what Sellars says about a certain approach to observational knowledge, which might account for how such knowledge can be immediate. This approach would say, roughly, that an observer has perceptual knowledge of an object's being

green if, in the presence of green objects, he reliably forms beliefs that they are green. Sellars, arguing to the contrary, says that Jones does not *know* that the object is green unless he is able to reflect on his own performance and take the formation of his statement or belief as a reason for supposing that a green object is present. Alston diagnoses the maneuver as follows:

> Sellars is clearly denying that observational knowledge is or can be immediate knowledge, as that term was explained above [in (IJ)]. His reason for denying it clearly falls under our Level Ascent rubric. One's belief counts as knowledge only if one knows something about the epistemic status of that belief, viz., that it counts as a reliable sign of the fact believed. And, equally clearly, this move could be used against *any* claim to immediate knowledge. (1989, 66)

As further support for his interpretation and rejection of Sellars's level ascent argument, Alston accuses Sellars (and others) of confusing two subtly different concepts that need to be kept separate. The two concepts are, first, the concept of a belief's being justified and, second, engaging in some kind of activity vis-à-vis the belief that reflects on it and establishes its legitimacy. In the concert ticket case, there is a difference between the ticket's being legitimate and the person's having an additional document to prove or establish the ticket's legitimacy. Similarly, Alston argues, it is one thing for a belief to have a certain epistemic property (e.g., being justified); it is quite another for the believer to have a higher-level justification for believing or asserting that it is legitimate. Immediate justification is a first-level property like having legitimacy. When one insists, as Sellars appears to do, that having legitimacy also involves having a higher-level justification about the belief's legitimacy, this inevitably brings in additional beliefs and keeps the original belief from being immediately justified. For the foundationalist, the proper solution is to reject the insistence on any extra, higher-level requirement.

Another epistemologist who tries to debunk basic justification by deploying a level ascent maneuver is Laurence BonJour (1985), who at the time was a coherentist. BonJour begins with the usual assumption that foundationalism can succeed only if there is a certain class of beliefs that are immediately justified. Moreover, given his discussion of justifiedness, a belief's being justified implies that it is highly likely to be true. Thus, whatever feature you might specify that promises to legitimize a belief as immediately justified, possession of that feature must imply that the belief is highly likely to be true. Call your favorite candidate for the correct feature Φ. Then for a belief to qualify as basic, says BonJour, the premises of the following argument must themselves be at least justified:

i. Belief B has feature Φ.
ii. Beliefs having feature Φ are highly likely to be true.

Therefore, B is highly likely to be true.

BonJour then continues as follows:

> And if we now assume, reasonably enough, that for B to be justified for a particular person (at a particular time) it is necessary, not merely that a justification for B exist in the abstract, but that the person in question be in cognitive possession of the justification, we get the result that B is not basic after all since its justification depends on that of at least one other empirical belief. (1985, 5–6)

In other words, BonJour is saying that it's not enough for the premises of his argument to be true—for example, that B has feature Φ (being likely to be true). It is also necessary that the believer be justified in believing that B is likely to be true. But this is an extra, higher-level belief that is said to be needed for first-order justification. If this were indeed necessary, what is initially claimed to be a basic belief depends for its justification on the justifiedness of a higher-level belief as well. If this were so, it would obviously undermine foundationalism. The argument is again a level ascent style of argument. It says that the justifiedness of any first-level belief B requires the epistemic agent also to have a justified higher-level belief B* to the effect that B is likely to be true. It is not enough for the first-level belief, B, genuinely to be likely to be true. In addition, the agent must have a different belief, B*, that justifies him or her in believing that B is likely to be true. This automatically implies that B is not immediately justified, at most only mediately justified.

Alston's response to BonJour is now predictable: He challenges BonJour's crucial step of requiring that for S's first-level belief to be justified there must be a justified higher-order belief to the effect that the first-level belief is highly likely to be true. "In other words, in order that I be justified in accepting B, I must know, or be justified in believing, the premises of the above argument. And why should we suppose that?" (1989, 73). By Alston's lights, BonJour just begs a crucial question, and doesn't really argue for it. Alston concludes, "Thus in BonJour, as in Sellars, the contention that putatively immediate knowledge really rests on higher-level reasons itself rests on a foundation of sand" (77). Alston offers no positive defense of foundationalism by specifying a particular criterion for immediate justifiedness, but he does poke a large hole in the indicated attempts to refute the very possibility of foundationalism.[2]

1.8 STRENGTH, CONTENTS, DEFEATERS, AND BOOSTERS

Our discussion has focused almost entirely on our earlier question (1): What are the necessary and sufficient conditions for obtaining immediate justification? Let us now say a few words about questions (2) and (3) in our original list of questions (Section 1.5).

Descartes was responsible for jump-starting the foundationalist picture of knowledge and justification. He also led the way with the idea that the starting point of the epistemic justification project resides in our own states

of mind. His famous maxim "I think, therefore I am" and the struggles he inaugurated about how to proceed from mental states to states of the external world pervaded several centuries of epistemology. In the twentieth century, however, foundationalists started to rethink matters and decided that even physical object propositions ("There's a pear on the table") might be an immediately justified belief. A resolution of this debate is still up in the air.[3]

Roughly the same story can be told for debates over question (2). Descartes again set things in motion by identifying the epistemological project with a quest for certainty. He argued that certainty could be attained in the case of first-person current mental propositions, and certainty was just the strength of justification to demand for basic beliefs. In the last century, however, epistemologists began to have second thoughts: Perhaps immediate justifiedness need not require the highest level of justifiedness. There is no obvious connection between the type of belief content and this required justificational strength, although some such connections have long been defended.

A less explored topic is the impact of so-called defeaters on the problem of justificational types. In foundationalist lore, there are two kinds of (positive) justificational statuses, either of which can characterize a justified belief. There is immediately justified belief and mediately justified belief. This is an exhaustive and exclusive typology. All justified beliefs are justified in one manner or the other; no justified belief has its justification in both ways. But this cannot be an adequate portrait of the epistemic landscape. In fact, the justificational source for many beliefs can be both immediate and mediate, both direct and inferential.

How can this transpire? First, let's introduce some terminology. When a proposition gets an "injection" of justification from some source, let us say it is *prima facie* justified for the epistemic agent. Prima facie justification is provisional, or tentative, justification, as opposed to justification on balance, or all things considered. Whether the J-status of the proposition (for the given agent) changes to on-balance justified may depend on other sources. Some of those other sources might cut against the proposition, thereby reducing its initial positive justifiedness. If that happens (to a sufficient extent), we say that the undercutting or undermining sources *defeat* the prima facie justification arising from the original source. The beliefs (or whatnot) that perform the undermining are called *defeaters*. We might also introduce a term for factors that add justificational support to a proposition: let us call them *boosters*. Both pieces of terminology can also be used in association with the language of "reasons." Doxastic or nondoxastic reasons may be boosters for believing p. If there are both positive and negative sorts of reasons and the negative ones are strong enough, they defeat the prima facie reasons in favor of p. In that case the *ultima facie* (all things considered) J-status of the proposition is unjustified. Alternatively, though an initial injection of justifiedness may be insufficient to bring a proposition across the threshold of justifiedness, getting a booster from a new source may add enough to bring the proposition across the justificational threshold.

Suppose you are looking at the sky and spot a plane with an unfamiliar contour. Is it one of those new-fangled 829s, you wonder? That's what they look like, but the plane is partly in the clouds, it's getting dark, and your distance vision was never so good anyway. But you now recall reading that this new model was going to be flight-tested in your area today. So you conclude, firmly, that what you are seeing is an 829. Your belief is justified but its source of justification is twofold: both noninferential and inferential. The first source is visual, presumably falling in the category of immediate justifiedness. The second source is background information, falling in the category of mediate justifiedness. The two sources jointly supply enough justificational "juice" to make the belief ultima facie justified. But neither source alone would reach across that threshold. A satisfactory theory of justification, obviously, must be capable of handling cases of this kind. If foundationalism is to be such a theory, it had better be able to handle them. Fortunately for foundationalism, this seems feasible with only a slight adjustment. Foundationalism can simply take the properties of immediate and mediate justifiedness to pertain not only (or even primarily) to beliefs, but to sources, factors, or components of justifiedness. Thus, the same belief can have both a mediate source of justifiedness and also an immediate source.

Coherentists will here jump in eagerly. They will point out that for *any* belief one is tempted to form by an immediate source, its justificational status can always in principle be influenced by other beliefs one already holds. "This is just what we have been saying all along," coherentists will exclaim. A careful epistemological should resist this claim. This is not *exactly* what coherentists have said right along; at least it's not the entirety of what they say. Another, stronger thesis they endorse is that inference, or inferential relations to other beliefs, is the *only* source of justification. Coherentism denies the existence of any source of justifiedness other than inference. So what was conceded in the previous two paragraphs is not tantamount to coherentism. But it does give ground to coherentism at least to some degree. It also seems to require foundationalism to countenance the fact that many basic beliefs owe their justification, in part, to inferential sources, albeit these basic beliefs would also have *some* justifiedness arise from noninferential sources.

1.9 MOVING FROM THE FOUNDATIONS TO HIGHER FLOORS (QUESTION 4)

Cartesian foundationalism is the variant of foundationalism that confines basic beliefs to beliefs about first-person current mental states. Setting aside the question of how beliefs about mental states gain immediate justifiedness, we now ask how a subject can proceed from such beliefs to propositions about material things in the external world that also qualify as justified. Clearly, this requires there to be legitimate methods for inferring such external-world propositions from propositions about current experience. Are there such methods? Obviously, we cannot logically deduce external-world

propositions from ones about our own inner experiences. According to many epistemologists and philosophers of science, however, our most important and interesting inferences are not deductive anyway. A standard kind of inference—used both in everyday life and in science—is *inference to the best explanation* (also called *abductive* inference).

After this morning's snowfall, I see some tracks in the snow. I infer that the tracks were made by a deer because they look like deer tracks and deer frequent my neighborhood. This looks like a good inference and it's clearly an explanatory inference. I regard the deer hypothesis as a good explanation of the observed tracks in the snow. But my neighbor tells me: "A company called 'Toys for Tricksters' recently marketed a machine that simulates deer tracks. Its tracks in the snow are indistinguishable from genuine deer tracks (and leave no telltale wheel tracks). These machines have become very popular among jokesters." My neighbor claims that my tracks were produced by such a machine; or, when pushed a bit, he says that it is perfectly possible that they were so caused. In either case he concludes that I am not justified in believing it was a deer.

How should we adjudicate this issue? Some inference principle is needed here, one for explanatory inference in particular. Here is a candidate principle:

> (JEI) *Justification by Explanatory Inference.* If hypothesis H purports to explain S's evidence E, and there is no incompatible hypothesis H' that provides a better or equally good explanation of E, then S is justified in believing H on the basis of E. [OR: . . . then S is justified in believing hypothesis H on the basis of E if and only if there is no incompatible hypothesis H' that provides a better or equally good explanation of E.]

If we substitute for "H" the hypothesis that the deer tracks were caused by a deer, and substitute for "E" the observed deer tracks, does JEI sanction my believing the deer hypothesis? Or does the deer-track simulator hypothesis constitute a better or equally good alternative explanation of the tracks? If my believing the deer hypothesis does not satisfy JEI, but JEI is retained as the relevant inference principle, then I should withhold judgment rather than flat-out believe it.

What is involved in one hypothesis being explanatorily superior to another? We shall briefly address this question here and return to it in greater depth in Chapter 4 (Section 4.5). Here are a few principles that have won some degree of acceptance among philosophers (Beebe 2009):

- *Explanatory Simplicity:* Other things being equal, a theory that posits fewer primitive explanatory notions should be preferred to one that posits more.
- *Coherence with Background Knowledge:* Other things being equal, a theory that fits better with other widely accepted theories and background knowledge should be preferred to a theory that fits less well.

Explanatory Depth: Other things being equal, a theory that provides a more illuminating explanation of the relevant data should be preferred to a theory that provides a less illuminating explanation.

Avoidance of Ad Hoc Elements: Other things being equal, a theory that has fewer ad hoc elements should be preferred to a theory that has more ad hoc elements.

How might these ideas be applied to the problem of inferring propositions about the external world from propositions about our own inner experiences? Most philosophers would say that the existence of genuinely material things (chairs, rocks, buildings) is the best explanation, for each of us, of our collection of mental experiences. But George Berkeley disagreed. As a better explanation he offered the hypothesis of a nonmaterial thing—namely God (an infinite mind)—that causes our sequences of experience. He defended this partly by appeal to the greater simplicity of the God hypothesis and partly by appeal to the principle of greater coherence with background knowledge. We already know that minds can cause experiences, but we don't already know that material substances can cause experiences. So it is unclear whether foundationalism can use an inference-to-the-best-explanation principle (whether JEI or another such principle) to show how beliefs about the external world can justifiedly be based on beliefs about our own mental experiences (the foundations). We shall return to this question in Chapter 4.

QUESTIONS

1. The first section of the chapter draws a parallel between moral rightness and wrongness, on the one hand, and epistemic rightness and wrongness on the other. The latter is usually described in terms of being justified or unjustified, rational or irrational in holding certain beliefs. At the same time it was said that moral and epistemic normativity are different species of normativity. Can we really hold both of these things? How can there be a strong commonality between moral and epistemic normativity if they are two different species of rightness and wrongness? How would you explain what is common to the two realms and what is different?
2. Descartes seemed to regard the foundations of his belief system as beliefs he could trust as the basis for other beliefs, on the analogy of what a house-owner does in relying on the foundations of the house to support its higher stories. Does this mean that the foundations must always be stronger than the higher floors? How would this analogy play out in the realm of belief and justification? Must a foundational, or "basic," belief be better justified than any nonbasic justified belief? Descartes said that his foundational belief was "I think." Is that better justified than any other (nonbasic) belief one might have? What property does this belief have that sets it apart from nonbasic justified beliefs and make its justification so strong? Does this property (or properties) hold for all basically justified beliefs? What should a wise foundationalist hold?

3. Suppose you are a member of a jury that is trying a murder charge against a defendant. As often happens in such cases (at least on television), much seems to depend on whether the prosecution's case "hangs together" better than the rival story of what transpired presented by the defense. The first story, of course, includes a proposition about how the defendant did the deed. The second story includes propositions about how the defendant was nowhere near the scene of the crime. Aren't you, as a juror, better justified in believing whichever story is the more coherent one, the one that "hangs together" better? If this is right, doesn't it prove the truth of coherentism? How might a foundationalist reply to this argument?
4. Sticking with the murder case example, suppose the prosecution contends that its own story of what transpired is simpler than the story told by the defense and it has fewer ad hoc elements. Suppose you agree with these claims. Is it clear that this should convince you of the prosecution's story? Why is a simpler story and one with fewer ad hoc elements more worthy of belief? Either explain why it is more belief-worthy, or indicate why it is a mistake to assign simplicity and minimization of ad hoc elements such heavy weight in determining justification.
5. At the end of Section 1.4 it was objected that according to coherentism all members of a highly coherent system will be justified and all members of a weakly coherent system will be unjustified. In other words, for any system, all of its member beliefs have the same justificational status. This is an objection because it's clear that this is not how justification usually works: Normally people have some justified and some unjustified beliefs, despite having just one current system. This appears to be a serious problem given the formulations of coherentism presented in the chapter. But maybe a new formulation of coherentism could avoid this problem. Can you suggest one?

FURTHER READING

Alston, William (1980). "Level Confusions in Epistemology." *Midwest Studies in Philosophy* 5, 135–150. Reprinted in Alston (1989). *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.

BonJour, Laurence (1985). "The Elements of Coherentism." Chapter 5 in *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

Haack, Susan (1999). "The Foundherentist Theory of Epistemic Justification." In Louis Pojman (ed.), *The Theory of Knowledge: Classical and Contemporary Readings* (2nd ed.). Belmont, CA: Wadsworth.

Huemer, Michael (2001). *Skepticism and the Veil of Perception* (pp. 98–118). Lanham, MD: Rowman and Littlefield.

Klein, Peter (2005). "Infinitism Is the Solution to the Regress Problem." In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (pp. 131–140). Oxford: Blackwell.

Lyons, Jack C. (2009). *Perception and Basic Beliefs* (Chapters 1, 2, and 4). New York: Oxford University Press.

Pollock, John, and Joseph Cruz (1999). *Contemporary Theories of Knowledge* (2nd ed., Chapter 2, "Foundations Theories"). Lanham, MD: Rowman and Littlefield.

Williams, Michael (2001). *Problems of Knowledge*. New York: Oxford University Press.

NOTES

1. Note that Alston's official target in his paper is knowledge rather than justification. But since the present chapter is devoted to justification rather than knowledge, our exposition will be an adaptation of Alston's discussion.
2. By contrast, Jack Lyons does specify a clearly stated criterion for immediate justifiedness. See Lyons (2009).
3. We return to this topic in Chapter 6.

2

Two Debates About Justification:

Evidentialism vs. Reliabilism and Internalism vs. Externalism

Alvin I. Goldman

2.1 JUSTIFICATION AND EVIDENCE

As we have seen, foundationalism and coherentism are the leading theories of the structure of justification. They also constitute the dominant players in what might be called traditional epistemology. Although they disagree sharply on the question of structure, they share a number of assumptions. Much of this chapter focuses on contrasts between traditional justification theories, mainly represented by *mentalist evidentialism*, and untraditional (or less traditional) theories, mainly represented by *process reliabilism*. A lively opposition between these two theories is revealed both by a head-to-head debate between them and by an associated debate between the broader approaches they exemplify, *internalism* and *externalism*.

The term *justifier* refers to anything that helps make a belief state justified or unjustified. In other words, it's anything that contributes, positively or negatively, to the justificational status (*J-status*) of a target belief. Foundationalism and coherentism, even between them, offer only two types of justifiers. The first type is (other) *belief states* of the subject, beliefs from which the target belief can be inferred. These are the things to which we usually appeal when asked or challenged to defend a specified belief. When I reply, "I believe P because of X, Y, and Z," the reasons cited are normally propositions I *believe*. (Sometimes I also know them, but that implies belief in them.) A second type of justifier is *experiential* states, such as perceptual and memory experiences. (What is meant by a memory experience is a conscious seeming-to-remember episode.) As we saw in Chapter 1, coherentism restricts justifiers to beliefs (or other doxastic states) exclusively. It holds, in effect, that all justification takes place by inference, or inferential relations. By contrast, foundationalism allows justification to be conferred by experiences as well as by (other) beliefs.

One thesis shared by traditional foundationalism and coherentism is that all justifiers are mental states of one sort or another—more specifically,