4
EXTERNALISM

1. Motivation

All of the theories discussed so far have been internalist theories. Doxastic theories take the justifiability of a belief to be a function exclusively of what else one believes. Internalist theories in general loosen that requirement, taking justifiability to be determined more generally by one's internal states. Beliefs are internal states, but so are perceptual states, memory states, and so on. Internal states have been only vaguely characterized as those to which we have "direct access". This vagueness must eventually be remedied, but further clarification will have to await later chapters. Externalist theories loosen the requirement for justifiability still further, insisting that more than the believer's internal states is relevant to the justifiability of a belief. For instance, reliabilism takes the reliability of the cognitive process generating the belief to be relevant to its justifiability.

The primary motivation for most externalist views proceeds in two stages. The first consists of the rejection of all doxastic theories. That rejection was defended above on the grounds that doxastic theories cannot handle perceptual input, which is basically a nondoxastic process that is nevertheless subject to rational evaluation. Thus we must select an epistemological theory from among nondoxastic theories. The externalist proposes that this selection be driven by a particular intuition. This is the intuition that we want our beliefs to be probable—we should not hold a belief unless it is probably true. Probability, however exactly it is brought to bear on the selection of beliefs, is an external consideration. The probability of one's belief, or the reliability of the cognitive process producing it, is not something to which one has direct access. Thus we are led to an externalist theory that evaluates the justification of a belief at least partly on the basis of external considerations of probability.

There is also a secondary motivation for the particular kinds of externalist theories that have been proposed in the literature. The doxastic theories discussed above propose very elaborate criteria for the justification of a belief. For instance, both the foundations theory of chapter two and the version of direct realism sketched at the end of the last chapter proceed by laying down a complex array of epistemic rules governing the justification of various kinds of judgments. A linear coherence theory proceeds similarly, adopting a structure of reasons similar to those involved in foundations theories (although in the case of negative coherence theories the rules only concern defeaters). And a similar point can be made about extant holistic coherence theories. For example, although Lehrer's theory does not proceed in terms of linear reasons, he too proposes a very complicated criterion for justification. (The full complexity of his criterion is not indicated by the brief sketch of his theory given in chapter three.) All of these theories proceed by initially taking the concept of epistemic justification for granted, and then using our intuitions about epistemic justification to guide us in the construction of a criterion that accords with those intuitions. A telling objection can be raised against all of these internalist theories—they are simultaneously ad hoc and incomplete. They are ad hoc in that they propose arrays of epistemic rules without giving any systematic account of why those should be the right rules, and they are incomplete in that they propose no illuminating analysis of epistemic justification.[1] The methodology of internalism has been to describe our reasoning rather than to justify it or explain it. These two points are connected. As long as we take the concept of epistemic justification to be primitive and unanalyzed, there is no way to *prove* that a particular epistemic rule is a correct rule. All we can do is collect rules that seem intuitively right, but we are left without any way of justifying or supporting our intuitions.

It might be responded that internalist theories do not leave the concept of epistemic justification unanalyzed. The criteria of justifiedness that they propose could be regarded as analyses of justification.[2] But even if one of these criteria correctly described which beliefs are justified, it would not explain what epistemic justification is all about. The criterion would not provide an *illuminating* analysis. Because of its very complexity we would be left wondering why we should employ such a concept of epistemic justification. Its use would be unmotivated. What we would have is basically an ad hoc theory that is contrived to give the right answer but is unable to explain in any deep way why that is the right answer. What we really want is an analysis of epistemic justification that makes it manifest why we should be interested in the notion. The analysis could then be used to generate a principled account of epistemic norms. An account of this sort would not proceed by simply listing the epistemic rules that seem to be required to license those beliefs we regard as actually being justified. Instead, it would *derive* the rules from the concept of epistemic justification. No epistemological theory yet discussed in this book has that character.

Most doxastic theorists have had remarkably little to say about the analysis of epistemic justification. This is not due to lack of interest—it is due to lack of ideas. Within the context of an internalist theory, the analysis of epistemic justification has proven to be an extremely difficult problem. Here is a respect in which externalist theories appear to have a marked advantage. If it is granted that the justifiedness of a belief can be determined partly by external considerations, then it becomes feasible to

---

1. Ernest Sosa (1980) raises this objection.

2. For example, Lehrer casts his theory as an analysis of epistemic justification.

try to analyze justification in terms of probability. That seems like a very hopeful approach. Lehrer (1974) expresses this idea succinctly in his critical discussion of foundationalism:

> What I object to is postulation without justification when it is perfectly clear that an unstated justification motivates the postulation. The justification is that people are so constructed that the beliefs in question, whether perceptual beliefs, memory beliefs, or whatnot, have a reasonable probability of being true. (77)

It is worth noting, however, that Lehrer's own theory diverges from this basic idea and thereby becomes subject to the same sort of criticism. This is because technical considerations lead him to an analysis no less complex than other internalist analyses. The simple intuition that justified beliefs must be probable provides no explanation at all for this complicated structure. Lehrer has merely replaced the complicated structure of foundations theories by another complicated structure, without any explanation for why justification should have either structure. He is as guilty of postulating epistemic rules as the foundations theorist.[3]

The hope that some simple analysis can be given of epistemic justification in terms of probability has had a powerful influence on epistemologists and has made externalist theories seem attractive. But it must be emphasized that such an analysis has to be simple. If the externalist is led to a complicated analysis, he will have been no more successful than the internalist in explaining why epistemic justification is a notion of interest to us. His theory will be no less ad hoc. The only virtue of the externalist theory will be the same one claimed (perhaps falsely) by all other theories, namely, that they correctly pick out the right beliefs as justified. Of course, that itself would be no mean feat.

To summarize, there are two sources for the appeal externalism has exercised on recent epistemologists. On the one hand, externalist theories seek to capture the common intuition that there is an intimate connection between epistemic justification and probability. On the other hand, there is the hope that an externalist analysis can explain what justification is all about rather than merely providing a correct criterion for justifiedness.

Externalist theories promise dividends not provided by any of the theories thus far discussed. But to evaluate these promissory notes, we must look more carefully at actual externalist theories. Two major kinds of externalist theories can be found in the literature—probabilism and reliabilism. Probabilism attempts to characterize the justifiedness of a belief in terms of its probability and the probability of related beliefs. Reliabilism seeks instead to characterize the justifiedness of a belief in terms of more general probabilities pertaining not just to the belief in

---

3. Lehrer's theory diverges from his basic intuition in another respect as well. His theory is formulated in terms of *beliefs about* probabilities rather than in terms of the probabilities themselves. Otherwise, his theory would not be a doxastic theory.

question but to the cognitive processes responsible for the belief. Among epistemologists, reliabilism is a more familiar version of externalism. Since reliabilism is partially predicated on probabilities, however, we will discuss probabilism first. Thus, probabilism and reliabilism will be the subjects of sections three and four, respectively. Before we can discuss them, however, we must lay some groundwork.

2. Varieties of Probability

Philosophers tend to make too facile a use of probability. They throw the word around with great abandon and often just assume that there is some way of making sense of their varied pronouncements, when frequently there is not. The main difficulty is that there is more than one kind of probability and philosophers tend to conflate them. A reasoned assessment of externalist theories requires us to make some careful distinctions between different kinds of probability. One important distinction is that between physical probability and epistemic probability. *Physical probability* pertains to the structure of the physical world and is independent of knowledge or opinion. For instance, the laws of quantum mechanics state physical probabilities. Physical probabilities are discovered by observing relative frequencies, and they are the subject matter of much of statistics. Physical probability provides the stereotype in terms of which most philosophers think of probability. But another important use of the word 'probable' in ordinary speech is to talk about degree of justification. For instance, after looking at the clues the detective may decide that it is probable that "the butler did it". Probability in this sense is directly concerned with knowledge and opinion and has no direct connection to the physical structure of the world. The *epistemic probability* of a proposition is a measure of its degree of justification. Epistemic probability is relative to a person and a time. It is an open question whether numerical values can be assigned to epistemic probabilities, and even if they can it is not a foregone conclusion that they will conform to the same mathematical principles (the probability calculus) as physical probabilities.

In addition to physical and epistemic probabilities, it is arguable that there are mixed physical/epistemic probabilities that are required for decision theory, weather forecasting, and so on. These probabilities appeal both to general physical facts about the world and to our knowledge of the present circumstances and how they relate to those general physical facts.[4] We will say more about these mixed probabilities below.

There is a second distinction that is related to but not identical with the distinction between physical and epistemic probabilities. We can

---

4. For more on the interrelationships between these various kinds of probability, see Pollock (1984a), and for a complete account see Pollock (1989).

distinguish between *definite probabilities* and *indefinite probabilities*. Definite probabilities are the probabilities that particular propositions are true or that particular states of affairs obtain. Indefinite probabilities, on the other hand, concern concepts or classes or properties rather than propositions. We can talk about the indefinite probability of a smoker contracting lung cancer. This is not about any particular smoker—it is about the class of all smokers, or about the property of being a smoker and its relationship to the property of contracting lung cancer. Some theories of probability take definite probabilities to be basic, and others begin with indefinite probabilities. Epistemic probabilities are always definite probabilities because they reflect a degree of belief in a particular proposition. Physical probabilities might be either. Theories taking physical probabilities to be closely related to relative frequencies make them indefinite probabilities. This is because relative frequencies concern classes or properties rather than single individuals. But there is also an important class of theories—"propensity theories"—that take the basic physical probabilities to pertain to individual objects. For example, we can talk about the probability that a particular coin will land heads on the next toss. Such "propensities" are definite probabilities.

With these preliminary distinctions before us, let us turn to some developed proposals regarding probability. There are three broad categories of probability theories to be found in the current literature, and externalists could in principle appeal to any of them, so we turn now to a brief sketch of each kind of theory.

2.1 Subjective Probability

Theories of subjective probability begin with the platitude that belief comes in degrees, in the sense that I may hold one belief more firmly than another, or that I can have varying degrees of confidence in different beliefs. The subjectivist is quick to explain, however, that he is using 'degree of belief' in a technical way. What he *means* by 'degree of belief' is something measured by betting behavior.[5] Officially, to say that a person has a degree of belief 2/3 in a proposition *P* is to say that he would accept a bet with 2:1 odds that *P* is true but would not accept a bet with less favorable odds. Given this technical construal of degrees of belief, the subjective probability of a proposition (relative to a person and a time) is identified with either (a) the person's degree of belief in that proposition, or (b) the degree of belief he rationally should have in the proposition given his overall situation. We can distinguish between these two conceptions of subjective probability as *actual degree of belief* and *rational degree of belief*.[6] It is generally claimed that subjective probability

---

5. See, for example, Rudolf Carnap (1962).

6. Leonard Savage (1954) proposed to distinguish between these by calling them "subjective probability" and "personalist probability", but this terminology is used only infrequently.

is a variety of epistemic probability.

There are problems for both conceptions of subjective probability. The principal difficulty for subjective probability as actual degree of belief is that the degrees of belief of real people will not satisfy the probability calculus. According to the probability calculus, probabilities satisfy the following three conditions:

<u>THE PROBABILITY CALCULUS</u>:

(1) $0 \leq \text{prob}(P) \leq 1$.

(2) If P and Q are logically incompatible with each other then $\text{prob}(P \vee Q) = \text{prob}(P)+\text{prob}(Q)$.

(3) If P is a tautology then $\text{prob}(P) = 1$.[7]

A person's degrees of belief are said to be *coherent* (in a sense unrelated to coherence theories of knowledge) if and only if they conform to the probability calculus. It is generally granted by all concerned that real people cannot be expected to have coherent degrees of belief. If there was ever any doubt about this, contemporary psychologists have delighted in establishing this experimentally. For some purposes the lack of coherence would not be a difficulty, but for the uses to which probability is put in epistemology it is generally essential that probability satisfy the probability calculus. Recall, for example, its use in Lehrer's theory. To carry out the kinds of calculations required by his theory, he must assume that probabilities conform to the probability calculus. It follows that subjective probability as actual degree of belief is of little use in epistemology.

Mainly because of the failure of actual degrees of belief to satisfy the probability calculus, most subjectivists adopt the "rational degree of belief" construal of subjective probability. But this construal is beset with its own problems. The first concerns whether rational degrees of belief satisfy the probability calculus any more than do actual degrees of belief. There is a standard argument that is supposed to show that they do. This is the Dutch book argument. In betting parlance, a "Dutch book" is a combination of bets on which a person will suffer a collective loss no matter what happens. For instance, suppose you are betting on a coin toss and are willing to accept odds of 2:1 that the coin will land heads and are also willing to accept odds of 2:1 that the coin will land tails. I could then place two bets with you, betting 50 cents against the coin landing heads and also betting 50 cents against the coin landing tails, with the result that no matter what happens I will have to pay you 50

---

7. It is customary to add a fourth axiom, to the effect that logically equivalent propositions have the same probability. This axiom implies that necessary truths have probability 1. We will not assume this axiom in the current discussion because it only exacerbates the problem of making fruitful use of probability within epistemology.

cents on one bet but you will have to pay me one dollar on the other. In other words, you have a guaranteed loss—Dutch book can be made against you. The Dutch book argument consists of a mathematical proof that if your degrees of belief (which, remember, are betting quotients) do not conform to the probability calculus then Dutch book can be made against you.[8] It is alleged that it is clearly irrational to put yourself in such a position, so it cannot be rational to have incoherent degrees of belief.

A number of objections can be raised to this argument. First, there is a familiar philosophical distinction between epistemic rationality and practical rationality. Epistemic rationality is concerned with what to believe, and falls within the purview of rationality. Practical rationality is concerned with what *to do*. Practical rationality deals with prudential concerns rather than epistemic concerns. As we saw in chapter one, these are distinct concepts. The Dutch book argument seems to be concerned with practical rationality—not epistemic rationality. It may be practically irrational to put yourself in a situation in which you are guaranteed to lose, but what has that to do with the epistemic rationality of belief? This is connected with the definition of subjective probability. Subjective probability is defined to be the degree of belief it is rational to have in a proposition, but this overlooks the distinction between practical and epistemic rationality. Which should be employed in the definition of subjective probability? The degree of belief it is epistemically rational to have in a proposition looks initially like what was defined above as epistemic probability, but this does not fit well with the technical notion of 'degree of belief' defined in terms of betting behavior. It does not seem to make any sense to say that certain betting behavior is or is not epistemically rational. Only practical rationality is applicable to betting behavior, and it seems that subjective probability must be understood in this way.

Two stances are now possible. It could be that the subjectivist has simply confused these two kinds of rationality and that the confusion pervades his entire theory. A more charitable reading of subjective probability theory would take it as an explicit attempt to reduce epistemic rationality to practical rationality. On this understanding, subjective probability is defined as the degree of belief it is practically rational to have, and so understood it may be used to explicate epistemic rationality. Subjective probability might be used in different ways to explicate epistemic rationality. The simplest proposal would be to identify the epistemic probability of a proposition with its (practical) subjective probability, but other more complicated alternatives are also possible and we will say more about them in the next section.

Adopting the charitable construal of subjective probability, does the Dutch book argument establish that subjective probabilities must conform

---

8. This was first proven by Bruno de Finetti (1937).

to the probability calculus? We do not think that it does. Contrary to the argument, it is not automatically irrational to accept odds allowing Dutch book to be made against you. If you are considering a very complicated set of bets (as you would be if you were betting on all your beliefs at once), it may be far from obvious whether the odds you accept satisfy the probability calculus. If you have no reason to suspect that they do not, and could not be expected to recognize that they do not without undertaking an extensive mathematical investigation of the situation, then surely you are not being *irrational* in accepting incoherent odds. You may be making a mistake of some sort, but you are not automatically irrational just because you make a mistake.

The Dutch book argument will not do it, but perhaps there is some other way of arguing that the degrees of belief it is practically rational to have must conform to the probability calculus. Let us just pretend for the moment that this is the case. Thus one constraint on rational degrees of belief becomes satisfaction of the probability calculus. Are there any other constraints? Some probability theorists write as if this were the only constraint. But others acknowledge that there must be further constraints. For example, so-called "Bayesian epistemologists" (discussed below) adopt constraints regarding how our degrees of belief should change as we acquire new evidence. But these are still rather minimal constraints.

The question we want to raise now is whether subjective probability as *the* degree of belief one should rationally have in a proposition is well defined. Specifically, is there any reason to think that, in each specific case, there is a *unique* degree of belief it is rationally permissible to have? Consider someone who has actual degrees of belief that do not satisfy the probability calculus (as we all have). If rational degrees of belief must be unique then there must be a unique way of transforming his actual degrees of belief into ideally rational degrees of belief. Whether this is so will depend upon what constraints there are. If the only constraint is that rational degrees of belief satisfy the probability calculus, there will be infinitely many ways of changing our actual degrees of belief so that the resulting degrees of belief are rational. The coherence constraint gives us no way to choose between them, because it tells us nothing at all about how our rationally changed degrees of belief should be related to our initial degrees of belief. The coherence constraint only concerns the product of changing our degrees of belief to make them rational; it does not concern how those resulting degrees of belief are gotten from the original incoherent degrees of belief. In fact, no constraints that have ever been proposed are of any help here. Some, like the Bayesian constraints, *sound* as if they should be helpful because they concern how degrees of belief should change under various circumstances, but they are of no actual help because they assume that the degrees of belief with which we begin satisfy the probability calculus.

The possible confusion between epistemic and practical rationality is relevant here. If we understand degree of belief as degree of confidence and we are allowed to bring all the resources of epistemology to bear, it seems likely that there will be a unique degree of belief (in the sense of degree of confidence) that we should have in any particular proposition in any fixed epistemic setting. However, characterizing subjective probability in this way involves giving up its characterization in terms of practical rationality. If we proceed in this way, then it will obviously be circular to turn around and use subjective probability to analyze epistemic justification, and the latter is the avowed purpose of the externalist endeavor.

If subjective probability is to be useful to the externalist, it must be defined in terms of practical rationality rather than epistemic rationality, so we cannot appeal to epistemic constraints to guarantee that there is a unique degree of belief we rationally ought to have in each proposition. The constraints can only be practical. It might be supposed that although no one has been able to enumerate them, there are some practical constraints on rational degrees of belief that will enable us to get from incoherent actual degrees of belief to unique rational degrees of belief satisfying the probability calculus. Is this at all plausible? We think not, because we think that unlike epistemically rational degrees of confidence, rational betting quotients are not always uniquely determined by the epistemic situation. For example, suppose I hold one ticket from each of two lotteries—lottery A and lottery B. One of the lotteries is a 100-ticket lottery, and the other is a 1000-ticket lottery, but I do not know which is which. I am now required to bet on whether I will win lottery A and whether I will win lottery B. Is there a unique rational bet that I should make? It does not seem so. I know that *either* my chances of winning A are .01 and my chances of winning B are .001, *or* vice versa, but I have no way of choosing between these two alternatives. Perhaps it is irrational for me to bet in accordance with any combination of odds other than one of these two, but there can be no rational constraint favoring one of them over the other. Alternatively it might be insisted that I should regard it as equally likely that (a) lottery A has 100 tickets and lottery B has 1000, and (b) lottery A has 1000 tickets and lottery B has 100, and so I should weigh these possibilities equally and arrive at a degree of belief of $.01 \times .5 + .001 \times .5 = .0055$ for winning either lottery. But this seems wrong. Intuitively, there would be nothing irrational about betting at odds of 1:99 and 1:999 instead.[9] It certainly seems as though there is no unique rational bet in this case, and it follows that the subjective probability does not exist. Furthermore, although this is a contrived case, reflection indicates that it is not atypical of most of the situations in which we actually find ourselves. So it must be concluded that unique practically

9. A further difficulty with the weighting strategy is that it seems to presuppose something like the Laplacian principle of insufficient reason, but as intuitive as that principle is it is also well known that it is inconsistent. In this connection, see Wesley Salmon (1966), 66ff.

rational degrees of belief rarely, if ever, exist.

To summarize, we regard the entire theory of subjective probability as being pervasively confused, turning upon a conflation of epistemic and practical rationality. If we define subjective probability in terms of practical rationality, subjective probabilities do not exist. If instead we define subjective probability in terms of epistemic rationality and forgo the characterization of degrees of belief in terms of betting behavior, then the notion makes sense but it becomes identical with epistemic probability defined as 'degree of justification'. In the latter case, none of the technical apparatus of subjective probability theory can be brought to bear any longer. The Dutch book argument becomes inapplicable, and there is no reason to attribute any particular mathematical structure to epistemic probabilities. In fact, reasons will be given shortly for denying that epistemic probabilities satisfy the probability calculus.

This conclusion will be unpopular among externalists, because subjective probability has been the favored kind of probability for use in probabilist theories. We stand by the negative conclusions we have drawn regarding subjective probability, but a number of philosophers are too firmly wedded to subjective probability to be dissuaded by such arguments, and accordingly we will occasionally pretend that the notion makes sense in order to discuss popular versions of probabilist theories of knowledge.

2.2 *Indefinite Physical Probabilities*

The most popular theories of physical probability relate physical probabilities to relative frequencies. Where A and B are properties, the relative frequency, freq[A/B], is the proportion of all actual B's that are A's. For example, given a coin that is tossed four times and then destroyed, if it lands heads just twice then the relative frequency of heads in tosses of that coin is 1/2. Some theories identify the physical probability, prob(A/B), with the relative frequency freq[A/B]. Others take the prob(A/B) to be the limit to which freq[A/B] would go if the set of all actual B's were hypothetically extended to an infinite set. Another alternative is to take prob(A/B) to be a measure of the proportion of B's that are A's in all physically possible worlds (rather than just in the actual world). On the latter theory, the connection between freq[A/B] and prob(A/B) is only epistemic—observation of relative frequencies in the actual world gives us evidence for the value of the probability.[10] The physical probabilities described by all of these theories are indefinite probabilities. They relate properties rather than attaching to propositions or states of affairs.

There is something quite commonsensical about the idea that the most fundamental kind of physical probability is an indefinite probability.

10. Pollock's theory is of the latter sort. It is sketched in his (1984a), and worked out in detail in his (1984d) and (1989).

At the very least our epistemological access to physical probabilities is by way of observed relative frequencies, and these always concern general properties. But it cannot be denied that for many purposes we require definite probabilities rather than indefinite probabilities. This is particularly true for decision theoretic purposes. For example, if I am betting on whether Blindsight will win the third race, my bet should be based on an assessment of the probability of Blindsight winning the third race. The latter is a definite probability. It is incumbent upon any theory of physical probability to tell us how such definite probabilities are related to the more fundamental indefinite probabilities. The traditional answer has been that definite probabilities are inferred from indefinite probabilities by what is called “direct inference”. The details of direct inference are problematic, and there are competing theories about how it should go, but in broad outline it is fairly simple. The basic idea is due to Hans Reichenbach (1949) who proposed that the definite probability, PROB(Aa), should be identified with the indefinite probability prob(A/B) where B includes as much information as possible about a, subject to the constraint that we have statistical information enabling us to evaluate prob(A/B). To illustrate, suppose we want to know the probability that Blindsight will win the third race. We know that he wins 1/5 of all the races in which he participates. We know many other things about him, for instance, that he is brown and his owner’s name is “Anne”, but if we have no information about how these are related to a horse’s winning races then we will ignore them in estimating the probability of his winning this race, and we will take the latter probability to be 1/5. On the other hand, if we do have more detailed information about Blindsight for which we have statistical information, then we will base our estimate of the definite probability on that more detailed information. For example, I might know that he is injured and know that he wins only 1/100 of all races in which he participates when he is injured. In that case I will estimate the definite probability to be 1/100 rather than 1/5.

Perhaps the best way to understand what is going on in direct inference is to take the definite probability PROB(Aa) to be the indefinite probability prob($Ax/x = a$ & $\mathbf{K}$), where $\mathbf{K}$ is the conjunction of all our justified beliefs. In direct inference we are estimating this indefinite probability on the basis of the known indefinite probability prob(Ax/Bx), where B includes all the things we are justified in believing about a and for which we know the relevant indefinite probabilities.[11]

For present purposes, the important thing to be emphasized about those definite probabilities at which we arrive by direct inference is that they are mixed physical/epistemic probabilities. We obtain definite probabilities by considering indefinite probabilities conditional on

11. For a detailed account of direct inference based upon this idea, see Pollock’s (1984d) and (1989).

properties we are justified in believing the objects in question to possess. The epistemic element is essential here. In decision theoretic contexts we seek to take account both of the probabilistic structure of the world and of our justified beliefs about the circumstances in which we find ourselves.

It is worth noting that an analogous account of direct inference to non-epistemic definite probabilities cannot work. Such an account would have us estimating the definite probability of *Aa* by considering all *truths* about *a* rather than all justified beliefs about *a*. In other words, in direct inference we would be trying to ascertain the value of prob($Ax/x = a$ & **T**) where **T** is the conjunction of all truths. But among the truths in **T** will be either *Aa* or ~*Aa*, so prob($Ax/x = a$ & **T**) is always either 1 or 0 depending upon whether *Aa* is true or false. Direct inference to such non-epistemic probabilities could never lead to intermediate values. Such an account of definite probabilities and direct inference would be epistemologically useless.

2.3 Definite Propensities

An approach to physical probabilities that is less popular but by no means moribund takes the fundamental physical probabilities to be definite probabilities that pertain to specific individuals. These are propensities. We can, for example, talk about the propensity of a particular coin to land heads on the next toss. These are supposed to be purely physical probabilities, untinged by any epistemic element, and are supposed to reflect ineliminable chance relationships in the world.[12] The defenders of propensities usually agree that propensities would always be either 0 or 1 in a deterministic world.[13]

Propensity theories have not been developed to the same extent as subjective theories and frequency theories, and most philosophers remain suspicious of propensities, but it would be premature to reject them outright. We must at least keep them in the back of our minds in considering what kinds of probabilities to use in formulating externalist theories of knowledge.

3. Probabilism

The distinction between probabilism and reliabilism can now be made precise by saying that probabilism seeks to characterize epistemic justification in terms of the definite probabilities of one's beliefs, while

---

12. Propensity theories have been proposed by Ian Hacking (1965), Isaac Levi (1967), Ronald Giere (1973), (1973a), and (1976), James Fetzer (1971), (1977), and (1981), D. H. Mellor (1969) and (1971), and Patrick Suppes (1973). A good general discussion of propensity theories can be found in Ellory Eells (1983).

13. Ronald Giere (1973), p. 475.

reliabilism seeks to characterize epistemic justification in terms of more general indefinite probabilities pertaining to such things as the reliability of the cognitive processes that produce the beliefs. Probabilism represents the most straightforward way of trying to capture the intuition that in acquiring beliefs we should adopt only probable beliefs.

3.1 The Simple Rule and Bayesian Epistemology

The simplest form of probabilism endorses what we call *the simple rule*:

> A person is justified in believing *P* if and only if the probability of *P* is sufficiently high.

This rule seems intuitively quite compelling, and at various times it has been endorsed by a wide spectrum of philosophers.[14] The simple rule, if acceptable, admirably satisfies the requirement that an analysis of epistemic justification must explain why the notion should be of interest to us. What could be more intuitive than the claim that in deciding what to believe we should be trying to ensure that our beliefs are probably true? The endorsement of the simple rule allows us to bring all of the mathematical power of the probability calculus to bear on epistemology, and the results have often seemed extremely fruitful. The simple rule pertains most directly to what we believe in a fixed epistemic setting, but what happens when our epistemic situation changes through the acquisition of new data (e.g., in perception)? Probabilists typically appeal to what is known as Bayes' Theorem,[15] according to which

$$\text{prob}(P/Q) = \text{prob}(Q/P) \times \frac{\text{prob}(P)}{\text{prob}(Q)}$$

Taking *P* to be the proposition whose epistemic status is to be evaluated and *Q* to be the new evidence, prob(*P* / *Q*) is interpreted as the probability of *P* given the new evidence; prob(*P*) is the prior probability of *P* (i.e., the probability prior to acquiring the evidence); prob(*Q*) is the prior probability of acquiring that evidence; and prob(*Q*/*P*) is the prior probability of acquiring the evidence given the specific assumption that *P* is true. This principle then tells us how to alter our probability assignments in the face of new evidence.

Epistemology based upon the simple rule and Bayes' Theorem is known as *Bayesian epistemology*.[16] It has exerted a strong influence on

---

14. It was endorsed, for example, by Roderick Chisholm (1957), p. 28, and Carl Hempel (1962), p. 155. Its most ardent recent defender is probably Henry Kyburg (1970) and (1974). See also Richard Jeffrey (1970), Rudolf Carnap (1962) and (1971), David Lewis (1980), and Isaac Levi (1980).

15. After Thomas Bayes.

16. Sometimes the term 'Bayesian epistemology' is reserved for theories

technically minded philosophers, partly because of the intuitiveness of its basic principles and partly because of its mathematical elegance and power. It has spawned an extensive literature and has seemed to be extremely fruitful when applied to problems like the problem of induction or the analysis of the confirmation of scientific theories. Unfortunately, Bayesian epistemologists have concentrated more on the mathematical elaboration of the theory than on its foundations. There are major problems with the foundations. These concern the very idea of acquiring new data. Note that it follows from the proposed interpretation of Bayes' theorem that when we acquire new data Q, it will come to have probability 1. This is because $\text{prob}(Q/Q) = 1$. But what is it to acquire new data through, for instance, perception? This is just the old problem of accommodating perceptual input within an epistemological theory. The beliefs we acquire through perception are ordinary beliefs about physical objects, and it seems most unreasonable to regard them as having probability 1. Furthermore, it follows from the probability calculus that if $\text{prob}(Q) = 1$ then for any proposition R, $\text{prob}(Q/R) = 1$. Thus if perceptual beliefs are given probability 1, the acquisition of further data can never lower that probability. But this is totally unreasonable. We can discover later that some of our perceptual beliefs are wrong.[17]

The idea that "data" should receive probability 1 is reminiscent of the appeal to epistemologically basic beliefs. What is happening here is that although Bayesian epistemology is a nondoxastic theory, it is nondoxastic in the wrong way. Doxastic theories fail to handle perception correctly because the only internal states to which they can appeal are beliefs. Specifically, they do not appeal to perceptual states. But beliefs are also the only internal states to which Bayesian epistemology appeals. Bayesian epistemology is nondoxastic because it appeals to probability, not because it appeals to internal states other than beliefs. Consequently, Bayesian epistemology encounters precisely the same sort of problem as do doxastic theories in accommodating perception. Contrary to both doxastic theories and Bayesian epistemology, the justifiability of a perceptual belief is partly a function of nondoxastic internal states.

It should be emphasized that the use of Bayes' theorem in describing perception and belief change is not required by the simple rule. These are independent principles. Thus we can reject Bayesian epistemology without rejecting the simple rule. The simple rule might be combined with a more sophisticated account of perception and memory without robbing it of the power inherent in its use of probability. And the simple rule would still retain its intuitive appeal in capturing the idea that what we should be doing in the epistemological evaluation of beliefs is choosing

---

proceeding in terms of subjective probability.

17. Richard Jeffrey (1965) has proposed a variant of the Bayesian rule that avoids this problem by allowing new data to have probability less than 1, but it does not tell us how to assign probability to the new data.

beliefs that are probable. Unfortunately, this elegant rationale for the simple rule begins to crumble when we examine the rule more closely. In evaluating the simple rule we must decide what kind of probability is involved in it. It is a definite probability, but we have taken note of the (at least putative) existence of four distinct kinds of definite probability: epistemic probability, subjective probability, mixed physical/epistemic probability, and propensities.

3.1.1 Epistemic Probabilities

The simple rule is a truism when interpreted in terms of epistemic probabilities. So understood it claims no more than that a belief is justified if and only if its degree of justification is sufficiently high. This claim cannot be faulted as long as it is understood that there is no presupposition either that epistemic probabilities are quantifiable or that if they are quantifiable then they satisfy the probability calculus. But, of course, understood in this way the simple rule is trivial and unilluminating. In particular, it does not constitute an analysis of epistemic justification, because epistemic probability is itself defined in terms of epistemic justification.

3.1.2 Subjective Probabilities

What is no doubt the favorite interpretation of the simple rule in contemporary philosophy is the one proceeding in terms of subjective probability. So construed, the simple rule becomes the claim that epistemic probabilities are the same as subjective probabilities. We have argued that no sense can be made of subjective probability, and we regard this as the most serious objection to the endorsement of the simple rule construed in terms of subjective probability. But suppose we waive this difficulty, pretending that there is always a unique betting quotient that one practically should accept for each given proposition. This construal of the simple rule then yields an analysis of epistemic justification that is not obviously circular, and an immense literature has grown up around it. Is this a plausible analysis?

The simple rule, construed in terms of subjective probabilities as actual degrees of belief, would tell us that a belief is justified if and only if it is firmly held. That obviously has nothing to recommend it, so let us confine our attention to subjective probabilities as rational degrees of belief. A simple objection to the use of such subjective probabilities in the simple rule consists of questioning whether practical rationality can be understood without first understanding epistemic rationality. It certainly seems that what we practically should do is a function in part of what we (epistemically) reasonably believe. Whether I should bet that Blindsight will win the next race is going to be determined in part by what I am justified in believing about Blindsight. If it is correct that practical rationality presupposes epistemic rationality, then it becomes circular to analyze epistemic rationality in terms of practical rationality, and hence it becomes circular to analyze epistemic rationality in terms of subjective probabilities.

3.1.3 Mixed Physical/Epistemic Probabilities

We might interpret the simple rule in terms of the mixed physical/epistemic probabilities that are obtained from indefinite physical probabilities by direct inference. This is the proposal of Henry Kyburg.[18] But any such proposal is subject to a simple objection—the resulting analysis of epistemic justification is circular. Recall that these physical/epistemic definite probabilities are obtained by direct inference from indefinite probabilities, and direct inference proceeds by considering indefinite probabilities conditional on *what we are justified in believing* about the objects in question. For example, what makes it true that the probability (for me) is 1/2 of Jamie having an accident while driving home this evening is that I am *justified in believing* that he is inebriated, driving on busy streets, and so on, and the indefinite probability of someone having an accident under those circumstances is 1/2. It is not the mere *fact* that Jamie is inebriated that makes the probability high. Only what I am justified in believing about Jamie can affect the mixed physical/epistemic probability. Thus mixed physical/epistemic probabilities cannot be used non-circularly in the analysis of epistemic justification.

3.1.4 Propensities

An analysis of epistemic justification in terms of propensities would not be obviously circular. Can we take a proposition to be justified if and only if it has a sufficiently high propensity to be true? It is hard to evaluate this suggestion without a better understanding of propensities. Notice, however, that according to most propensity theorists, nontrivial propensities only exist in nondeterministic worlds. If the world were deterministic, all propensities would be either 0 or 1, depending upon whether the proposition in question were true or false. The simple rule would then reduce to the absurd principle that we are justified in believing something if and only if it is true. We might avoid this objection by insisting that nontrivial propensities exist even in deterministic worlds, but to defend this we need a better understanding of propensities than anyone has yet provided. It is rather difficult to say much specifically about the propensity interpretation of the simple rule without a better theory of propensities.

3.1.5 General Difficulties

We have raised objections to each of the interpretations of the simple rule in terms of different kinds of definite probabilities. Some of those

18. In Kyburg (1974), and elsewhere.

objections are more telling than others. We turn now to some more general objections that apply simultaneously to all versions of the simple rule.

(a) Tautologies

A reasonably familiar objection is that it follows from the probability calculus that every tautology has probability 1.[19] It would then follow from the simple rule that we are justified in believing every tautology. Such a conclusion is clearly wrong. If we consider some even moderately complicated tautology such as

$$[P \leftrightarrow (R \vee \sim P)] \rightarrow R$$

it seems clear that *until we realize that it is a tautology*, we are not automatically justified in believing it. The only way to avoid this kind of counterexample to the simple rule is to reject the probability calculus, but that is a very fundamental feature of our concept of probability and rejecting it would largely emasculate probability. It is because probability has the nice mathematical structure captured by the probability calculus that it has proven so fruitful, and that mathematical structure has played an indispensable role in the employment of probability in epistemology.

(b) Epistemic Indifference

The preceding difficulty illustrates one respect in which epistemic justification seems to have a more complicated structure than can be captured by the probability calculus. Supposing that epistemic probability satisfies the probability calculus forces us to regard different propositions (different tautologies) as equally justified when it seems clear that we want to make epistemic distinctions between them. The converse problem also arises. There are cases in which we want to regard propositions as equally justified when the probability calculus would preclude that. Consider a pair of unrelated propositions, *P* and *Q*, regarding which we know essentially nothing. Under the circumstances, we should neither believe these propositions nor disbelieve them—we should withhold belief. We can express this by saying that we should be *epistemically indifferent* with respect to *P*, and also with respect to *Q*. We are precisely as justified (or unjustified) in believing *P* and in believing ~*P*, and similarly for *Q*. The simple rule and the probability calculus would require that in

19. We have avoided endorsing the standard axiom requiring logically equivalent propositions to have the same probability. That axiom must be added for a reasonable axiomatization of probability, but we have not endorsed it here in order to avoid begging questions against probabilism. Probabilism would encounter even more severe difficulties in the face of that axiom and its consequence that all necessary truths (not just tautologies) have probability 1.

this sort of case, prob(P) = prob(~P) = 1/2, and prob(Q) = prob(~Q) = 1/2. But now consider the disjunction ($P \vee Q$). If we are completely ignorant regarding P and Q and they are logically unrelated, then we are also completely ignorant regarding ($P \vee Q$). We should withhold belief with respect to ($P \vee Q$) just as we did with respect to P and with respect to Q. Thus we should conclude, as above, that prob($P \vee Q$) = prob(~($P \vee Q$)) = 1/2. The difficulty is that the probability calculus commits us to regarding ($P \vee Q$) as more probable than either P or Q individually, so we cannot assign probability 1/2 to all three of P, Q, and ($P \vee Q$).[20] Thus the simple rule precludes our being epistemically indifferent to all three of P, Q, and ($P \vee Q$), and yet intuition seems to indicate that there are circumstances under which such indifference is epistemically prescribed. Once again, the structure of epistemic justification is not properly reflected by the probability calculus.

(c) The Simple Rule Is Self-Defeating

The preceding objections to the simple rule are familiar ones in the philosophical literature, but for reasons that escape us, they have failed to make many confirmed probabilists repent. So consider a third objection, which seems to us show that probabilism based on the simple rule is a hopeless theory. This third objection argues that the theory is "self-defeating", in the sense that if it were true it would make it impossible for us to have any interesting justified beliefs. The argument is as follows. According to the simple rule, to be justified a belief must be highly probable. Many of the beliefs that we regard as justified are obtained by inference from other justified beliefs. We would like to be able to apply logical inference "blindly", simply assuming that if a conclusion is deduced from other highly probable conclusions, then it is itself highly probable. For this to be possible, the inference must proceed in terms of a "probabilistically valid" inference rule, where this notion is defined as follows:

> <u>DEFINITION:</u>
> An inference rule is *probabilistically valid* if and only if it follows from the probability calculus that whenever a conclusion can be inferred in accordance with it from a set of premises, the probability of the conclusion is at least as great as the probability of the least probable premise.

It is a consequence of the probability calculus that if P logically entails Q then PROB(Q) ≥ PROB(P). So deductive inference from a single premise always preserves high probability, and such inferences are

---

20. Technically, prob($P \vee Q$) = prob(P) + prob(Q) – prob(P & Q), and because P and Q are unrelated, prob(P & Q) = prob(P) × prob(Q) = 1/4, with the result that prob($P \vee Q$) = 3/4.

probabilistically valid. This seems to lend credence to the claim that inferences should always be made in accordance with probabilistically valid inference rules. However, many important inference rules proceed from multiple premises. Here are two such rules:

ADJUNCTION

$$\frac{P, Q}{(P \& Q)}$$

MODUS PONENS

$$\frac{P, (P \rightarrow Q)}{Q}$$

A problem now arises. Neither of these rules is probabilistically valid. The probability calculus implies that $\text{PROB}(P \ \& \ Q) \leq \text{PROB}(P)$ and $\text{PROB}(P \ \& \ Q) \leq \text{PROB}(Q)$. If P and Q are independent, $\text{PROB}(P \ \& \ Q) = \text{PROB}(P) \times \text{PROB}(Q)$, and so unless $\text{PROB}(P) = 1$ or $\text{PROB}(Q) = 1$, $\text{PROB}(P \& Q)$ will be less than either $\text{PROB}(P)$ or $\text{PROB}(Q)$. So adjunction is not probabilistically valid. The situation is similar for modus ponens. Nor is this just a problem for these two inference rules. Let us say that a premise *occurs essentially* in an inference rule if the rule becomes invalid when we delete the premise. It then turns out that *no* deductive inference rule is probabilistically valid if it has multiple premises all of which occur essentially. Thus the probabilist is committed to asserting that no such inference rules can be used blindly in drawing new justified conclusions. For instance, if a person justifiably believes P and justifiably believes Q, she cannot automatically infer (P & Q). Instead, she must somehow compute the probability of (P & Q) and then decide on that basis whether to believe it.

This is extremely counterintuitive. For instance, consider an engineer who is designing a bridge. She will combine a vast amount of information about material strength, weather conditions, maximum load, costs of various construction techniques, and so forth, to compute the size a particular girder must be. These various bits of information are, presumably, independent of one another, so if the engineer combines 100 pieces of information, each with a probability of .99, the conjunction of that information has a probability of only $.99^{100}$, which is approximately .366. According to the probabilist, she would be precluded from using all of this information simultaneously in an inference—but then it would be impossible to build bridges.

As a description of human reasoning, this seems clearly wrong. Once one has arrived at a set of conclusions, one does not hesitate to make further deductive inferences from them. But an even more serious difficulty for the probabilist is that the simple rule turns out to be self-defeating; if it were correct, it would be impossible to perform the very calculations required by the simple rule for determining whether a belief ought to be held. This arises from the fact that the simple rule would require a person to decide what to believe by computing probabilities. The difficulty is that the probability calculations themselves cannot be performed by a probabilist who endorses the simple rule. To illustrate the difficulty, suppose the reasoner has the following beliefs:

$\text{PROB}(P \vee Q) = \text{PROB}(P) + \text{PROB}(Q) - \text{PROB}(P \,\&\, Q)$
$\text{PROB}(P) = .5$
$\text{PROB}(Q) = .49$
$\text{PROB}(P \,\&\, Q) = 0.$

From this we would like the reasoner to compute that $\text{PROB}(P{\vee}Q) = .99$, and perhaps go on to adopt $(P{\vee}Q)$ as one of her beliefs. However, the probabilist cannot do this. The difficulty is that this computation is an example of a "blind use" of a deductive inference, and as such it is legitimate only if the inference is probabilistically valid. To determine the probabilistic validity of this inference, it must be treated on a par with all the other inferences performed by the reasoner. Although the premises are about probabilities, they must also be assigned probabilities ("higher-order probabilities") to be used in their manipulation. Viewing the inference in this way, we find that although the four premises do logically entail the conclusion, the inference is not probabilistically valid for the same reason that *modus ponens* and adjunction fail to be probabilistically valid. It is an inference from a mutliple premise set, and despite the entailment, the conclusion can be less probable than any of the premises.

The upshot of this is that the simple rule is self-defeating. On the one hand, it precludes a person from drawing conclusions on the basis of probabilistically invalid deductive inferences, requiring one instead to decide whether to believe putative conclusions by computing their probabilities. But on the other hand it makes it impossible to compute those probabilities because the computations involved are probabilistically invalid. The inescapable conclusion is that the simple rule is bankrupt.

We have raised three "formal" objections to the simple rule. It should be pointed out that these objections to the simple rule are also objections to a more general principle. A number of philosophers reject the identification of epistemic justification with probability, but nevertheless maintain that degree of epistemic justification "work like" probabilities, in the sense that they satisfy the probability calculus.[21] That claim is made equally untenable by the formal objections.

We regard these objections to the simple rule as decisive. The structure of epistemic justification is too complicated to be captured by the probability calculus. Why then does the rule seem so intuitive? We think that there is a twofold explanation for this. First, a very common use of 'probable' in English is to express epistemic probability, and the simple rule understood in terms of epistemic probability is a truism. Philosophers confuse this truism with more substantive versions of the simple rule that proceed in terms of other varieties of probability. Only those more substantive versions can hope to provide a noncircular account of epistemic justification, but those more substantive versions have fatal flaws and are not themselves directly supported by our intuitions.

The second part of the explanation amounts to observing that we also

---

21. Richard Fumerton (1995).

have the intuition that we should not believe something if it is improbable, where the kind of probability involved is the mixed physical/epistemic probability involved in decision theory. This intuition is correct, but it lends no support to probabilism. It is correct because mixed physical/epistemic probabilities are conditional on the conjunction of all justified beliefs, and hence the probability of any justified belief will automatically be 1.[22] But this lends no support to probabilism, because mixed physical/epistemic probabilities are defined in terms of justified belief and hence cannot be used noncircularly to analyze epistemic justification.

3.2 Other Forms of Probabilism

The simple rule is the most natural form of probabilism, but it fails due to the fact that the mathematical structure of epistemic justification cannot be captured by the probability calculus and hence degree of justification cannot be identified with any kind of probability conforming to the probability calculus. However, it is possible to construct more sophisticated forms of probabilism that escape this objection. These theories characterize epistemic justification in terms of probabilities, but they do not simply identify degree of justification with degree of probability. For example, recall Keith Lehrer's coherence theory. Its central thesis is:

> *P* is justified for *S* if and only if for each proposition *Q* competing with *P*, *S* believes *P* to be more probable than *Q*.

This is a doxastic theory because its appeal to probability is only via beliefs about probability, but it could be converted into a nondoxastic theory by appealing to the probabilities themselves:

> *P* is justified for *S* if and only if *P* is more probable than each proposition competing with *P*.

Supplementing this central criterion with a definition of competition will yield a version of probabilism more complicated than the simple rule. We might adopt Lehrer's definition of competition, or we might adopt another definition. Marshall Swain (1981) follows essentially this course. A slightly simplified version of his definition of competition is as follows:

Q is a competitor of *P* for *S* if and only if either:

(1) (a) *P* and *Q* are contingent,
(b) *Q* is negatively relevant to *P*, and
(c) *Q* is not equivalent to a disjunction of propositions one of whose disjuncts, *R*, is both (i) irrelevant to *P* and (ii) such that the probability of *R* is greater than or equal to the probability of *P*; or

(2) *P* is noncontingent and *Q* is ~*P* (p. 133).

---

22. It is a theorem of the probability calculus that prob(*P*/*P* & *Q*) = 1.

Swain's theory is just one example of a sophisticated kind of probabilism that escapes the objection to the simple rule that the structure of epistemic justification is too complicated to be captured by the probability calculus. Thus the formal objections to the simple rule do not refute probabilism in general.

Still, two general points can be made about all versions of probabilism. First, there appears to be no appropriate kind of probability for use in probabilist theories of knowledge. Such theories require a definite probability. They are circular if formulated in terms of epistemic probability. Only three other kinds of definite probability have been discussed in the literature: subjective probability, propensities, and mixed physical/epistemic probability. We have argued that the very concept of subjective probability is ill-defined—subjective probabilities do not exist. We regard mixed physical/epistemic probabilities as unproblematic, but they cannot be used in the analysis of epistemic justification because they already presuppose epistemic justification. Propensities are not sufficiently well understood to be of much use anywhere. Furthermore, a common view among propensity theorists is that nontrivial propensities only exist in nondeterministic worlds, but it seems pretty clear that any probabilist analysis of epistemic probability must proceed in terms of a variety of probability that can take values intermediate between 0 and 1 even in deterministic worlds. It seems inescapable that there is no appropriate kind of probability for use in probabilist theories of knowledge.

This point applies to any theory that has a probabilistic component embedded in it, no matter what other elaborations are proposed. For instance, William Alston (1988) attempts to combine elements of internalism and externalism by defending a view that has an internal access constraint while at the same time demanding that the grounds for a belief make it probable that the belief is true. Most of his efforts have been in developing the internalist side of his view. He acknowledges that it is extremely difficult to state with any precision just how the probabilistic part of the theory should be developed. We have offered some indication of why that is, and have pointed out some reasons for being pessimistic.

The second point to be made about probabilist theories concerns not their truth but their motivation. The original motivation for probabilism came from the intuition that what the epistemic evaluation of beliefs should be trying to ensure is that our beliefs are probable. If the simple rule were defensible, it would capture that intuition. But the simple rule is not defensible, and more complicated versions of probabilism do not capture this intuition. In fact, they are incompatible with the intuition—insofar as they diverge from the simple rule, they will have the consequence that we can be justified in believing improbable propositions and unjustified in believing probable propositions.

The intuition that we should only believe things when they are probable is a powerful one. It seems to lead directly to the simple rule, but there are overwhelming objections to the simple rule. What, then, should we make of this intuition? As we have indicated, it is epistemic probability that is involved in this intuition. In ordinary non-philosophical English, "probable" is used to express epistemic probability at least as often as it is used to express other kinds of probability, and it is a truism that a belief is justified if and only if its epistemic probability is sufficiently high. But, of course, epistemic probability is defined in terms of epistemic justification, so this provides no analysis of epistemic justification and no support for probabilism.

What should we conclude about probabilism at this point? Decisive objections can be raised against existing probabilist theories of knowledge, and they militate strongly against there being any defensible kind of probabilism. At the very least, the probabilist owes us an account of the kind of probability in terms of which he wants his theory to be understood, and there is good reason for being skeptical about there being any appropriate kind of probability. The most rational attitude to adopt towards probabilism at this point is healthy skepticism, but it cannot be regarded as absolutely certain that no probabilist theory can succeed. In chapter five, we will eventually present arguments that purport to show that no externalist theory of any kind can be correct, and those arguments should lay probabilism to rest for good.