6

The concept of acceptance

6.1 Definition

In everyday life, and also in science, opinions are often expressed by making categorical assertions, rather than by citing personal probabilities. The categorical assertion of H, when sincere and intentional, expresses a mental state of the assertor that I refer to as *acceptance* of H. This chapter will examine the concept of acceptance just defined; the following chapter will argue that it is an important concept for the philosophy of science.

What I am here calling *acceptance* is commonly called *belief.* Consider, for example, G. E. Moore's (1942, p. 543) observation that it is paradoxical to say 'H, but I do not believe that H.' The paradox is explained by the fact that sincere and intentional assertion of H is taken to be a sufficient condition for believing H.

It is, then, *part* of the folk concept of belief that it is a mental state expressed by sincere, intentional assertions. However, I think the folk concept involves other aspects too, and I am going to argue that the various aspects of the folk concept of belief do not all refer to the same thing. That is why I am calling the concept I have defined 'acceptance', rather than 'belief'.

My definition of acceptance assumes that you understand what a sincere assertion is. If that is not so, we are in trouble. For I cannot define sincerity except by saying that an assertion is sincere iff it was intended to express something the assertor accepts; and you won't understand this definition unless you already understand acceptance. But I think that everyone knows what counts as evidence for and against the sincerity of an assertion, and this shows that everyone has the sort of understanding of sincerity that we require here. (Note that my aim here is not to reduce the concept of acceptance to other concepts, but merely to identify what I mean by

'acceptance'. Thus the fact that the notion of sincerity cannot be explained without invoking the notion of acceptance does not mean the definition of acceptance cannot achieve its purpose.)

I say that acceptance is expressed by assertions that are both sincere and intentional. What I mean by an intentional assertion is that it asserts what the person intended to assert; thus slips of the tongue are not intentional assertions. Unintentional assertions may very well be sincere; sincerity requires only the intention to assert something that is accepted and not that this intention is successful. Thus you may sincerely utter something you do not accept; however, you cannot make a sincere and *intentional* assertion of something you do not accept, as acceptance is here understood.

There is no infallible test for whether or not a person accepts H. For one thing, there is no infallible test of whether an assertion is sincere and intentional. And even if we had an infallible test of sincerity and intentionality, that test would not determine whether a person who does not assert H accepts H. But verificationism is now in disrepute, so I suppose I do not have to argue that a concept can be legitimate, even though there is no infallible test for its application. What is more important is that we be able to make reasonable inductive inferences about what a person accepts; this we can do, as the following examples will illustrate.

Categorical assertions predominate in science textbooks. Because these assertions appear to be sincere, and because they usually do not appear to be slips, we can say that the authors of science texts accept the assertions that occur in these books. Presumably their colleagues would also sincerely endorse most of these assertions, for example in teaching; and so they too accept most of what is in textbooks of their subject.[1]

In their research articles, scientists often mix categorical assertions with more guarded statements. Here is an example,

[1] Some texts present an outmoded theory (e.g., classical mechanics), because it is easier for students than the theory that is actually accepted (e.g., relativity theory). Often this is done by describing the theory, without asserting it. Where the outmoded theory is asserted, the assertion must be deemed insincere, even if pedagogically sound.

taken almost at random, from an article in a recent issue of *Nature* (15 June 1989 issue).

> We found that this enzyme RNA catalyses the site-specific attack of guanosine on the isolated $P1$ stem, but that the K_m for free $P1$ was very high (> 0.1mM). This weak interaction probably reflects the fact that there are few sequence or size requirements for the recognition of $P1$ by the core intron.

Here is another example, from the same issue of *Nature*.

> In conclusion, Greenland ice cores reveal abrupt and radical changes in the North Atlantic region during the Younger Dryas–Pre-Boreal transition, including decreased storminess, a 50% increase in the precipitation rate, a 7°C warming, and probably a temporary decrease of the evaporation temperature in the source area of moisture ultimately precipitated as snow at high elevation in the arctic.

In each of these quotations, categorical assertions are followed by statements that do not make a categorical assertion but rather express a judgment of probability. From these statements we can reasonably infer that the authors accept the assertions they have made categorically, but we cannot conclude that they accept the hypotheses they merely say to be probable. When scientists refrain from categorically asserting a hypothesis in a publication, this does not necessarily mean that they do not accept the hypothesis. They may accept the hypothesis themselves, but refrain from categorically asserting it in a publication because the evidence is not yet enough to convince the scientific community at large. In such a case, the scientists may say in private that they believe the hypothesis, and this would be strong evidence that they accept it. The point is that if there is any context in which a person would assert a hypothesis, then we can reasonably infer that the person accepts that hypothesis (provided it is reasonable to think that the person is sincere and not making a mistake in that context, and that the change of context would not change what the person accepts).

Acceptance of H is the state expressed by sincere intentional assertion of H. But what sort of thing is H here? Although assertions are made by means of sentences, I do not intend that H should be understood to be a sentence. I would say that what

a German speaker asserts by saying 'Der Schnee ist weiss' is the same as what an English speaker asserts by saying 'Snow is white', though the sentences are different. I would also say that the state of acceptance these assertions express (if sincere and intentional) is the same. Conversely, a sentence like 'This is an electron' may be used to make quite different assertions on different occasions (possibly assertions with different truth values), even though the sentence is the same; here the state of acceptance that the assertions express (if sincere and intentional) is also different.

I therefore take H to be a proposition, rather than a sentence. Propositions are here understood as sets of states (they are also referred to as events). Thus we identify what a person has asserted on a given occasion with the set of states that are consistent with the assertion. For example, an utterance of 'Snow is white' (or 'Der Schnee ist weiss') asserts that the true state is a member of the class of all states in which snow is white. And an utterance of 'This is an electron' asserts that the true state is a member of that class of states in which the thing denoted by 'this' is an electron.

6.2 Acceptance and probability

6.2.1 Probability 1 not necessary

What is the relation between acceptance and probability? One suggestion would be to identify acceptance of a hypothesis with assignment of probability 1 to that hypothesis. But this view is untenable. For to give hypothesis H probability 1 is to be willing to bet on it at any odds; for example, a person who gave H probability 1 would be willing to accept a bet in which the person wins a penny if H is true, and dies a horrible death if H is false. I think it is clear that scientists are not usually this confident of the hypotheses they sincerely categorically assert, and thus that probability 1 is not a necessary condition for acceptance.[2]

[2] I argued that this is so in (Maher 1986b) and (1990c, n. 13). Because I think few readers will need convincing, I do not repeat those arguments here.

6.2.2 High probability not sufficient

Having discarded the idea that acceptance can be identified with probability 1, we might try identifying it with "high" probability, that is, probability greater than (or not less than) some value $r < 1$. But this also would be a mistake. For example, consider a lotto game in which six numbers between 1 and 50 are drawn without replacement. If you purchase a ticket in such a game, your chance of winning is less than one in ten million. Now suppose you are thinking of purchasing such a ticket; you have selected six numbers, and you ask me whether I think this combination will win. I certainly would not say it *will* win, but I also would not say it *won't* win. I would tell you the odds against it winning are enormous, but no greater than for any other number. Nor would I be being less than forthcoming here; I simply would not accept the proposition that your numbers will not win, even though I give this proposition an enormously high probability. Hence even an extremely high probability is not sufficient for acceptance.

The claim of the preceding paragraph is that it is logically possible to give a proposition an extremely high probability, yet not accept it. Even if this is conceded, it might still be claimed that it is irrational not to accept a proposition with an extremely high probability. But I think this claim should also be rejected. I do not think the stance I adopted in the preceding paragraph is irrational. Furthermore, the claim conflicts with what I take to be a more compelling principle of rationality.

If I were to accept everything that I think extremely probable, then for every set of six lotto numbers, I would accept that this set would not win. However, I also accept that some set of six lotto numbers would win. Thus I accept an inconsistent set of propositions. Hence if rationality required me to accept all the propositions I regard as highly probable, it would require me to accept an inconsistent set of propositions. But we naturally suppose that rationality requires consistency.

That natural supposition needs some qualification to be defensible. It may be that there is some nonobvious inconsistency in the propositions I accept, but the only way for me to remove it would be to spend years investigating the logical relations among the things I accept, or else to give up much of what I

accept. Either way, the cost of achieving consistency seems more than it is worth, and so it is rational for me to continue with the inconsistency.

However, the lotto example I was discussing is not like this. Here I have a clearly identified inconsistency, which can easily be removed. In such a situation, I think rationality requires consistency. Our everyday practice suggests that this endorsement of consistency is widely shared. Show people that various things they have said are inconsistent, and they will feel obliged to retract one of those propositions. But if rationality requires consistency, even if only in cases where inconsistency is easily avoidable, then rationality cannot require accepting all propositions with high probability.[3]

6.2.3 Probability 1 not sufficient

I have argued that no probability short of 1 is sufficient for acceptance, while a probability of 1 is not necessary for acceptance. Together, these results show that acceptance cannot be identified with any level of probability.

Still, it may seem that if a proposition is given a probability of 1, then it must be accepted. But this also is false. Suppose Professor Milli's probabilities concerning the charge of the electron, e, are as depicted in Figure 6.1. Here $f(e)$ denotes Milli's "probability density" for e. What this means is that for any numbers a and b, Milli's probability that e lies between a and b is the area under the curve $f(e)$ from a to b. As we can see from the figure, for any distinct values a and b, Milli gives a positive probability to e being between a and b. However, for any number a, Milli's probability that $e = a$ is zero. (The "area" under the curve $f(e)$ from a to a has zero width, and hence is zero.) Since Milli's probabilities satisfy the axioms of probability, it follows that for every number a, Milli gives probability 1

[3] The clash of principles here is the core of the "lottery paradox", which seems to have been introduced to the literature by Kyburg (1961, p. 197). Kyburg took the opposite stance to the one I am taking; he upheld the principle that high probability is sufficient for acceptance, and so rejected consistency. The difference between him and me may derive from the fact that his conception of acceptance differs from mine, being tied to practical action. For a critique of Jonathan Cohen's (1977) attempt to resolve the paradox with "inductive" probabilities, see Maher (1986b, p. 375f.).

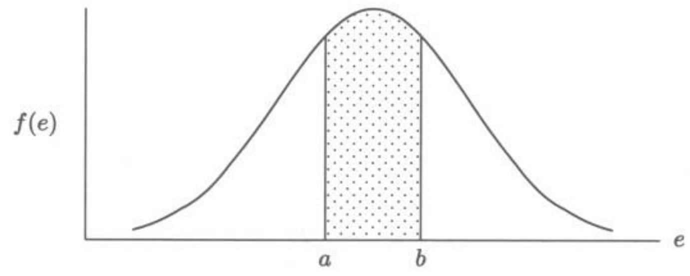

Figure 6.1: Milli's probability distribution for e

to the proposition that $e \neq a$. So if probability 1 were sufficient for acceptance, then for each number a, Milli must accept that $e \neq a$.

But suppose Milli says: "My conclusion is that the charge on the electron is $4.774 \pm .009 \times 10^{-10}$ electrostatic units." Pressed further, he refuses to rule out any value of e in this interval. Then absent any indications that Milli is being less than forthcoming, I think we should allow that for values of a in the stated interval, Milli does not accept that $e \neq a$. And this need not undermine our attribution to Milli of the probability density function f, that gives probability 1 to all propositions of the form $e \neq a$. Thus probability 1 is not sufficient for acceptance.

If you have gone this far with me, you concede that it is logically possible for a person to give a proposition probability 1, yet not accept it. This still leaves open the possibility that probability 1 rationally obliges one to accept a proposition, even if it does not logically force it. But I think there is also no such rational obligation. Indeed, the idea that there is such an obligation, like the parallel idea considered in the preceding section, conflicts with the principle of consistency.

The conflict is evident in Milli's case. If Milli were to accept every proposition to which he gives probability 1, then for all a, he would accept that $e \neq a$. Yet Milli is also sure that the electron has some value or other, and hence must accept this too. Thus if probability 1 is sufficient for acceptance to be rationally required, Milli is rationally obliged to accept an inconsistent set of propositions. Yet there is nothing pathological about Milli's

probability function; probability functions like this are called "normal" in statistics. The problem lies rather in the idea that probability 1 is sufficient for rational acceptance.

6.2.4 High probability not necessary

For all that I have said so far, it could be that high probability is necessary for acceptance. Specifically, it may be thought that a proposition cannot be accepted without giving it a probability of at least 1/2. While I would agree that accepted propositions commonly do have a probability greater than 1/2, I do not think that this must always be the case.

It has often been noted that, in view of the regularity with which past scientific theories have been overthrown, we can reasonably infer that current scientific theories will also be overthrown. Thus anyone reflecting on the history of science ought to give a low probability (less than 1/2) to any given significant current theory being literally correct. Yet scientists continue to sincerely assert significant scientific theories. If high probability were necessary for acceptance, we would have to say that either

(a) These scientists have not drawn the appropriate lesson from the history of science, and accord their theories an unreasonably high probability; or
(b) Contrary to appearances, these scientists do not actually accept their theories.

While (a) might be true in many cases, we could have good evidence that it is false for some scientists. Suppose Einstein were offered his choice of

(1) World peace if general relativity is completely correct; or
(2) World peace if general relativity is false in some way.

I think Einstein probably would have chosen (2). In any case, let us suppose that he makes this choice. Then it would be reasonable to conclude (using the preference interpretation of probability) that his probability for general relativity is less than 1/2.

If high probability is necessary for acceptance, we then would have to say that Einstein did not accept general relativity. But suppose he does categorically assert the theory, at the same time that he is choosing (2) over (1); this is certainly possible. It is also possible that we could satisfy ourselves that no slip of the tongue occurred; his assertion was intentional. Then someone who holds that high probability is necessary for acceptance must say that Einstein's assertion is not sincere. "Actions speak louder than words," it will be said.

But suppose Einstein defends himself against this charge of insincerity. "General relativity is simple in conception, follows from attractive assumptions, and its empirical predictions have been successful. Thus although the theory is probably incorrect in some way, I am confident that it will live on as a limiting case of a future theory. And in the meantime, it is the only theory that fits the evidence. Thus for the time being I accept that the world is as the theory says, though I realize that the theory is likely to be corrected in the future." With such an explanation, and absent any other reasons to question sincerity, I think we should accept that Einstein is sincere in asserting general relativity, even though giving it a low probability of being completely correct. If someone thinks otherwise, their conception of sincerity is different from mine.

There may be a temptation to say that what Einstein really accepts is not general relativity itself, but rather the proposition that general relativity is *approximately* correct; that it will "live on as a limiting case of a future theory." But I am supposing that, as scientists usually do, Einstein is categorically asserting the theory itself, not merely that the theory is approximately correct. That being so, what he accepts, on my definition of acceptance, is the theory itself, not merely the claim that the theory is approximately correct.

The example I have been using is largely fictitious, though I hope it is not too implausible a fiction. In any case, the mere possibility of the story is enough to establish that it is possible to accept a proposition without giving it a probability as high as 1/2.

But even if acceptance is possible in this case, can it be rational? I submit that the example at hand supports an affirmative answer. It is certainly rationally permissible (if not obligatory) to give major scientific theories a low probability of

being literally correct. But we also pretheoretically suppose that it is rational to accept our best current scientific theories. If these things seem in conflict, I would suggest the cause is an inadequate theory of rational acceptance. In the next section, I will describe a theory that lets us say all the things we want to say here.

We don't need to go to grand scientific theories to find examples in which acceptance of propositions, while giving them probability less than 1/2, can be rational. I think that for almost everyone, there are ten propositions that the person accepts, but to whose conjunction the person would give probability less than 1/2. The propositions can be mundane things like 'My desk is made of oak', 'My wife is in the living room', 'Tom is a psychology major', and so on. In such cases, giving the conjunction a probability less than 1/2 seems reasonable; this follows from the assumption that we are not much more than 90 percent confident of each, and their probabilities are independent. Accepting the individual propositions is also pretheoretically reasonable. So someone who maintains that a probability greater than 1/2 is necessary for rational acceptance must say that rationality does not require people to accept the logical consequences of what they accept. But this logical principle is one on which we rely every day, so abandoning it is a high price to pay.

6.3 Rational acceptance

6.3.1 Theory

I have argued that high probability is neither necessary nor sufficient for rational acceptance of a hypothesis. If this is right, rational acceptance must depend on something other than probability.

To see what this something else is, consider the conclusion Cavendish drew from an experiment he conducted in 1773. The experiment was to determine how the electrostatic force between charged particles varies with the distance between the particles. Cavendish states his conclusion this way:

> We may therefore conclude that the electric attraction and repulsion must be inversely as some power of the distance between that of the $2+1/50$th and that of the $2-1/50$th, and there is no reason to think

that it differs at all from the inverse duplicate ratio. (Cavendish 1879, pp. 111–2)

This statement indicates that Cavendish accepted

H_C: The electrostatic force falls off as the nth power of the distance, for some n between 1.98 and 2.02.

Why wouldn't Cavendish have accepted only a weaker conclusion, for example by broadening the range of possible values of n, as in

H'_C: The electrostatic force falls off as the nth power of the distance, for some n between 1.9 and 2.1.

Or he could have made his conclusion conditional, as in

H''_C: If the electrostatic force falls off as the nth power of the distance, for some n, then n is between 1.98 and 2.02.

Both H'_C and H''_C are more probable than the conclusion H_C that Cavendish actually drew, as are infinitely many other weaker versions of Cavendish's hypothesis. The obvious suggestion is that although these weaker hypotheses are more probable than H_C, they are also considerably less informative, and that is why Cavendish did not limit himself to these weaker hypotheses. But if informativeness is what is wanted, why not accept a stronger hypothesis, a natural one here being

H'''_C: The electrostatic force falls off as the second power of the distance.

And again there is an obvious answer: Although H'''_C is more informative than H_C, Cavendish felt that it was not sufficiently probable to accept.

These considerations suggest that acceptance involves a trade-off of two competing considerations: the concern to be right (which would lead one to accept hypotheses of high probability), and the desire for informative hypotheses (which tends to favor hypotheses of low probability). Thus a theory of acceptance needs to take into account the scientist's goals or values, and specifically the relative weights put on the goals of truth and informativeness.

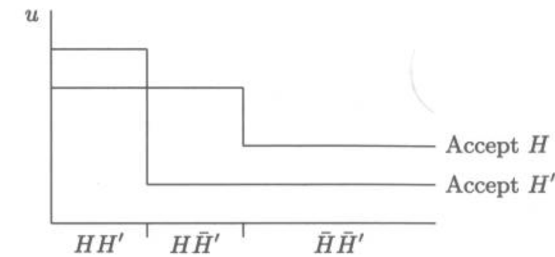

Figure 6.2: A simple conception of the utility of accepting hypotheses

Since the Bayesian way to represent goals is with a utility function, a Bayesian theory of acceptance thus requires there to be a utility function representing the weights that the scientist puts on the competing goals of truth and informativeness.

What would such a utility function be like? First we need to identify the consequences that are the domain of the utility function. A simple suggestion (Hempel 1960, 1962; Levi 1967) is that acceptance of a hypothesis H has two possible consequences, which we could describe as "accepting H when it is true," and "accepting H when it is false." The utility function would assign utilities to consequences such as these, and the goal of truth would be represented by giving a higher utility to the former consequence than to the latter. Then the utility of accepting H would be as in Figure 6.2.

If H' is a logically stronger hypothesis than H, then the goal of accepting *informative* true hypotheses would be represented by giving higher utility to the consequence "accepting H' when it is true" than to "accepting H when it is true." This is also diagrammed in Figure 6.2. (In this figure, the utility of accepting a false hypothesis is represented as being higher for the less informative hypothesis; but I do not insist that this must be the case.)

Figure 6.2 assumes that the utility of accepting a given hypothesis depends only on the truth value of the hypothesis, and thus has only two possible values. But it has often been said that false theories can be more or less close to the truth; or as Popper puts it, some false theories have greater *verisimilitude* than others. And it is held that the utility of accepting

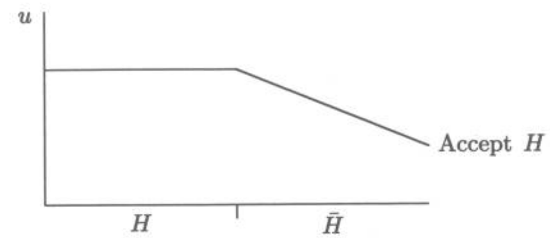

Figure 6.3: Utility of acceptance depending on verisimilitude

a false theory depends on this distance from the truth, being higher the closer the hypothesis is to the truth. On this view, not only will the utility of accepting a hypothesis depend on whether the hypothesis is true or false, but also, in case the hypothesis is false, it will depend on how far from the truth the hypothesis is. So a plot of the utility of acceptance might look like Figure 6.3.

Up to this point I have been discussing the acceptance of a single hypothesis. But scientists accept many hypotheses, and the utility of accepting two hypotheses will not in general be the sum of the utilities of accepting each. However, we can reduce the general case to the case of accepting a single proposition by focusing on the total corpus of accepted propositions. For this corpus can be represented by a single proposition K, namely the conjunction of all propositions that the person accepts. The acceptance of a new proposition H, without any change in those previously accepted, can then be represented as a replacement of the total corpus K by the logically stronger corpus consisting of the conjunction of K and H. The abandonment of a previously accepted proposition, and the replacement of one accepted proposition by another, can also be represented as shifts from one corpus to another. From this holistic perspective, a scientist's goals can be represented by assigning utilities to the possible consequences of accepting K as corpus, for each possible corpus K. I will refer to a function of this sort as a *cognitive utility function*, since it assigns utilities to cognitive consequences.

Chapter 8 will show how a cognitive utility function can be defined, consistently with the preference interpretation of probability and utility. In the meantime, I will anticipate that result

and assume that a cognitive utility function has been defined. We can then define the *expected cognitive utility* of accepting a corpus in the usual way, as the probability-weighted sum of the utilities of the possible consequences. So if $u(K, x)$ denotes the cognitive utility of accepting K as corpus in state x, and if the set of all possible states is X, then the expected cognitive utility of accepting K as corpus is $\sum_{x \in X} p(x)u(K, x)$, assuming X is countable. We can then say that acceptance of corpus K is rational just in case the expected cognitive utility of accepting this corpus is at least as great as that of accepting any other available corpus.

6.3.2 *Example*

To illustrate this theory of rational acceptance, consider the problem of estimating the true value of some real-valued parameter. For example, the problem might be to estimate the true value of the charge on the electron. The problem can be phrased, in our terminology, as: For what set A of real numbers should I accept (as corpus) that the true value of the parameter is in A? Since the hypotheses in which we are interested differ only in the set A, it will be convenient to identify the set with the corresponding proposition.

Let us define the *content* of (the proposition that the true value is in) A as[4]

$$c(A) = \begin{cases} \frac{1}{1+\sup(A)-\inf(A)} & \text{if } A \neq \emptyset \\ 1 & \text{if } A = \emptyset. \end{cases}$$

So, for example, the content of a set containing just a single point is 1, and it declines as the set A is enlarged, reaching 0 when A is the whole real line $(-\infty, \infty)$. Also, let A's *distance from the truth*, when the true value of the parameter is r, be defined as

$$d_r(A) = \begin{cases} \frac{\inf_{x \in A}|x-r|}{1+\inf_{x \in A}|x-r|} & \text{if } A \neq \emptyset \\ 1 & \text{if } A = \emptyset. \end{cases}$$

[4]Here $\sup(A)$ denotes the least upper bound, or supremum, of A; and $\inf(A)$ denotes the greatest lower bound, or infimum, of A. If A is the interval (a, b) or $[a, b]$, then $\sup(A) = b$ and $\inf(A) = a$.

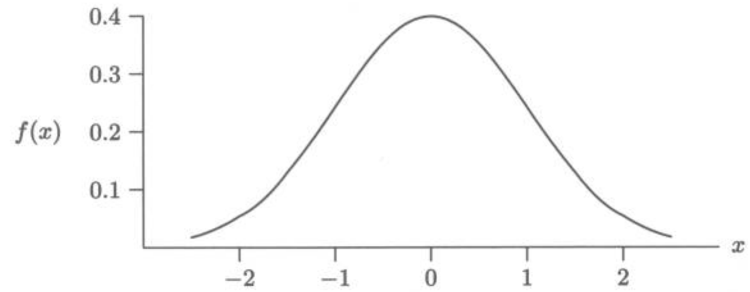

Figure 6.4: The standard normal distribution

If A is true (i.e., $r \in A$), then $d_r(A) = 0$. As the distance between r and the closest point in A increases, $d_r(A)$ increases, approaching 1 in the limit.

Now for any real number k, define

$$u_k(A, r) = kc(A) - d_r(A).$$

For each k, u_k represents a possible measure of the cognitive utility of accepting A when the true state is r; and k represents the relative weight this utility function puts on the desideratum of informativeness, as compared with the competing desideratum of avoidance of error. I believe that there are other possible cognitive utility functions besides the u_k; I define these functions here merely to give an example.

To complete the stipulations of this example, suppose that some scientist's probability distribution for the parameter is the standard normal distribution; this means that for any set A, the probability that the true value r is in A equals the area above the set A under the curve $f(x)$ in Figure 6.4.[5] Milli might have this distribution if r is a multiple of $e - 4.774 \times 10^{-10}$.

If cognitive utility is measured by u_k, for some $k < 1$, then expected cognitive utility can always be maximized by accepting as corpus a closed interval of the form $[-a, a]$, for some a.

[5]For those who want the formula, it is

$$p(A) = \frac{1}{\sqrt{2\pi}} \int_A e^{-x^2/2}\, dx.$$

The set A must be measurable.

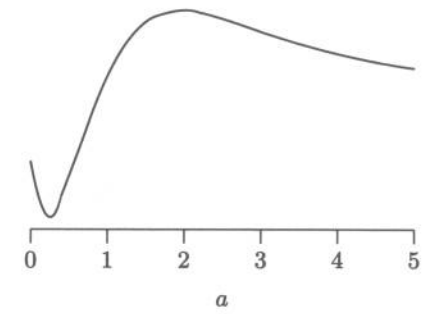

Figure 6.5: $\mathcal{E}(u_{.37}[-a, a])$ plotted as a function of a

Proof. For any $A \neq \emptyset$, the interval $[\inf(A), \sup(A)]$ has the same content as A, and cannot have a larger error. Also,

$$u_k(\emptyset, r) = k - 1 < 0 = u_k((-\infty, \infty), r),$$

and so expected utility is not maximized by accepting $\emptyset$. Thus expected utility can always be maximized by accepting a corpus that is a closed interval. Since the standard normal distribution is symmetric about 0, the closed intervals that maximize expected utility have the form $[-a, a]$.

Thus we lose no real generality if we assume that the sets A that are candidates for acceptance are all closed intervals of the form $[-a, a]$.

Figure 6.5 shows the expected utility of accepting $[-a, a]$, plotted as a function of a, when $k = .37$.[6] Expected utility is maximized at $a = 1.95$, and thus it is rational to accept as corpus that the true value of the parameter is in the interval $[-1.95, 1.95]$. The probability that the true value is indeed in

[6]I use $\mathcal{E}$ to denote expected value. The formula is

$$\begin{aligned} \mathcal{E}(u_k[-a, a]) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} u_k([-a, a], x) e^{-x^2/2}\, dx \\ &= \frac{k}{1 + 2a} - \sqrt{\frac{2}{\pi}} \int_a^{\infty} \frac{x - a}{1 + x - a} e^{-x^2/2}\, dx. \end{aligned}$$

this interval is .95, and thus the interval $[-1.95, 1.95]$ can be thought of as the Bayesian analog of a 95 percent confidence interval.

Note that in accepting $[-1.95, 1.95]$ as corpus, one suspends judgment on propositions that are more probable than what is accepted. For example, to accept $[-1.95, 1.95]$ as corpus is to suspend judgment on whether the value is outside the interval $[-.01, .01]$; and the probability of the latter is .994. Despite its higher probability, acceptance of the latter proposition (either as corpus, or in conjunction with $[-1.95, 1.95]$) would reduce expected cognitive utility. Similarly for the negation of even smaller intervals, which have even higher probability. In accepting $[-1.95, 1.95]$, one also fails to accept that the value of the parameter is not precisely 0, though that proposition has probability 1. Unlike the interval cases, addition of this proposition to $[-1.95, 1.95]$ would not *reduce* expected utility, because it does not reduce the content of what is accepted or increase the possible error; however, it does not increase expected utility either, so there is no positive reason to accept it. Furthermore, since the set of all propositions with probability 1 is inconsistent, to accept all propositions with probability 1 is to accept the empty set $\emptyset$ as corpus, which does not maximize expected utility (see preceding proof). Thus the present theory of rational acceptance supports my earlier claim that no level of probability is sufficient for acceptance.

As the value of k is increased, the cognitive utility function u_k puts more weight on content. Figure 6.6 shows $\mathcal{E}(u_{.5}[-a, a])$, plotted as a function of a, assuming again that the parameter has a standard normal probability distribution. Here the maximum expected utility is attained when $a = 0$, and thus expected utility is maximized by taking as corpus the degenerate interval $[-0, 0]$, that is, the singleton set $\{0\}$. Since this set has probability 0, one sees that *expected cognitive utility can be maximized by accepting a proposition with zero probability.* The reason is that although there is no chance of the corpus being literally correct, there is a very good probability that it will be quite close to the truth, and this together with the desire for an informative corpus makes acceptance of the corpus optimal. This dramatically supports my earlier intuitive argument, that

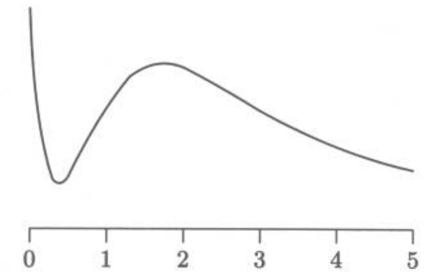

Figure 6.6: $\mathcal{E}(u_{.5}[-a, a])$ plotted as a function of a

high probability is not a necessary condition for accepting a hypothesis.[7]

6.3.3 *Objections*

Epistemologists have stock objections to decision-theoretic accounts of rational acceptance. One is that we don't have probabilities and utilities; the answer to this is in Sections 1.3 and 1.5. Another is that expected utility calculations are difficult to do; the answer to this is in Section 1.2. Here I wish to answer two other objections that are often raised, and have not been addressed earlier in this book.

The first of these objections is this: The acts that are evaluated in decision theory must be options that you could choose if you wanted to. But a doxastic state like acceptance is not something you can choose at will, like deciding to take an umbrella. Hence the application of decision theory to acceptance is fundamentally misguided.

To see the answer to this objection, we need to think about the role norms of rationality play. These norms get their point because our acceptance of them tends to influence what we do. Thus there is no point having norms requiring us to do f, if f is so far beyond our power that acceptance of this norm could have no tendency to make us do f. For example, it would be

[7] Niiniluoto (1986) gives a similar Bayesian account of interval and point estimation.

futile for me to accept the norm 'Have blue eyes,' given that doing so would have no tendency to make me have blue eyes. Typically, the things that acceptance of a norm cannot get us to do are the things that are not subject to our will; and hence we get the principle that 'ought' implies 'can'. However, the case of acceptance is not typical in this respect. While acceptance is not usually directly subject to the will, it nevertheless is the case that acceptance of norms governing acceptance influences what propositions we accept. We see this every day in science, for example where scientists determine a 95 percent confidence interval and then accept that the true value of the variable of interest lies in this interval. If the scientists accepted a different normative statistical theory, they would draw a different conclusion. Hence there is a point to saying that we *ought* to accept this or that proposition, even if acceptance is not directly subject to the will.

Thus for the purposes of cognitive decision theory, we take the options to be not alternatives that are directly subject to the will but alternatives that we would (or could) take if our norms required this. It seems that we could accept any proposition we can formulate, if we accepted a norm requiring this. Hence there is no shortage of options to which to apply cognitive decision theory.

The second objection I wish to address is this: Suppose acceptance of A maximizes expected utility, and that I do accept it. Then according to the theory I have outlined, I have done the rational thing. But suppose that I accepted A for some irrelevant reason; perhaps I just picked it randomly and it was a fluke that I happened to pick the proposition that maximizes expected cognitive utility. Then, the objection claims, my acceptance of A was irrational, even though it maximizes expected utility.

This objection has an air of unreality about it, since I am not likely to be able to accept something for a reason that I think is irrelevant. And if we want to make the case one in which I do not think my reason is irrelevant, but ought to, then my irrationality can be attributed to having irrational probabilities and/or utilities (Section 1.9).

But putting aside such cavils, let us suppose that I can choose what to accept by a process that I regard as random. I hold

that this would not be irrational unless the random process had some positive chance of resulting in me accepting something with lower expected utility than A. But in that case, the expected utility of randomly choosing what to accept is less than the expected utility of directly accepting A. This is so even if, as luck would have it, the random method results in me accepting A. For the expected utility of randomly deciding what to accept is a mixture of the expected utilities of the various propositions that I might accept by this method, and in the case we are dealing with, this must be less than the expected utility of A. Thus my theory of rational acceptance adequately explains the irrationality involved in accepting A as the result of a random process.

6.4 Acceptance and Action

My definition of acceptance linked acceptance of H to assertion of H. Is there also a relation between acceptance and action? Only an indirect one, I shall argue.

Note first that acceptance of H does not commit you to act in all cases as if H were true. For to act as if H were true is to act in the way that would be optimal if H were indeed true (where what counts as "optimal" is determined by your utilities). Thus if you were committed to acting in all cases as if H were true, you would be committed to accepting any bet on H, since you will win such a bet if H is true. But I have already noted that you can accept a proposition without being willing to bet on it at any odds.

Conversely, you can be willing to act in all cases as if H were true yet not accept H. For we have seen that you can give H a probability of 1 yet not accept H; and if you give H a probability of 1, you are willing to act in all cases as if H were true.

Nor is it the case that acceptance of H commits you to betting on H when the odds exceed some threshold. For to be willing to bet on H at odds greater than $m : n$ is to have a probability for H greater than $m/(m+n)$; and we have seen that no positive level of probability is necessary for acceptance.

There is thus no necessary link between what a person accepts and how the person acts in practical circumstances. In any situation, acceptance of H is consistent both with acting as if H were true and also with acting as if H were false, and nonacceptance of H is also compatible with acting either way. This reflects the fact that rational action is determined by probabilities (plus utilities), and acceptance is not identifiable with any level of probability.

The lack of linkage between acceptance and action may also be brought out by imagining two rational individuals who have the same probability distributions, and whose utility functions agree on practical consequences, but who assign different utilities to some cognitive consequences. Then both individuals will have identical preferences regarding all practical actions but may nevertheless accept different propositions. Conversely, if their utility functions agree on cognitive but not practical consequences, they will accept the same propositions but have different preferences regarding practical actions.

Furthermore, the decision to accept a theory (or to not accept it) normally should produce no change in one's willingness to act as if the theory were true in practical contexts. If I rationally accept H, then H maximizes expected cognitive utility, relative to my current probability and cognitive utility functions. The fact that I have accepted H would not normally give me any reason to further increase the probability of H. For example, the reasons that make me confident the theory of evolution is true do not include the fact that I accept the theory. Thus the decision to accept a theory (or to not accept it) normally should produce no change in the probability of that theory – and hence no change in one's willingness to act as if the theory were true in practical contexts.

If this were not so (i.e., if the acceptance of a theory produced a change in one's willingness to act in accordance with the theory), then acceptance of a theory would have an influence on practical utilities, and this would need to be taken into account in considering whether or not it is rational to accept the hypothesis. The fact that there is normally no such influence means that, normally, practical utilities are irrelevant to the rationality of accepting a hypothesis. That is why, in presenting my theory of acceptance in Section 6.3, I supposed that rational acceptance maximizes expected *cognitive* utility, and ignored practical utilities.

Having just mentioned my acceptance of the theory of evolution, let me indicate what role acceptance has played in our evolution. I have been arguing that acceptance of a hypothesis normally has no influence on practical action; and our survival, insofar as it is up to us, is completely determined by our practical actions. Why then would evolution produce creatures like us, who tend to place considerable importance on accepting truths, and avoiding accepting falsehoods? The reason, I suggest, is that our intellectual curiosity provides us with a motivation to reason, observe, and experiment, and these activities often tend ultimately to favor our practical success. The desire for practical success can also motivate such information gathering, but intellectual curiosity provides an additional, and often more immediate, stimulus.

This does *not* mean that intellectual curiosity is reducible to a desire for practical success, and hence that cognitive utility is a species of practical utility. Since the point is somewhat subtle, let me illustrate it with an unsubtle analogy: the desire for sex. It is clear that we have this desire because individuals who have it are more likely to reproduce. However, the desire for sex is not the same as a desire to reproduce, or else there would be no market for contraceptives. I am suggesting that intellectual curiosity (cognitive utility) is like the desire for sex: Just as the desire for sex increases our motivation to reproduce, so our intellectual curiosity increases our motivation to acquire information; and both of these tend to increase evolutionary fitness – but what the desire is a desire for is not the same as the evolutionary function that the desire serves. It may be that astrophysics is the intellectual equivalent of a contraceptive: something that breaks the connection between our desires and the evolutionary function these desires serve.

Although there is no necessary linkage between acceptance and practical action, we may sometimes be able to make inferences about what people accept on the basis of their practical actions. For people's practical actions give us some information about their probabilities; and from that, together with some assumptions about their cognitive utilities, we can make an inference about what they accept. This is the same process we use to predict people's actions in one practical context, on the basis of their actions in a different practical context.

There is, then, some connection between acceptance and practical action, due to the fact that both partially reflect a person's probability function. But this indirect connection is the only connection there is.

6.5 Belief

At the beginning of this chapter, I mentioned my reservation about the folk concept of belief. We have now reached the point where I can explain the problem with this concept.

It is standardly assumed that you believe H just in case you are willing to act as if H were true. But under what circumstances is this willingness to act supposed to occur? As we observed in the preceding section, to be committed to *always* acting as if H were true is to be willing to bet on H at any odds. However, the usual view seems to be that you do not need to be absolutely certain of H (give it probability 1) in order to believe it. For one thing, it is usually supposed that there is very little we can be rationally certain of, but that we can nevertheless rationally hold beliefs on a wide range of topics. Thus belief in H does not seem to imply a willingness to act as if H were true under all circumstances.

Two responses to this difficulty suggest themselves. One is to say that you believe H just in case you are willing to act as if H were true, *provided the odds are not too high* (where what counts as "too high" remains to be specified). On this account, belief in H would be identified with having a probability for H exceeding some threshold. The other approach is to abandon the idea that belief is a qualitative state that you either have or you don't, and instead say that it comes in degrees. Your degree of belief in H could then be measured by the highest odds at which you would be willing to bet on H. This second suggestion effectively identifies belief with probability; and in fact, Bayesians since Ramsey (1926) have referred to subjective probability as "degree of belief."

Whichever of these suggestions is adopted, one will still have to deal with another aspect of the concept of belief. As I

remarked at the beginning of this chapter, it is standardly assumed that belief in H is the mental state expressed by sincere intentional assertion of H – in other words, that belief is the same as acceptance. But if belief is acceptance, then it cannot be related to probability in either of the ways just considered.

The reason why acceptance cannot be identified with probability greater than some threshold has already been given in Section 6.2. No matter where the threshold is set, you can give H a probability that high but still not accept H; and you also can accept something with a probability lower than the threshold. The reason why acceptance cannot be identified with probability itself is that acceptance (as I have defined it) is a qualitative state that one either has or lacks, while probability is a matter of degree. One might think of dealing with the latter problem by introducing a notion of "degree of acceptance," and attempting to identify that with probability; but this also will not work. On any reasonable understanding of degree of acceptance, a person who (qualitatively) accepts H but not K, ought to count as having a higher degree of acceptance for H than for K; yet we have seen that such a person may well give a lower probability to H than to K.

The upshot, then, is that the folk concept of belief appears to regard belief in H as a single mental state that is expressed both by a willingness to act as if H were true and also by sincere intentional assertion of H; and these are in fact two distinct states.

Stich (1983, pp. 230–7) has claimed that the mental states underlying sincere assertion and practical action need not be the same, and has inferred that the folk concept of belief may not refer to anything. Obviously I agree with Stich that these states need not be the same; for I have argued that they *are not* the same.

Stich cites some psychological research that, he claims, supports the conclusion that the mental states underlying assertion and action are distinct. I would like to be able to cite some empirical support of this kind, for the conclusion I have reached on more a priori grounds. But unfortunately, the research cited by Stich gives no real support to this conclusion (though it also does not disconfirm it).

Consider, for example, a study by Wilson, Hull, and Johnson (1981), discussed by Stich. In this study, subjects were induced to volunteer to visit senior citizens in a nursing home, and then later were asked if they would volunteer to help former mental patients. For one group of subjects, overt pressure from the experimenter to visit the senior citizens was made salient; for another group, it was not. There was a negative correlation between being in the first group and agreeing to help the mental patients. The experimenters hypothesized that the reason for the correlation was that subjects in the second group, feeling less pressured, would tend to infer that they were visiting the nursing home because they were helpful people, and that this inference made them more likely to volunteer to help the mental patients.

After agreeing to visit the nursing home, half the subjects in each group were asked to list all the reasons they could think of that might explain why they agreed to go and to rate their importance. Later, all subjects were asked to rate themselves on various traits relevant to helpfulness. Those who had been asked to list reasons tended to rate themselves as more helpful than those who had not, but they were not more likely to volunteer to help the mental patients.

This study does show that subjects' actual and self-reported helpfulness can be independently manipulated. But it does not follow from this that different types of belief state underlie assertion and action; for subjects who act in helpful ways are not acting on the belief that they are helpful. What follows is rather that *being* helpful and accepting that you are helpful are two different states.

A suitable experiment to show the difference between action-producing and assertion-producing belief states would be to have subjects consider a million-ticket fair lottery. If they are not willing to assert categorically that a given ticket will not win, but are willing to bet on this at high odds, then we have established that (a high degree of) the kind of belief that underlies action is not sufficient for the kind of belief that underlies assertion. Conversely, if we can find propositions that subjects will categorically assert to be true but will not bet high odds on, we have shown that the kind of belief that underlies

assertion is not sufficient for (a high degree of) the kind of belief that underlies action. I have already suggested that, at least for reflective people, scientific theories often fall in the latter category.

6.6 Other concepts of acceptance

The term 'acceptance' is widely used, especially in the philosophy of science; and other authors have given it definitions that differ from the one I have given in this chapter. I will conclude this chapter by comparing my definition with a few of these alternatives.

6.6.1 Kaplan

Mark Kaplan (1981) has advocated a conception of acceptance that is similar in spirit to mine. Like me, Kaplan views rational acceptance as maximizing expected cognitive utility, and he asserts that improbable propositions can be rationally accepted, while probable ones need not be. But he and I define acceptance in different ways.

Kaplan defines 'S accepts H' as meaning

S would defend H were S's sole aim to defend the truth.

By 'defend' Kaplan means categorical assertion; so his definition and mine are related in connecting acceptance to assertion. Also, if S's sole aim were to defend the truth, then S would be sincere; and so the notion of sincerity figures in both definitions.

The main difference between Kaplan's definition and mine is that his definition uses a counterfactual conditional. This results in his definition's having an extension different from mine. For example, on Kaplan's account, only very confused persons could ever accept

A: Defending the truth is not my sole aim.

For if your sole aim were to defend the truth, and you were not confused, you would know A was false, and hence would not defend it. On my account, you could accept A because you could be in the mental state expressed by sincere intentional assertion of A (even though you might not wish to actually assert it).

So far this difference is just a different choice of how to use a word. But when combined with the decision-theoretic account of rational acceptance that Kaplan and I share, it becomes a substantive difference. Suppose that you are in the mental state expressed by sincere intentional assertion of A, but would not defend A were your sole aim to defend the truth, since then you would know A was false. My theory says that you accept A, and hence that your present cognitive utility is higher if A is true than if A is false. Kaplan's theory says that you accept $\bar{A}$, and hence that your present cognitive utility is higher if A is false than if A is true. I think my theory better fits our sense of what your cognitive utility would be in this situation.

Another difference is that Kaplan does not take account of the possibility of unintentional assertions. Suppose that if your sole aim were to defend the truth, then you would attempt to assert H, but you would make a slip, and unintentionally assert H'. Then on Kaplan's account you accept H', but on my account you probably[8] accept H. So on Kaplan's account, your cognitive utility, if you would slip and assert H', is the same as if you would sincerely and intentionally assert H'; while on my account, it is the same as if you were to sincerely and intentionally assert H. Again, I think my account better fits our sense of what your cognitive utility would be in this situation.

6.6.2 *Levi*

Levi (1967) discussed two notions of acceptance, which he called 'acceptance as true' and 'acceptance as evidence'. Levi (1967, p. 25) says that 'accept as true' means the same as 'believe'. However, Levi allows that you can accept H as true, without being willing to act on H in all contexts; he even allows that you can accept H while giving it a low probability. Thus acceptance-as-true departs from the common concept of belief by not being linked with action. Levi does seem to assume, however, that sincere intentional assertion of H expresses a person's acceptance of H as true (1967, pp. 10, 223). In both these

[8]If you are attempting to assert H, and if you are sincere, then on my account you accept H. But in the case we are considering here, this only shows that if your sole aim were to defend the truth, then you would (on my account) accept H. It remains possible that you do not in fact accept H.

respects, acceptance as I have defined it appears to agree with Levi's concept of acceptance-as-true.

Levi also offered a decision-theoretic account of when it is rational to accept a hypothesis as true. Like the account of rational acceptance I have sketched, Levi's account proposed that the cognitive goals of accuracy and content could be represented by a cognitive utility function. However, Levi held that a person might simultaneously have different demands for information, these demands being representable by different utility functions. And for Levi, acceptance-as-true was relative to these demands for information; on his theory, one could accept H as true to satisfy one demand for information, and at the same time accept $\bar{H}$ as true to satisfy another demand for information. This would be rational, according to Levi, provided each acceptance-as-true maximized expected cognitive utility, relative to its own particular cognitive utility function.

My notion of acceptance differs from Levi's concept of acceptance-as-true in not being relative to demands for information. My main reason for not adopting a relativized notion is that rationality requires keeping the various propositions we accept consistent with one another; thus someone who accepts H to satisfy one of their demands for information is obliged to also accept H as an answer to any other demand for information that H may satisfy. We therefore do not need a question-relative notion of acceptance in a theory of rational acceptance.

In Levi's later writings (1976, 1980), the notion of acceptance-as-true has virtually disappeared, and acceptance-as-evidence occupies center stage. It is now acceptance-as-evidence for which Levi gives a decision-theoretic account involving cognitive utility. I have the impression that after (1967), Levi came to view acceptance-as-true as not a real bearer of cognitive utility. On this later view, as I interpret it, to accept H as true relative to some demand for information is merely to be committed to accept H as evidence if (possibly contrary to fact) that demand for information were one's only demand.[9] This interpretation of Levi implies that he too does not now believe that acceptance, in the sense that influences

[9]This interpretation is suggested by (Levi 1976, p. 35, and 1984, p. xiv).

cognitive utility, is relative to demands for information.

But while Levi's notion of acceptance-as-evidence is like my notion of acceptance in not being relativized, it differs from my notion in another crucial respect: To accept a hypothesis as evidence, Levi says, is to assume its truth in all (practical and theoretical) deliberation. Thus if one accepts H as evidence, one must give H probability 1. By contrast, my notion of acceptance, like Levi's notion of acceptance-as-true, allows that one can accept hypotheses without giving them probability 1, and indeed can do so while giving them a probability as low as you like.

I find the concept of acceptance-as-evidence less important than Levi does, because I think people accept-as-evidence much less than he supposes. Levi imagines that observation reports, scientific theories, and other things are often accepted-as-evidence. This means that people should be willing to bet at any odds on a wide range of observation reports and scientific theories. I am not willing to do that, and people I have talked to say they are not willing either. I do not deny that some propositions, even contingent ones, are assigned probability 1; if one is dealing with a real-valued parameter, for example, some contingent propositions must get probability 1. But I take it that the propositions given probability 1 are typically very uninformative ones, like "the force between charged particles does not vary precisely as the rth power of the distance." Observation reports and scientific theories are typically regarded as fallible.

In any event, acceptance as I have defined it is ubiquitous, and is not identifiable with acceptance-as-evidence. For as Levi has observed (1967, p. 10), people can and do sincerely assert propositions they are not willing to bet on at all odds and hence do not take to be certainly true; the propositions thus asserted are accepted in my sense, but not accepted-as-evidence. Thus a theory of acceptance in my sense is needed, whatever one thinks of Levi's concept of acceptance-as-evidence.

6.6.3 Van Fraassen

In *The Scientific Image*, van Fraassen maintained that acceptance of a theory in science involves belief that the theory is

empirically adequate (it agrees with observable phenomena) but not that the theory is true. He also said that this belief in empirical adequacy is only a necessary condition for acceptance, not a sufficient condition. Acceptance, he said, also involves a certain commitment. For scientists, the relevant commitment entails the adoption of a particular kind of research program; and for nonscientists, it entails "willingness to answer questions *ex cathedra*" (van Fraassen 1980, p. 12).

For van Fraassen (1985, pp. 247ff.), belief is the same as probability. So his statement, that acceptance of a theory in science involves belief in its empirical adequacy, must mean that acceptance of a theory in science involves giving a high probability to the theory being empirically adequate. And similarly, the statement that belief in the truth of the theory is not necessary for acceptance in science must mean that one can accept a theory in science without giving a high probability to the theory being true.

Van Fraassen's reference to ex cathedra pronouncements points to a similarity between his concept of acceptance and mine: Both are connected with categorical assertion. However, the connection is less direct on my account. People, including scientists, are sometimes secretive; they can be in the mental state expressed by sincere intentional assertion of H, without being willing to assert H. Thus I would not agree that acceptance entails "willingness to answer questions *ex cathedra*."

On both van Fraassen's view and mine, a person can accept a theory without giving a high probability to the theory being true. However, I disagree with his statement that acceptance of a theory in science involves belief that the theory is empirically adequate. Van Fraassen's reason for saying this is that he thinks (a) science aims at accepting empirically adequate theories, and (b) (rational?) acceptance involves a high probability that the aim of acceptance is served. Proposition (a) is about the aim of science, not the concept of acceptance; I will discuss it in Section 9.6. Here I focus on (b), which is a claim about the nature of (rational) acceptance.

One problem with (b) – and also with (a) – is that the notion of an aim seems to presuppose that there is only one desirable outcome; whereas in fact there may be a continuum of possible outcomes, of varying degrees of desirability. But let us put that problem aside, and suppose that we have a simple case in which we can say that the aim in accepting H is to accept an empirically adequate proposition. Still, rational acceptance of H does not require that H have a high probability of being empirically adequate. For example, suppose that accepting H increases utility by 10 if H is empirically adequate, and reduces utility by 1 otherwise. Then the probability of H being empirically adequate need only be 0.1 for acceptance of H to be rational. Thus (b) is false.

Let us turn now to the question of whether acceptance of a hypothesis commits a scientist to a research program. Van Fraassen says that this commitment to a research program comes about because accepted theories are never complete (1980, p. 12); this suggests that the research program that he takes to be entailed by acceptance is one of extending the scope of the theory. But most scientists extend the scope of only a few theories, if any, and I doubt that van Fraassen wishes to keep the class of theories accepted by scientists this small. Perhaps the commitment to extend the scope of the theory is meant to apply only to those scientists who extend the scope of some theory in the subject area of the accepted theory: The position would be that if you accept H, then if you extend the scope of some theory in the subject area of H, you must extend H. But even this is not a condition that must hold for acceptance as I have defined it. For example, we can easily imagine a scientist Poisson′, who sincerely asserted a corpuscular theory of light, but extended the scope of Fresnel's wave theory of light by deriving a previously unnoticed consequence of the theory, and extended the scope of no other theory of light.[10] Poisson′ would accept a corpuscular theory on my account of acceptance but not on van Fraassen's account.

Even if your sole concern was the scientific one of accepting the right theory, it could still be rational to accept (in my sense) one theory but work on an incompatible theory. To see this, suppose that acceptance of H currently has a higher expected cognitive utility than does the acceptance of H', but there is

[10]The real Poisson did extend the scope of Fresnel's theory, without accepting it (Worrall 1989). But I do not know whether Poisson accepted the corpuscular theory, or whether he extended the scope of that theory.

a chance H' could be developed further into a theory whose acceptance would have a higher expected utility than H. In such a case, it could be rational to develop H' while accepting H.[11]

6.7 Summary

I defined acceptance of H as the mental state expressed by sincere intentional assertion of H. We saw that whether a person will accept H depends not only on the probability of H but also on other factors, such as how informative H is. In the light of this, I proposed a decision-theoretic account of rational acceptance, in which acceptance is viewed as having consequences with cognitive utility, and rational acceptance maximizes expected cognitive utility.

Acceptance, as I have defined it, captures one aspect of the folk notion of belief, while probability is a different, incompatible aspect of that concept. This notion of acceptance is related to the notions of acceptance employed by Kaplan, Levi, and van Fraassen; where it differs from the latter, I have given reasons for preferring my definition.

Two main questions remain to be answered. One is why (if at all) we should think that acceptance is an important concept for the philosophy of science. The other is the question of how to justify the assumption of a utility function for cognitive consequences. These questions will be considered, in this order, in the next two chapters.

[11]Essentially this point has been made by Laudan (1977, pp. 108–13). However, the point seems to be inconsistent with Laudan's notion of acceptance, which is *treating a theory as if it were true* (p. 108).