Bayesian confirmation function which he appropriated. (For a discussion and decisive criticism of Popper's account, see Grünbaum, 1976.)

The Bayesian position has recently been misunderstood to imply that if some evidence is known, then it cannot support any hypothesis, on the grounds that known evidence must have unit probability. That the objection is based on a misunderstanding is shown in Chapter 15, where a number of other criticisms of the Bayesian approach will be rebutted.

f THE RAVENS PARADOX

That evidence supports a hypothesis more the greater the ratio $\frac{P(e \mid h)}{P(e)}$ scotches a famous puzzle first posed by Hempel (1945) and known as the *Paradox of Confirmation* or sometimes as the *Ravens Paradox*. It was called a paradox because its premisses were regarded as extremely plausible, despite their counter-intuitive, or in some versions contradictory, implications, and the reference to ravens stems from the paradigm hypothesis ('All ravens are black') which is frequently used to expound the problem. The difficulty arises from three assumptions about confirmation. They are as follows:

1. Hypotheses of the form 'All *R*s are *B*' are confirmed by the evidence of something that is both *R* and *B*. For example, 'All ravens are black' is confirmed by the observation of a black raven. (Hempel called this Nicod's condition, after the philosopher Jean Nicod.)
2. Logically equivalent hypotheses are confirmed by the same evidence. (This is the Equivalence condition.)
3. Evidence of some object not being *R* does not confirm 'All *R*s are *B*'.

We shall describe an object that is both black and a raven with the term RB. Similarly, a non-black, non-raven will be denoted $\bar{R}\,\bar{B}$. A contradiction arises for the following reasons: an RB confirms 'All *R*s are *B*', on account of the Nicod condition. According to the Equivalence condition, it also confirms 'All non-*B*s are non-*R*s', since the two hypotheses are logically equivalent. But contradicting this, the third condition implies that RB does not confirm 'All non-*B*s are non-*R*s'.

The contradiction may be avoided by revoking the third condition, as is sometimes done. (We shall note later another reason for not holding on to it.) However, although the remaining conditions are compatible, they have a consequence which many philosophers have regarded as blatantly false, namely, that by observing a non-black, non-raven (say, a red herring or a white shoe) one confirms the hypothesis that all ravens are black. (The argument is this: 'All non-*B*s are non-*R*' is equivalent to 'All *R*s are *B*'; according to the Nicod condition, the first is confirmed by $\bar{R}\,\bar{B}$; hence, by the Equivalence condition, so is the second.)

If non-black, non-ravens support the raven hypothesis, this seems to imply the paradoxical result that one could investigate that and other generalisations of a similar form just as well by observing white paper and red ink from the comfort of one's writing desk as by studying ravens on the wing. However, this would be a non sequitur. For the fact that RB and $\bar{R}\,\bar{B}$ both confirm a hypothesis does not imply that they do so with equal force. Once it is recognised that confirmation is a matter of degree, the conclusion is no longer so counter-intuitive, because it is compatible with $\bar{R}\,\bar{B}$ confirming 'All *R*s are *B*', but to a minuscule and negligible degree.

Indeed, most people do have a strong intuition that an RB confirms the ravens hypothesis (h) more than an $\bar{R}\,\bar{B}$. We can appreciate why that might be by consulting Bayes's Theorem as it applies to the two types of datum:

$$\frac{P(h \mid RB)}{P(h)} = \frac{P(RB \mid h)}{P(RB)} \quad \& \quad \frac{P(h \mid \bar{R}\,\bar{B})}{P(h)} = \frac{P(\bar{R}\,\bar{B} \mid h)}{P(\bar{R}\,\bar{B})}$$

These expressions can be simplified. First, $P(RB \mid h) = P(B \mid h \,\&\, R)P(R \mid h) = P(R \mid h) = P(R)$. We arrived at the last equality by assuming that whether some arbitrary object is a raven is independent of the truth of h, which seems plausible to us, at any rate as a good approximation, though Horwich (1982, p. 59) thinks it has no plausibility. By similar reasoning, $P(\bar{R}\,\bar{B} \mid h) = P(\bar{B} \mid h) = P(\bar{B})$. Also $P(RB) = P(B \mid R)P(R)$, and $P(B \mid R) = \sum P(B \mid R \,\&\, \theta)P(\theta \mid R) =$ (assuming independence between θ and R) $\sum P(B \mid R \,\&\, \theta)P(\theta)$, where θ represents possible values of the percentage of ravens in the universe that are black (according to h, of course, $\theta = 1$). Finally, $P(B \mid R \& \theta) = \theta$, for if the percentage of black ravens in the universe is θ, the probability of an arbitrary raven being black is also θ. (This is intuitively correct and is formalised in the so-called Principal Principle, which we shall discuss later.)

Combining all these considerations with the above forms of Bayes's Theorem yields

$$\frac{P(h \mid RB)}{P(h)} = \frac{1}{\Sigma\theta P(\theta)} \quad \& \quad \frac{P(h \mid \bar{R}\,\bar{B})}{P(h)} = \frac{1}{P(\bar{R} \mid \bar{B})}.$$

Consider first the term $P(\bar{R} \mid \bar{B})$. Presumably there are vastly more non-black things in the universe than ravens. So even if no ravens are black, the probability of some object about which we know nothing, except that it is not black, being a non-raven must be very high, indeed, practically 1. Hence, $P(h \mid \bar{R}\,\bar{B}) \simeq P(h)$, and, so, the observation that some object is neither a raven nor black provides very little confirmation for h.

According to the equation above, the degree to which RB confirms h is inversely proportional to $\Sigma\theta P(\theta)$. This means, for example, that if it is initially very probable that all or virtually all ravens are black, then $\Sigma\theta P(\theta)$ would be large and RB would confirm h rather little. While if it is initially relatively probable that most ravens are not black, confirmation could be substantial. Intermediate levels of uncertainty about the proportion of ravens that are black would bring their own levels of confirmation. By contrast, because the class of non-black objects is so much larger than the class of ravens, $\bar{R}\,\bar{B}$ confirms 'All ravens are black' to only a tiny extent, irrespective of $P(\theta)$. Mackie's well-known Bayesian solution to the ravens paradox, which is given in the Exercises section at the end of this chapter, is similar and also depends on an assumed large disparity in the number of non-black objects and ravens.

Our Bayesian working of the raven example appears to support the Nicod condition, with the minor limitation that no confirmation is possible, even with positive instances, when the hypothesis has a prior probability of 1. But a Bayesian approach anticipates the violation of Nicod's condition in other circumstances too. And numerous examples have been suggested as plausible instances of such violations. The first of these seems to be due to Good (1961). We shall use an example

that is taken, with some modification, from Swinburne (1971). The hypothesis under examination is 'All grasshoppers are located outside the County of Yorkshire'. The observation of a grasshopper just beyond the county border is an instance of this generalisation and, according to Nicod, confirms it. But it might be more reasonably argued that since there are no border controls or other obstacles restricting the movement of grasshoppers in that area, the observation of one on the edge of the county increases the probability that others have actually entered and hence undermines the hypothesis. In Bayesian terms, this is a case where, relative to background information, the probability of some datum is reduced by a hypothesis —that is, $P(e \mid h) < P(e)$—which is therefore disconfirmed—in other words, $P(h \mid e) < P(h)$.

A much more striking example where Nicod's conditions break down was invented by Rosenkrantz (1977, p. 35). Three people leave a party, each with a hat. The hypothesis that none of the three has his own hat is confirmed, according to Nicod, by the observation that person 1 has person 2's hat and by the observation that person 2 has person 1's hat. But since there are only three people, the second observation must *refute* the hypothesis, not confirm it.

Our grasshopper example provides an instance where a datum of the type $\bar{R}\,B$ confirms a generalisation of the form 'All Rs are B'. Imagine that an object which looked for all the world like a grasshopper had been found hopping about just outside Yorkshire and that it turned out to be some other sort of insect. The discovery that the object was not a grasshopper would be relatively unlikely unless the grasshopper hypothesis was true (hence, $P(e) < P(e \mid h)$); thus it would confirm that hypothesis. If the deceptively grasshopper-like object were within the county boundary, the same conclusion would follow, though the degree of confirmation would be greater. This shows that 'All Rs are B' may also be confirmed by a datum of the $\bar{R}\,\bar{B}$ type. Hence, the impression that non-Rs never confirm such hypotheses may be dispelled.

Horwich (1982) has argued that the raven hypothesis may be differently confirmed, depending on how the black raven was chosen, either by randomly selecting an object from the population of ravens or by making the selection from the population of black objects. (Horwich denotes the evidence that some object is a black raven as either R^*B or RB^*, depending on whether it was discovered by the first selection process or the second.) Prompted by an unpublished paper by Kevin Korb ("Infinitely Many Resolutions of Hempel's Paradox", 1993), we agree with Horwich that this is so; the Bayesian explanation which Horwich gives is recapitulated, in a slightly different context, in Chapter 14, section **f**.

But Horwich offers another explanation, which fits poorly with his Bayesian one. For he claims that the datum R^*B is always more powerfully confirming than RB^*, because, he says, only it subjects the raven hypothesis to the risk of falsification. But this surely conflates the process of collecting evidence, which may indeed subject the hypothesis to different risks of refutation, with the evidence itself, which either refutes the hypothesis or does not refute it, and in the case of R^*B and RB^*, it does not. (For a fuller discussion of this point, the reader is referred to Chapter 15, section **g**.)

Our conclusions are, first, that the supposedly paradoxical consequences of Nicod's condition and the Equivalence condition are not problematic, and, secondly, that there are separate reasons for rejecting Nicod's condition, which, moreover, conform to Bayesian principles.

g THE DESIGN OF EXPERIMENTS

Why should anyone go to the trouble and expense of performing new experiments and seeking more evidence for hypotheses? The question has been debated recently and is sometimes felt to be something of a problem. Maher (1990) argues that since evidence can neither conclusively verify nor conclusively refute a theory, Popper's scientific aims cannot be served by gathering new data. Since a large part of scientific activity is devoted to the acquisition of new evidence, if Mayer were right, there would appear to be a serious gap in Popper's philosophy. Miller (1991, p. 2) claims that the same difficulty appears in Bayesian philosophy:

> If e is the agent's total evidence, then $P(h \mid e)$ is the value of his probability and that is that. What incentive does he have to change it, for example by obtaining more evidence than he has already? He might do so, enabling his total evidence to advance from e to e+; but in no clear way would $P(h \mid e+)$ be a better evaluation of probability than $P(h \mid e)$ was.

There seems to us, on the contrary, a quite straightforward reason why a Bayesian might seek new evidence, namely, in order to diminish uncertainty about some aspect of the world, in a desire to find out the truth. Suppose, for instance, that the question of interest concerns a particular parameter. You might start out fairly uncertain about its value, in the sense that your probability distribution over its possible values is rather diffuse. A suitable experiment, if successful, would furnish evidence to lessen that uncertainty by changing the probability distribution, via Bayes's Theorem, so that it was now more concentrated in a particular region, the greater the concentration and the smaller the region the better. This criterion has been given a precise, quantitative expression by Lindley (1956), in terms of Shannon's characterisation of information. Lindley showed that in the case where knowledge of a parameter, θ, is sought, provided the density of x varies with θ, any experiment in which x is measured has an expected yield in information. But, of course, this result is compatible with a good experiment (with a high expected information yield) being relatively uninformative in a particular case; similarly, a poor experiment may to one's surprise be relatively informative.

In deciding which experiment to perform, one must also take at least three other factors into account: the cost of the experiment; the morality of carrying it out; and the value, both theoretical and practical, of the hypotheses one is interested in. Bayes's Theorem, of course, implies nothing about how these separate factors should be balanced.

h THE DUHEM PROBLEM

h.1 The Problem

The so-called Duhem (or Duhem-Quine) problem is a problem for theories of science of the type associated with Popper, which emphasise the power of certain evidence to refute a hypothesis. According to Popper's influential views, the characteristic of a theory which makes it 'scientific' is its falsifiabil-