13

Probability and Reliability

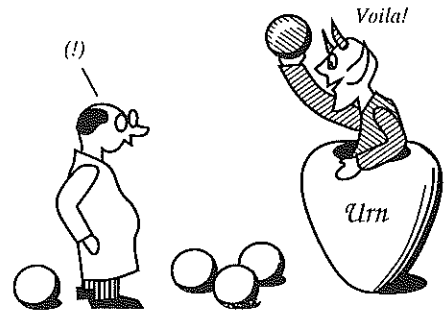

1. Introduction

The results of the preceding chapters make no reference whatever to probability. Since probability and induction have long been viewed as inseparable, it is interesting to relate the probabilistic perspective to the purely logical point of view of the preceding chapters. Of particular interest are the limiting reliability claims made for probabilistic methods.[1] For example, it is often said that the process of updating probabilities by Bayes' theorem will almost surely approach the truth in the limit:

> The person learns by experience. The purpose of the present section is to explore with a moderate degree of generality how he typically becomes almost certain of the truth, when the amount of his experience increases indefinitely. . . . It is to be expected intuitively, and will soon be shown, that under general conditions the person is very sure that after making the observation he will attach a probability of nearly 1 to whichever element of the partition actually obtains.[2]

In light of the many negative results in the preceding chapters, such claims sound too good to be true. Are they? Or do they illustrate the triumph of modern, probabilistic thinking over skepticism? The aim of this chapter is to address these important epistemological questions.

[1] I am indebted to my colleague Teddy Seidenfeld for many useful discussions and references concerning the material in this chapter.
[2] Savage (1972): 46.

2. Conditionalization

Let *Bo* denote the closure of the open sets of $\mathcal{N}$ under complementation and countable union. Thus, *Bo* contains the union of the complexity classes Σ_n^B. Let *P* be a real-valued function taking elements of *Bo* as arguments. *P* is a *probability measure* on *Bo* (or, I will say, on $\mathcal{N}$) just in case

(1) *for each* $S \in Bo$, $P(S) \geq 0$,

(2) $P(\mathcal{N}) = 1$ *and*

(3) If $S_0, S_1, \ldots, S_n$ *is a finite sequence of pairwise disjoint elements of Bo, then* $P\left(\bigcup_{i=1}^{n} S_i\right) = \sum_{i=1}^{n} P(S_i)$.

Whatever probability is, it is supposed to behave sort of like mud spread out on a table marked with areas representing subsets of $\mathcal{N}$. The first two axioms say that we start out with a unit block of mud. The third axiom says that the total mud spread over a finite number of disjoint regions can be found by summing the mud on each region. In other words, the probability of a whole region is the sum of the probabilities of the parts of the region when the region is divided into only finitely many parts. This property is called *finite additivity* (*Fig. 13.1*).

Countable additivity requires, further, that if an event is divided into a countable infinity of nonoverlapping parts, the probability of the whole is still the sum of the probabilities of its parts (*Fig. 13.2*). In other words, if $S_0, S_1, \ldots, S_i, \ldots$ is an ω-sequence of pairwise disjoint elements of *Bo*, then

$$P\left(\bigcup_{i \in \omega} S_i\right) = \sum_{i \in \omega} P(S_i).$$

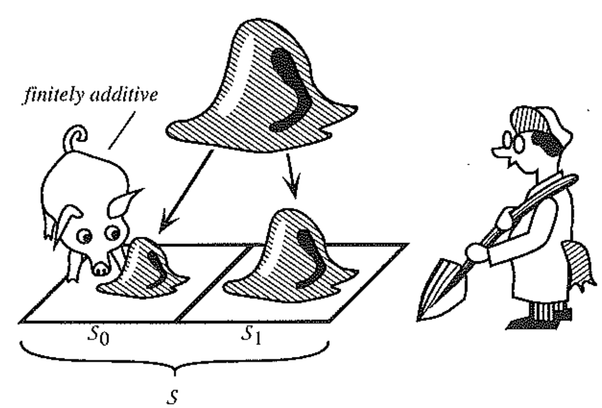

Figure 13.1

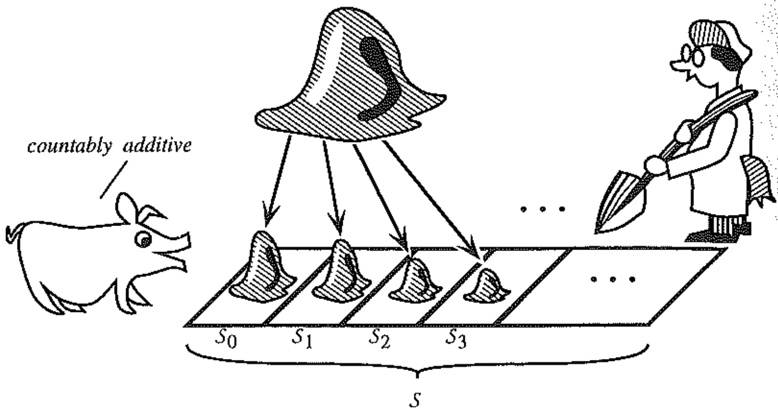

Figure 13.2

P is a *countably additive probability measure* just in case P is a probability measure that is also countably additive. When a probability measure is not countably additive, we say that it is *merely finitely additive*. Countable additivity can be expressed in a different form. Define (*Fig. 13.3*):

P is continuous $\Leftrightarrow$ for each downward nested sequence

$$Q_1 \subseteq Q_2 \subseteq \ldots \subseteq Q_i \subseteq \ldots \textit{ of Borel sets,}$$

$$P\left(\bigcap_{i \in \omega} Q_i\right) = Lim_i P(Q_i).$$

Proposition 13.1

If P is a finitely additive probability measure on $\mathcal{N}$ then P is countably additive $\Leftrightarrow$ P is continuous.

Proof: Exercise 13.1. ■

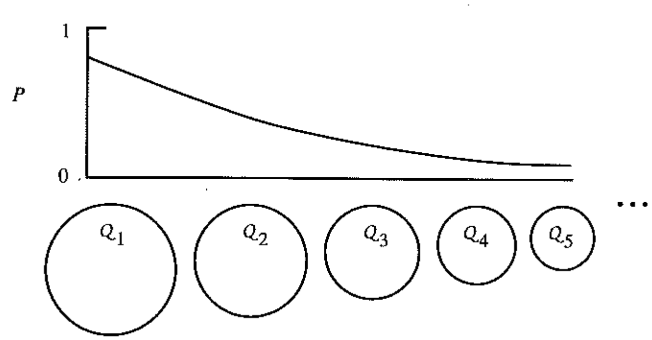

Figure 13.3

If P is a probability measure on $\mathcal{N}$, and $\mathcal{H}$, $\mathcal{E} \subseteq \mathcal{N}$ and $P(\mathcal{E}) > 0$, then the *conditional probability* of $\mathcal{H}$ given $\mathcal{E}$ is given by:

$$P(\mathcal{H}, \mathcal{E}) = \frac{P(\mathcal{H} \cap \mathcal{E})}{P(\mathcal{E})}.$$

It is an immediate consequence of this definition that:

Proposition 13.2 (Bayes' theorem)

If $P(\mathcal{E}) > 0$ then $P(\mathcal{H}, \mathcal{E}) = \dfrac{P(\mathcal{H})P(\mathcal{E}, \mathcal{H})}{P(\mathcal{E})}$ *and hence*

$$P(\mathcal{H} \cap \mathcal{E}) = P(\mathcal{H})P(\mathcal{E}, \mathcal{H}).$$

Proof: Exercise 13.2. ■

According to a popular conception of learning from experience, the scientist starts out with some initial probability measure P_0 and then updates this measure in response to new data by defining $P_n(\mathcal{H})$, the probability of $\mathcal{H}$ at time n, to be $P_0(\mathcal{H}, \mathcal{E})$, the conditional probability of $\mathcal{H}$ given the total evidence $\mathcal{E}$ encountered so far. The process of updating one's current degrees of belief by replacing them with probabilities conditional on all the evidence gathered so far is known as *temporal conditionalization*. Bayes' theorem is often used to compute $P_0(\mathcal{H}, \mathcal{E})$ because a statistical model typically provides an explicit formula for $P_0(\mathcal{E}, \mathcal{H})$ and the Bayesian statistician also concocts an explicit formula for $P(\mathcal{H})$. Thus conditionalization is sometimes referred to as *Bayesian updating*. When Bayes' theorem is used in temporal conditioning, $P(\mathcal{H}, \mathcal{E})$ is called the *posterior probability* of $\mathcal{H}$ given $\mathcal{E}$, $P(\mathcal{H})$ is called the *prior probability* of $\mathcal{H}$, and $P(\mathcal{E}, \mathcal{H})$ is called the *likelihood* of $\mathcal{H}$ given $\mathcal{E}$.

Since I have not defined the value of $P(\mathcal{H}, \mathcal{E})$ when $P(\mathcal{E}) = 0$, one might suppose that Bayesian updating goes into a coma whenever such evidence is encountered (*Fig. 13.4*). In fact, it is possible to extend conditional probability

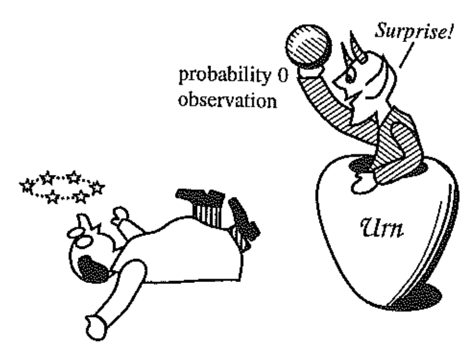

Figure 13.4

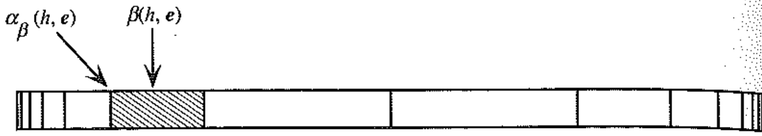

Figure 13.5

to such cases.[3] Having noted this possibility, I will sidestep the complications it raises by considering only conditionalizers that assign nonzero prior probabilities to all relevant hypotheses and to all finite data sequences logically consistent with background knowledge, so that $\mathcal{K}$ *entails* that no finite data sequence of probability zero will ever be seen. Accordingly, define:

β is a conditionalizer for $H, C, \mathcal{K} \Leftrightarrow$ there is a probability measure P on $\mathcal{N}$ such that

(a) $P(\mathcal{K}) = 1$,

(b) *for all* $e \in \omega^*$, *if* $[e] \cap \mathcal{K} \neq \varnothing$ *then* $P(e) > 0$, *and*

(c) $\beta(h, e) = P(C_h, [e])$.

Conditionalizers produce real-valued probabilities and hence are not assessment methods in the sense defined in chapter 3, since assessment methods were defined to conjecture only rationals. Nonetheless, every conditionalization method β induces a rational-valued assessment method α_β with the same limiting performance, as follows. First, digitize the $[0, 1]$ interval by two rational-valued sequences that start at the middle of the interval and converge to its end points, treating the end points as special cases (*Fig. 13.5*):

$$\gamma(r) = \begin{cases} 2^{-i}, \text{ where } i = \text{the least } n > 0 \\ \quad \text{such that } 2^{-(n+1)} \leq r < 2^{-n} & \text{if } 0 < r < 2^{-1} \\ 1 - 2^{-i}, \text{ where } i = \text{the least } n > 0 \\ \quad \text{such that } 1 - 2^{-n} \leq r < 1 - 2^{-(n+1)} & \text{if } 2^{-1} \leq r < 1 \\ r, & \text{if } r = 0 \text{ or } r = 1. \end{cases}$$

Then $\alpha_\beta = \gamma \circ \beta$ is a rational-valued assessment method. It is easy to see

[3] There are applications in which it is both natural and desirable to suppose that $P(\mathcal{H}, \mathcal{E})$ is uniquely determined, even though $P(\mathcal{E}) = 0$. For example, suppose that for each $\theta \in R$, P_θ is a probability measure on $\mathcal{N}$. Then Bayesian statisticians usually let $P(\mathcal{S}, \{\theta\}) = P_\theta(\mathcal{S})$ even though $P(\{\theta\}) = 0$ in the typical case in which P is continuous. The general theory of which this example is an instance is presented, for example, in Billingsly (1986).

that:

Proposition 13.3

Let β be a conditionalization method.

If $\beta \begin{bmatrix} \textit{approaches} \\ \textit{stabilizes to} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ *on ε then so does α_β.*

It follows immediately that:

Corollary 13.4

If h is $\begin{bmatrix} \textit{verifiable}_C \\ \textit{refutable}_C \\ \textit{decidable}_C \end{bmatrix}$ *gradually given $\mathcal{K}$ by a conditionalizer then*

$$C_h \in \begin{bmatrix} \Sigma[\mathcal{K}]^B_3 \\ \Pi[\mathcal{K}]^B_3 \\ \Delta[\mathcal{K}]^B_2 \end{bmatrix}.$$

If h is $\begin{bmatrix} \textit{verifiable}_C \\ \textit{refutable}_C \\ \textit{decidable}_C \end{bmatrix}$ *in the limit given $\mathcal{K}$ by a conditionalizer then*

$$C_h \in \begin{bmatrix} \Sigma[\mathcal{K}]^B_2 \\ \Pi[\mathcal{K}]^B_2 \\ \Delta[\mathcal{K}]^B_2 \end{bmatrix}. \quad \blacksquare$$

Conditionalization, therefore, cannot evade the negative results of the preceding chapters. Whatever conditionalization does, it cannot do the impossible.

It remains to ask whether conditionalization does what is possible. In fact, conditionalization can fail to solve inductive problems that are fairly trivial from the point of view of the preceding chapters. It follows from Bayes' theorem that if $P(C_h) = 0$, then for all evidence e such that $P(e) > 0$, $P(C_h, e) = 0$. So if the prior probability of C_h is 0, then a conditionalizer employing this measure will fail to approach the truth when C_h is true. But this sort of trivial failure can arise even when $P(C_h) > 0$, for example if $P(C_h, e) = 0$ for some e logically consistent with C_h and $\mathcal{K}$. Nobody would be surprised to see conditionalization fail when the underlying measure "closes the door" forever on a hypothesis by assigning it probability zero before it is refuted or that assigns unit probability to a hypothesis before it is logically entailed by the data and background knowledge. It would be more interesting, however, to find a probability measure

P that always "keeps the door open" such that conditionalization *still* fails to converge to the truth. Accordingly, define:

> *P is an open door probability measure for $\mathcal{K}$, C, H*
> *$\Leftrightarrow$ for each $h \in H$, $e \in \omega^*$ such that $[e] \cap \mathcal{K} \neq \varnothing$,*
>
> (i) $P(e) > 0$ *and*
>
> (ii) $P(C_h, e) = 0 \Leftrightarrow [e] \cap \mathcal{K} \subseteq \overline{C_h}$ *and*
>
> (iii) $P(C_h, e) = 1 \Leftrightarrow [e] \cap \mathcal{K} \subseteq C_h$.

But even when the measure P is assumed to be an open door measure, it is possible for a conditionalizer to remain stuck for eternity on a value less than one when the hypothesis in question is true, even when the hypothesis is trivially refutable with certainty by a finite state automaton that does not conditionalize. Indeed, the following example has the property that there are uncountably many possible data streams on which conditionalization never updates its initial probability on the hypothesis for eternity.

Proposition 13.5

There exists a $C_h \subseteq 2^\omega$ and a countably additive probability measure P on 2^ω such that

(a) *P is an open door measure for 2^ω, C, $\{h\}$,*

(b) *h is refutable$_C$ with certainty (by a five-state automaton) given 2^ω, and*

(c) *for uncountably many $\varepsilon \in C_h$, for each n, $P(C_h, \varepsilon|n) = P(C_h) = 0.5$.*

Proof: Let h be the hypothesis that there will never appear two consecutive zeros with the first zero appearing at an even position. In other words, $C_h = \{\varepsilon: \forall n, \varepsilon_{2n} = 0 \Rightarrow \varepsilon_{2n+1} \neq 0\}$. The shaded fans in *Figure 13.6* are all included in the *complement* of C_h.

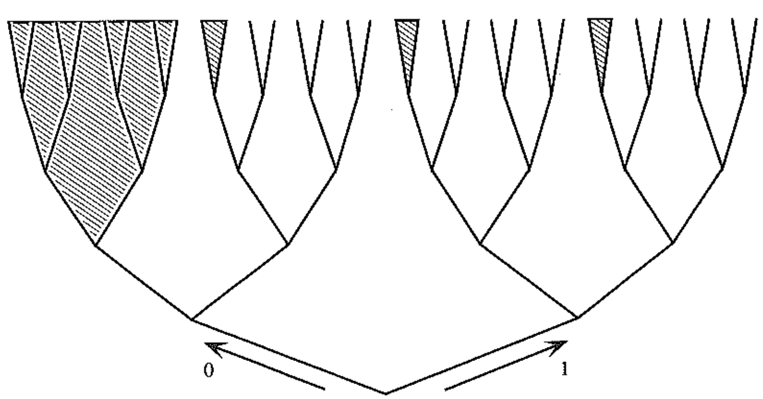

Figure 13.6

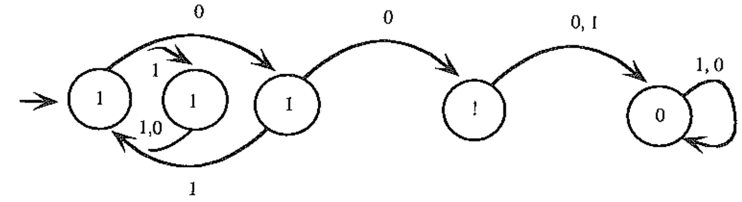

Figure 13.7

h is refutable$_C$ with certainty given 2^ω by the automaton depicted in *Figure 13.7*, which establishes (b).

Let $x \in \{0, 1\}$, $e \in 2^*$. Define $[x, i] = \{\varepsilon: \varepsilon_i = x\}$. If P is a probability measure, let us abbreviate $P([x, lh(e)], [e])$ as the more readable expression $P(x, e)$ and abbreviate $P([e])$ as $P(e)$. Now define measures P_+, P_- as follows:

$$P_+(x, e) = \begin{cases} 0 & \textit{if } [e^*x] \subseteq \overline{C_h} \\ 1 & \textit{if } [e^*|1-x|] \subseteq \overline{C_h} \\ 0.5 & \textit{otherwise.} \end{cases}$$

$$P_-(x, e) = 0.5.$$

The function $|1 - x|$ simply exchanges 0s and 1s. Intuitively, P_- thinks the next datum is generated by a fair coin toss, while P_+ thinks the next datum is generated by a fair coin toss if both outcomes are consistent with h; else the outcome logically excluded by h occurs with probability 0 and the other outcome occurs with probability 1. $P_+(e)$, $P_-(e)$ are extended, by induction, to each $e \in 2^*$, as follows:

$$P_+(\mathbf{0}) = P_-(\mathbf{0}) = 1.$$

$$P_+(e^*x) = P_+(e)P_+(x, e).$$

$$P_-(e^*x) = P_-(e)P_-(x, e).$$

Figure 13.8 illustrates the values of P_+ on data sequences of length ≤ 5.

Observe that,

$$\sum_{lh(e)=k} P_+(e) = 1,$$

and similarly for P_-. It then follows by standard extension techniques[4] that P_+ and P_- have unique, countably additive extensions to all Borel events on 2^ω, so that for each Borel set $S \subseteq 2^\omega$, $P_+(S)$ and $P_-(S)$ are defined. The following two lemmas solve for the priors assigned to C_h by P_+ and P_-, respectively.

[4] Halmos (1974): 212, theorem A.

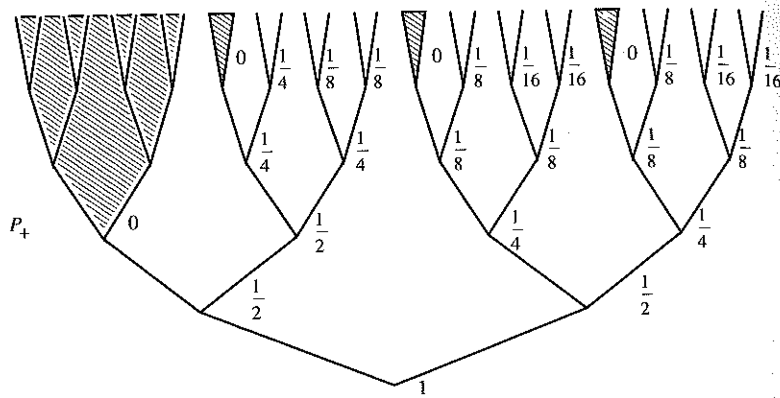

Figure 13.8

Lemma 13.6

$P_+(C_h) = 1.$

Proof: $\overline{C_h}$ is a disjoint, countable union of fans. Since each of these fans is assigned zero probability by P_+, the sum of these probabilities is zero so $P_+(C_h) = 0$ by countable additivity. Hence, $P_+(C_h) = 0$. ■

Lemma 13.7

$P_-(\overline{C_h}) = 1.$

Proof: For each $i \geq 1$, let G_i = the set of all Boolean sequences e of length $2i$ such that $[e] \subseteq \overline{C_h}$ and for no $e' \subset e$ is $[e'] \subseteq \overline{C_h}$. Let $\mathcal{G}_i = \bigcup \{[e]: e \in G_i\}$. A glance at *Fig. 13.8* reveals that $|G_i| = 3^{(i-1)}$. Also, for each $e \in G_i$, $P_-(e) = 1/2^{2i} = 1/4^i$. Since for all distinct $e, e' \in G_i$, $[e] \cap [e'] = \varnothing$, it follows by finite additivity that:

$$P_-(\mathcal{G}_i) = \sum_{e \in G_i} P_-(e) = |G_i| \frac{1}{4^i} = \frac{3^{(i-1)}}{4^i}.$$

By countable additivity and the fact that $\mathcal{G}_i \cap \mathcal{G}_{i+1} = \varnothing$, we have:

$$P_-(\overline{C_h}) = \sum_{i=1}^{\infty} P_-(\mathcal{G}_i) = \sum_{i=1}^{\infty} \frac{3^{(i-1)}}{4^i} = \frac{1}{4} \sum_{i=0}^{\infty} \frac{3^i}{4^i} = \frac{1}{4} 4 = 1.$$ ■

For each Borel set $\mathcal{S} \subseteq 2^\omega$, define P_0 to be the 50-50 *mixture* of P_+ and P_-:

$$P_0(\mathcal{S}) = 0.5P_+(\mathcal{S}) + 0.5P_-(\mathcal{S}).$$

The following lemma shows that P_0 satisfies (b).

Lemma 13.8

P_0 is an open door measure for $\mathcal{K}$, C, H.

Proof: (i) $P_0(e) \geq 0.5P_-(e) = 1/2^{lh(e)+1} > 0$. (ii) Suppose $[e] - \overline{C_h} \neq \varnothing$. Then $P_+(e) > 0$. But then:

$$P_0(C_h, e) = \frac{P_0(e \,\&\, C_h)}{P_0(e)} = \frac{0.5P_+(e \,\&\, C_h) + 0.5P_-(e \,\&\, C_h)}{P_0(e)} \geq \frac{0.5P_+(e \,\&\, C_h)}{P_0(e)}$$
$$= \frac{0.5P_+(e)}{P_0(e)} > 0.$$

The last identity is by lemma 13.6. A similar argument using proposition 13.7 establishes condition (iii). ■

Let β_P be a conditionalizer who starts out with measure P. Let us now consider how β_P fails. The following lemma is very useful in studying the asymptotic performance of conditionalization:

Lemma 13.9

Let P be an arbitrary probability measure and let $\mathcal{S} \subseteq 2^\omega$ be a Borel set. Then $Lim_{n\to\infty} P(\mathcal{S}, \varepsilon|n) = 1 \Leftrightarrow Lim_{n\to\infty} \dfrac{P(\varepsilon|n, \mathcal{S})}{P(\varepsilon|n, \overline{\mathcal{S}})} = \infty.$

Proof: Since $P(\mathcal{S}, \varepsilon|n) = 1 - P(\overline{\mathcal{S}}, \varepsilon|n)$,

$$Lim_{n\to\infty} P(\mathcal{S}, \varepsilon|n) = 1 \Leftrightarrow Lim_{n\to\infty} P(\overline{\mathcal{S}}, \varepsilon|n) = 0$$
$$\Leftrightarrow Lim_{n\to\infty} \frac{P(\mathcal{S}, \varepsilon|n)}{P(\overline{\mathcal{S}}, \varepsilon|n)} = \infty$$
$$\Leftrightarrow Lim_{n\to\infty} \frac{P(\mathcal{S})P(\varepsilon|n\,\mathcal{S})}{P(\overline{\mathcal{S}})P(\varepsilon|n, \overline{\mathcal{S}})} = \infty \text{ (by Bayes' theorem)}$$
$$\Leftrightarrow Lim_{n\to\infty} \frac{P(\varepsilon|n, \mathcal{S})}{P(\varepsilon|n, \overline{\mathcal{S}})} = \infty \text{ (since } P(\mathcal{S}), P(\overline{\mathcal{S}}) \text{ are positive constants).}$$ ■

The ratio on the right-hand side of the lemma is called the *likelihood ratio* for $\mathcal{S}$. It is also clear that if the likelihood ratio stabilizes to a fixed value, then $P_0(C_h, \varepsilon|n)$ stabilizes to a fixed value less than 1 in the limit. The following lemma solves for the likelihood ratio and priors of P_0.

Lemma 13.10

For each $e \in 2^*$,

(a) $\dfrac{P_0(e, C_h)}{P_0(e, \overline{C_h})} = \dfrac{P_+(e)}{P_-(e)}$.

(b) $P_0(C_h) = P_0(\overline{C_h}) = 0.5$

Proof of (a):

$$\frac{P_0(e, C_h)}{P_0(e, \overline{C_h})} = \frac{P_0(e \,\&\, C_h)/P_0(e)}{P_0(e \,\&\, \overline{C_h})/P_0(e)} = \frac{P_0(e \,\&\, C_h)}{P_0(e \,\&\, \overline{C_h})}$$

$$= \frac{0.5P_+(e \,\&\, C_h) + 0.5P_-(e \,\&\, C_h)}{0.5P_+(e \,\&\, \overline{C_h}) + 0.5P_-(e \,\&\, \overline{C_h})}$$

$$= \frac{0.5P_+(e \,\&\, C_h)}{0.5P_-(e \,\&\, \overline{C_h})} \text{ (by lemmas 13.6 and 13.7)}$$

$$= \frac{P_+(e)}{P_-(e)} \text{ (by lemmas 13.6 and 13.7).}$$

(b) Follows immediately from lemmas 13.6 and 13.7. □

Let us now consider where β_P fails to gradually decide$_C$ h in the limit. If $[e] \subseteq \overline{C_h}$, then $P_0(C_h, e) = 0$, so β_{P_0} is correctly certain of the falsehood of h and remains so, since $P_0(e) > 0$ whenever e is consistent with $\mathcal{K}$. So β_{P_0} can only fail by refusing to approach 1 when h is correct. To see how P_0 can fail to approach 1 when h is correct, consider the everywhere 1 data stream ζ. From lemma 13.10(a) and the definitions of P_+ and P_- we have for each n,

$$\frac{P_0(\zeta|n, C_h)}{P_0(\zeta|n, \overline{C_h})} = \frac{P_+(\zeta|n)}{P_-(\zeta|n)} = \frac{1/2^n}{1/2^n} = 1.$$

So by lemma 13.10(b), we have that for each n, $P_0(C_h, \zeta|n) = P_0(C_h) = 0.5$. So on the data stream ζ, β_{P_0} never updates its initial conjecture of 0.5 for eternity. Notice that the same phenomenon will arise on any data stream for which the likelihood ratio remains unity forever. This will occur on any data stream in the following set (*Fig. 13.9*):

$$\mathcal{Very\ Bad} = \{\varepsilon\colon \forall n, \{\varepsilon|n{*}0] \textit{ is not a subset of } \overline{C_h}\}.$$

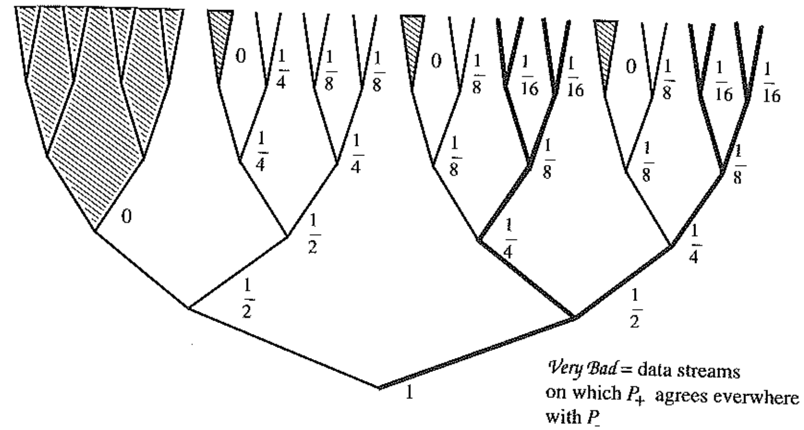

Figure 13.9

Moreover, there are data streams on which the likelihood ratio changes only finitely often before stabilizing to some value less than 1:

$$\mathcal{Bad} = \{\varepsilon\colon \exists n\, \forall m \geq n, [\varepsilon|m{*}0] \textit{ is not a subset of } \overline{C_h}\}.$$

If $\varepsilon \in \mathcal{Bad}$, then $P_0(C_h, \varepsilon|m)$ goes up for a finite time and then remains fixed at a value less than 1. Furthermore, $\mathcal{Very\ Bad}$ and hence $\mathcal{Bad}$ are both uncountably infinite, which may be seen as follows: let $\varepsilon \in 2^\omega$. Encode 0 as 10 and 1 as 11. Then the sequence τ such that for each $i \geq 1$, $(\tau_{2i}, \tau_{2i+1}) = code(\varepsilon_i)$ is an element of $\mathcal{Very\ Bad}$ and τ uniquely codes ε. By a similar argument, there are uncountably many more data streams in $\mathcal{Bad} - \mathcal{Very\ Bad}$ on which β_{P_0} fails. This establishes (c). ■

So no conditionalizer can do what is impossible, and some conditionalizers (even those who always "leave the door open" concerning alternative hypotheses) can fail in uncountably many possible worlds to arrive at the truth even when a trivial method is guaranteed to do so. It remains to consider whether every solvable inductive problem is solved by some conditionalizer or other. It is unreasonable to demand a strict conditionalizer to conjecture 0 or 1 unless he is certain, since a strict conditionalizer can never take back a 0 or a 1. The natural criterion of success for a conditionalizer is therefore gradual decidability. The measure defined in the preceding example does not succeed even in this sense. But the same problem can be solved in this sense by a strict conditionalizer whose initial probability measure is carefully tailored to the particular hypothesis at hand.

Proposition 13.11 (T. Seidenfeld)

Let C_h *be as in example 2.5. Then there exists a countably additive probability measure* P_1 *on* 2^ω *such that* β_{P_1} *gradually decides*$_C$ *h given* 2^ω.

Proof: Let $e \in 2^*$, $lh(e) = n$. Then define:

$$P_-(e) = 1/2^n.$$

$$P_+(e) = \begin{cases} 0 & \text{if } [e] \subseteq \overline{C_h} \\ \dfrac{1}{3^{(n+1)/2}} & \text{if } n \text{ is odd and } e \text{ ends with } 0 \\ \dfrac{2}{3^{(n+1)/2}} & \text{if } n \text{ is odd and } e \text{ ends with } 1 \\ \dfrac{1}{3^{n/2}} & \text{otherwise.} \end{cases}$$

P_- is just as in proposition 13.5, so once again we have that P_- induces a countably additive probability measure on 2^ω with $P_-(C_h) = 0$. As *Figure 13.10* should make clear, the values assigned by P_+ also add up to 1 over data segments of a fixed length, so just as in the proof of proposition 13.5, P_+ induces a unique, countably additive probability measure on 2^ω.

Since P_+ assigns each fan in $\overline{C_h}$ probability zero, we have by the proof of lemma 13.9 that $P_+(\overline{C_h}) = 0$, so $P_+(C_h) = 1$. For each Borel set $\mathcal{S} \subseteq 2^\omega$, define:

$$P_1(\mathcal{S}) = 0.5P_+(\mathcal{S}) + 0.5P_-(\mathcal{S}).$$

Then by lemmas 13.7 and 13.10 and an obvious analog of lemma 13.6, we have that:

$$(*)\ \frac{P_1(e, C_h)}{P_1(e, \overline{C_h})} = \frac{P_+(e)}{P_-(e)}.$$

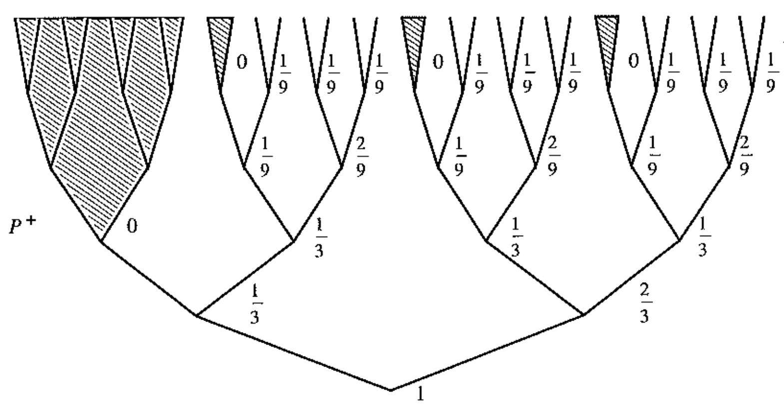

Figure 13.10

It remains to see that β_{P_1} gradually decides$_C$ h given 2^ω. Let $\varepsilon \in 2^\omega$. Suppose $\varepsilon \in \overline{C_h}$. Then there is an n such that $[\varepsilon|n] \subseteq \overline{C_h}$, since $\overline{C_h}$ is a union of fans. Then it is immediate that for each $m \geq n$, $P_1(C_h, \varepsilon|n) = 0$, so β_{P_1} stabilizes to 0 on ε, which is correct. Now suppose $\varepsilon \in C_h$. By lemma 13.9, it suffices to show that

$$Lim_{n\to\infty} \frac{P_1(\varepsilon|n, C_h)}{P_1(\varepsilon|n, \overline{C_h})} = \infty.$$

For this it suffices, in turn, to show that both the even and the odd subsequences of likelihood ratios go in infinity:

$$Lim_{n\to\infty} \frac{P_1(\varepsilon|(2n), C_h)}{P_1(\varepsilon|(2n), \overline{C_h})} = \infty \quad \textit{and} \quad Lim_{n\to\infty} \frac{P_1(\varepsilon|(2n+1), C_h)}{P_1(\varepsilon|(2n+1), \overline{C_h})} = \infty.$$

On the odd subsequence we have:

$$\frac{P_1(\varepsilon|(2n+1), C_h)}{P_1(\varepsilon|(2n+1), \overline{C_h})} = \frac{1/3^{((2n+1)+1)/2}}{1/2^{2n+1}} = \frac{2^{2n+1}}{3^{n+1}},$$

while on the even subsequence we have:

$$\frac{P_1(\varepsilon|(2n), C_h)}{P_1(\varepsilon|(2n), \overline{C_h})} = \frac{1/3^{(2n)/2}}{1/2^{2n}} = \frac{2^{2n}}{3^n}.$$

To see that both subsequences go to infinity, observe that $2^8/3^4 > 3$. Let $k \in \omega$ be given. Either ratio exceeds k by stage $4\log_3(k)$ and remains above k thereafter. So by lemma 13.10, $P_1(C_h, \varepsilon|n)$ correctly approaches 1 as n goes to infinity. ■

The examples show that not all initial, joint probability measures are equal when it comes to finding the truth about a deterministic hypothesis. Even an open door measure can result in failures on uncountable sets of data streams, while choosing another measure may guarantee success no matter what. Since a familiar criticism of personalism is that one's initial probabilities are arbitrary,[5] this is a consideration that might be taken into account in constraining one's choice.[6] The logical perspective on induction therefore has

[5] This advice pertains more to purely subjective personalists than to those who think that the likelihoods are given by an objective statistical model. The latter are not free to tinker at will with the probabilities $P(\mathcal{E}|\mathcal{H})$, and these are the probabilities that drive the convergence arguments in the examples.

[6] There is some evidence of such concerns in the statistical literature: e.g., Foster (1991). Diaconis and Freedman (1986) show that for some prior probability measures, Bayesian updating is not reliable in a classical sense (it does not have a unit likelihood of converging to the truth regardless of which likelihood parameter is the true one). If likelihoods are interpreted strictly as limiting relative frequencies so that parameters pick out determinate subsets of $\mathcal{N}$, then their question is of just the sort under discussion.

relevance for philosophies of induction and confirmation that are based on updating by Bayes' theorem. In fact, an important and general question for Bayesians is whether or not conditionalization actually prevents success by ideal agents when success is possible. In other words:

Question 13.12

If h is gradually decidable$_C$ given $\mathcal{K}$, does there exist a strict conditionalizer that gradually decides$_C$ h given $\mathcal{K}$?

Regardless of the outcome of the question for arbitrary, ideal methods, conditionalization is definitely restrictive for the class of all arithmetically definable methods. Recall from chapter 6 that this class includes highly idealized methods much more complex than computable methods.

Since the conjecture of a conditionalizer is a real number, we must specify what it is for such a method to be arithmetically definable. We will represent the conjecture as an infinite decimal expansion and require that for each given n, the arithmetical definition of β determines the nth position of $\beta(h, e)$. For β to be computable, we will require that there be a program M such that $M[h, e, n]$ returns the nth decimal position of $\beta(h, e)$.

Now we have the following result:

Proposition 13.13

Let $C_h = \{\delta\}$, as in proposition 7.19. Then

(1) *h is computably refutable with certainty but*

(2) *no arithmetically definable conditionalizer can even gradually verify$_C$ or refute$_C$ h.*[7]

Proof: If $C_h \cap [e] \cap \mathcal{K} = \varnothing$, then $P(h, e) = 0$. Hence, a conditionalizing method is consistent. Apply corollary 7.24. ■

So if an arithmetically definable agent elects to update probabilities by Bayes' rule, he may have to pay a price in terms of logical reliability.[8] Even under liberal standards of constructiveness, the respective aims of short-run coherence and long-run reliability compete rather than reinforce one another.

[7] In fact, no *hyper-arithmetically definable* method can succeed, either. Cf. Kelly and Schulte (1995). For a definition of the hyper-arithmetical functions, cf. Rogers (1987).

[8] For a different sort of restrictiveness result for conditionalization, cf. Osherson and Weinstein (1988). Their result concerns discovery rather than assessment, and its model of Bayesian discovery would be considered controversial by some Bayesians.

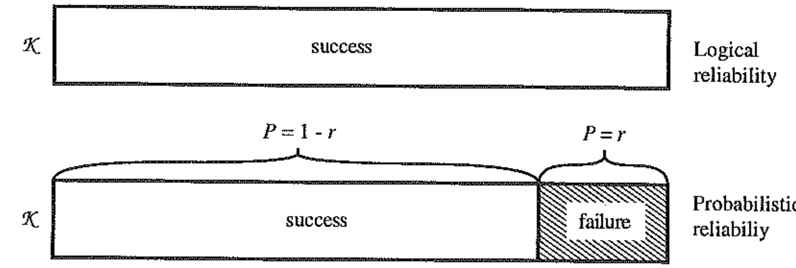

Figure 13.11

3. Probabilistic Reliability

In the preceding section, the reliability of conditionalization was analyzed from the logical point of view developed in the preceding chapters. But advocates of conditionalization as an inductive method rarely adopt the logical perspective just considered. Their sweeping convergence claims are based on *probabilistic* rather than *logical* reliability. Logical reliability demands convergence to the truth on each data stream in $\mathcal{K}$. Probabilistic reliability requires only convergence to the truth over some set of data streams that carries sufficiently high probability (*Fig. 13.11*).

Let P be a probability measure on $\mathcal{N}$.

α decides$_C$ H gradually with probability r in P
$\Leftrightarrow$ there exists $\mathcal{K} \in Bo$ such that

(1) *$P(\mathcal{K}) \geq r$ and*

(2) *α decides$_C$ H gradually given $\mathcal{K}$.*

Similar definitions can be given for the other notions of success, including the discovery case. Now we can state the sort of result that some probabilists have in mind when they say that conditionalization will arrive at the truth. The following proposition says that the conditionalization method β_P that starts out with measure P has a unit probability *according to the measure P* of approaching the truth value of an arbitrary Borel hypothesis. Observe that this result imposes no restrictions on the probability measure P other than countable additivity.

Proposition 13.14[9]

Let P be a countably additive probability measure on $\mathcal{N}$. Let $C \in Bo$. Then $P(\{\varepsilon$: β_P gradually decides$_C$ ω on $\varepsilon\}) = 1$.

Proof: Halmos (1974): section 49, theorem B, shows that β_P can identify in the limit each $h \in \omega$ with probability 1 in P. Let $\mathcal{L}_h$ be the set of probability 0

[9] For a detailed discussion of this similar results, cf. Schervish et al. (1990).

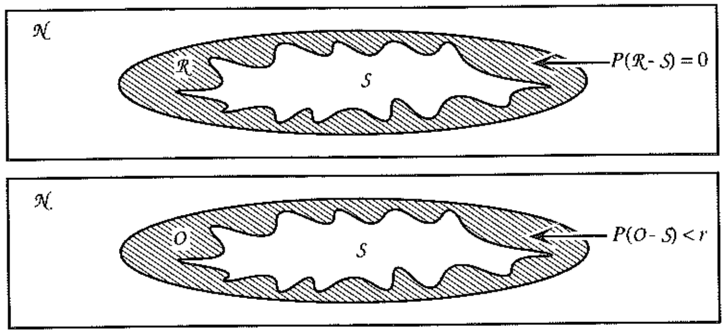

Figure 13.12

"neglected" in the case of hypothesis h. The union of all these sets still has probability 0 by countable additivity. ■

Proposition 13.14 entails that each Borel hypothesis is decidable in the limit with probability 1. It can also be shown that if we are willing to countenance an arbitrarily small but positive probability of error, then each Borel hypothesis is decidable with certainty. These facts are summarized as follows.

Proposition 13.15

For each countably additive P, for each $C \in Bo$

(a) *for each $r > 0$, ω is decidable$_C$ with certainty with probability $1 - r$ in P;*

(b) *ω is decidable$_C$ in the limit with probability 1 in P.*

Proof: (b) follows from proposition 13.14. (a) follows from proposition 13.16 below. ■

In light of the strong, negative results of preceding chapters, this is amazing. Problems unsolvable even with strong background knowledge in the logical sense are *all* solvable in the limit with unit probability and are decidable with certainty with arbitrarily high probability.[10] Arbitrarily low or zero probabilities of error hardly seem worth haggling over. Indeed, some probabilists refer to unit probability events as being *almost sure* or *practically certain*.

The usual proof of proposition 13.15 involves the following result as a lemma. It states that every Borel set is "almost surely" Π^B_2 and is a Σ^B_1 set with arbitrarily high probability. Moreover, the approximating sets can always be chosen as *supersets* (*Fig. 13.12*).

[10] In light of chapter 9, these results carry over to discovery problems as well.

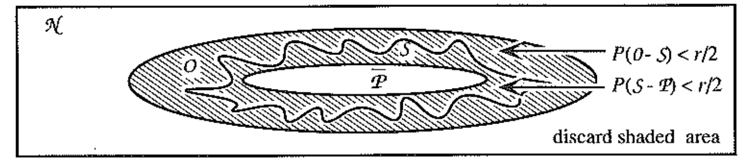

Figure 13.13

Proposition 13.16 The approximation theorem

For each countably additive P, for each $\mathcal{S} \in Bo$

(a) *for each $r > 0$ there is an $O \in \Sigma^B_1$ such that $\mathcal{S} \subseteq O$ and $P(O) - P(\mathcal{S}) < r$.*

(b) *there is an $\mathcal{R} \in \Pi^B_2$ such that $\mathcal{S} \subseteq \mathcal{R}$ and $P(\mathcal{R}) = P(\mathcal{S})$.*

Proof: Royden (1988): 294, proposition 7. ■

This fact implies the following proposition, which places the approximation theorems into the context of the characterization theorems of the preceding chapters. The proof illustrates how a probabilist can neglect the ragged "periphery"[11] of a complex hypothesis while leaving most of its probability intact.

Proposition 13.17

For each countably additive P, for each $\mathcal{S} \in Bo$,

(a) *for each $r > 0$, for each $h \in \omega$, there is $\mathcal{K} \in Bo$ such that $P(\mathcal{K}) > 1 - r$ and $C_h \in \Delta[\mathcal{K}]^B_1$;*

(b) *for each $h \in \omega$, there is $\mathcal{K} \in Bo$ such that $P(\mathcal{K}) = 1$ and $C_h \in \Delta[\mathcal{K}]^B_2$.*

Proof: (a) Let $\mathcal{S} \in Bo$ be given. We want to find $\mathcal{K}$ such that $P(\mathcal{K}) > 1 - r$ and $\mathcal{S} \in \Delta[\mathcal{K}]^B_1$. By proposition 13.16(a), we can choose open sets O, $\mathcal{P}$ such that $\mathcal{S} \subseteq O$, $\bar{\mathcal{S}} \subseteq \mathcal{P}$, $P(O) - P(\mathcal{S}) < r/2$ and $P(\mathcal{P}) - P(\bar{\mathcal{S}}) < r/2$ (*Fig. 13.13*). Thus, $P(\mathcal{S}) - P(\bar{\mathcal{P}}) < r/2$, since $\mathcal{P} - \bar{\mathcal{S}} = \mathcal{S} - \bar{\mathcal{P}}$.

So we have a closed $\bar{\mathcal{P}} \subseteq \mathcal{S}$ that differs in probability from $\mathcal{S}$ by less than $r/2$ and an open $O \subseteq \mathcal{S}$ that differs in probability from $\mathcal{S}$ by less than $r/2$. By finite additivity, $P(O - \bar{\mathcal{P}}) < r$. Now we may "neglect" this difference by choosing $\mathcal{K} = O - \bar{\mathcal{P}}$. Evidently, $P(\mathcal{K}) > 1 - r$, and since $\mathcal{K} \cap \mathcal{S} = \mathcal{K} \cap \bar{\mathcal{P}} = \mathcal{K} \cap O$, we have $\mathcal{S} \in \Delta[\mathcal{K}]^B_1$. The argument for (b) is similar, except that we end up with a Σ^B_2 approximation from within and a Π^B_2 approximation from without whose difference has probability zero. ■

[11] I did not say *boundary*, because in the measure 1 case, not all of the boundary is ignored. Otherwise, all hypotheses would be decidable with certainty with unit probability.

Probabilistic reliability is easier to achieve precisely because the probabilist entitles himself to augment his background knowledge by removing a set $\mathcal{Z}$ of probability zero from $\mathcal{K}$, whereas logical reliability demands success over all of $\mathcal{K}$. That the probabilist's "negligible" sets may be uncountably infinite even when the data, the hypothesis, and its negation all start out with nonzero probability is clear from propositions 13.5 and 13.7. It is a fine thing for the realist that alleged canons of inductive rationality entitle him to claim sufficient knowledge to solve all inductive problems in the limit. A competent skeptic, on the other hand, will reject this blanket entitlement to knowledge as a technically disguised form of dogmatism (*Fig. 13.14*).

We have seen in chapter 4 that local underdetermination and demonic arguments depend on missing limit points in the boundary of the hypothesis. The probabilistic convergence and approximation theorems just described imply that the boundary of a hypothesis (in the topological sense) is always a set of negligible probability. In other words, *arbitrarily high* topological complexity can be discounted in a set of *arbitrarily small* probability, and complexity higher than the Δ_2^B level can be neglected in a set of zero probability.

This discussion may shed some light on the stubborn persistence of debates between realists and skeptics. The antirealist looks at a problem and points out that it has high topological complexity and hence a high degree of underdetermination, subjecting it to demonic arguments against logical reliability, as we saw with Sextus Empiricus. The realist focuses on size rather than complexity. He observes that the boundary of the hypothesis, in which all of the antirealist's complexity resides, is a set of negligible probability. Neither position contradicts the other. The difference is one of emphasis on different, objective mathematical features of the problem at hand.

If, indeed, the scientific realism debate hinges to some extend on one's attitude toward neglecting the logically "rough" peripheries of hypotheses, it is of the utmost epistemological import to understand how it is that the axioms of probability guarantee that this periphery is always neglectable. In the next section, I will focus on the crucial role of countable additivity in this regard.

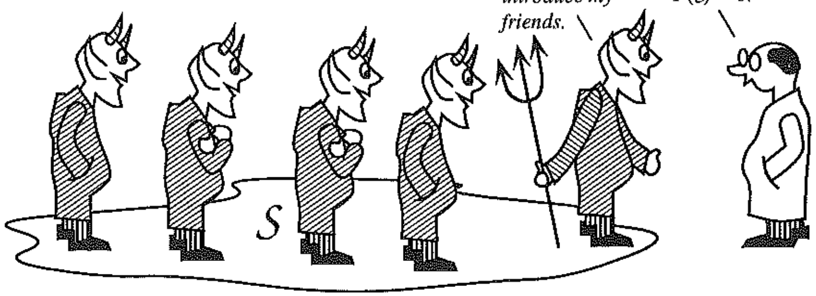

Figure 13.14

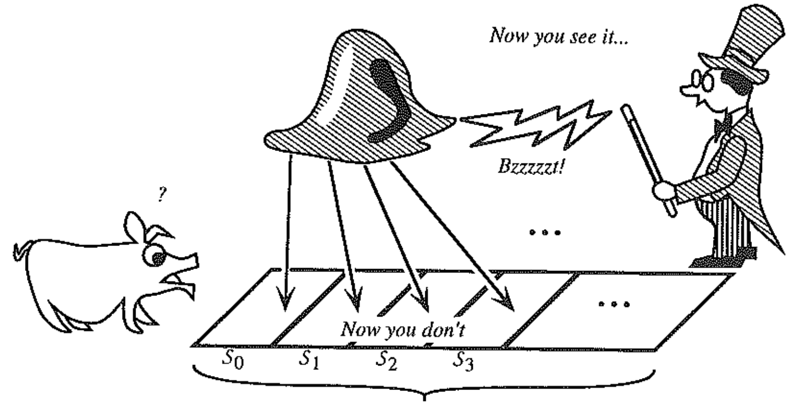

Figure 13.15

4. Countable Additivity

All the results reported in the preceding section assume countable additivity. Countable additivity was introduced by Kolmogorov[12] for expedience, and most probabilists have followed suit.

> The general condition of countable additivity is a further restriction . . . —a restriction without which modern probability theory could not function. It is a tenable point of view that our intuition demands infinite additivity just as much as finite additivity. At any rate, however, infinite additivity does not contradict our intuitive ideas, and the theory built on it is sufficiently far developed to assert that the assumption is justified by its success.[13]

Without countable additivity, we would be left with the apparently magical possibility that the probability of the whole could end up being more than the sum of its parts when it consists of a countable infinity of nonoverlapping parts. In particular, a countable infinity of probability zero parts could add up to a set with probability one (*Fig. 13.15*).[14] That seems strange.

On the other hand, countable additivity has the awkward consequence that probabilities on a countable partition of $\mathcal{N}$ must be "biased" in order to add

[12] Kolmogorov (1956).
[13] Halmos (1974): 187.
[14] This is an instance of what DeFinetti (1972: 143) calls *nonconglomerability*. Let $\Pi = \{\mathcal{H}_1, \ldots, \mathcal{H}_n, \ldots\}$ partition $\mathcal{N}$. Then P is *conglomerable* in Π just in case for all measurable events C, $P(C) = \Sigma_{i \in \omega} P(\mathcal{E}, \mathcal{H}_i)$. In the example, we have $P(\mathcal{S}) = 1$ but for each i, $P(\mathcal{S}, \mathcal{S}_i)P(\mathcal{S}_i) = P(\mathcal{S}_i) = 0$ so $\Sigma_{i \in \omega} P(\mathcal{S}, \mathcal{S}_i)P(\mathcal{S}_i) = 0$. Schervish et al. (1984) show that subject to some conditions concerning the existence of conditional probabilities, every merely finitely additive P is nonconglomerable in *some* partition (depending on P). The "strangeness" of nonconglomerability is sometimes cited as an argument for countable additivity.

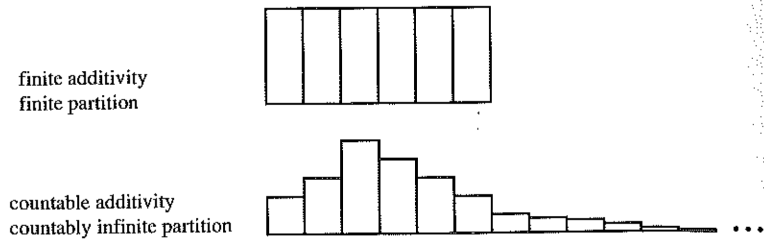

Figure 13.16

up to 1, since if the same probability is assigned to each cell, then the sum must be either 0 or infinity (*Fig. 13.16*). DeFinetti puts the issue this way:

> Suppose we are given a countable partition into events E_i, and let us put ourselves into the subjectivistic position. An individual wishes to evaluate the $p_i = P(E_i)$: he is free to choose them as he pleases, except that, if he wants to be coherent, he must be careful not to inadvertently violate the conditions of coherence. Someone tells him that in order to be coherent he can choose the p_i in any way he likes, so long as the sum $= 1$ (it is the same thing as in the finite case, anyway!).
>
> The same thing?!!! You must be joking, the other will answer. In the finite case, this condition allowed me to choose the probabilities to be all equal, or slightly different, or very different; in short, I could express any opinion whatsoever. Here, on the other hand, the content of my judgments enter into the picture: I am allowed to express them only if they are unbalanced. . . . Otherwise, even if I think they are equally probable . . . I am obliged to pick "at random" a convergent series which, however I choose it, is in absolute contrast to what I think. If not, you call me incoherent! In leaving the finite domain, is it I who has ceased to understand anything, or is it you who has gone mad?[15]

In keeping with this spirit, DeFinetti and L. J. Savage[16] have proposed foundational interpretations of probability in which countable additivity is not guaranteed. For frequentists, the issue is more readily settled, since limiting relative frequencies do not always satisfy countable additivity.[17]

From the logical reliabilist perspective, there is another reason to question countable additivity. It will be seen in this section how countable additivity is invoked in probabilistic convergence and approximation theorems like

[15] De Finetti (1990): 123.
[16] Savage (1972).
[17] For example, we can have each natural number occur in a collective with limiting relative frequency 0 by having it occur finitely often, but the limiting relative frequency that some natural number will occur is unity, since only natural numbers occur. De Finetti chided "frequentist" practice for failure to respect this failure of countable additivity on their interpretation: "So far as I know, however, none of [the frequentists] has ever taken this observation into account, let alone disputed it; clearly it has been overlooked, although it seems to me I have repeated it on many occasions" (De Finetti 1990: 123).

propositions 13.14, 13.15, and 13.16. We will see that the objectionable bias imposed by countable additivity is involved directly in these arguments. In section 13.5, it will be shown, moreover, that countable additivity is necessary for obtaining the probabilistic convergence and approximation theorems at issue. If probabilistic convergence theorems are to serve as a philosophical antidote to the logical reliabilist's concerns about local underdetermination and inductive demons, then countable additivity is elevated from the status of a mere technical convenience to that of a central epistemological axiom favoring scientific realism. Such an axiom should be subject to the highest degree of philosophical scrutiny. Mere technical convenience cannot justify it. Neither can appeal to its "fruits," insofar as they include precisely the convergence theorems at issue.

Let us consider in a concrete example how the bias imposed by countable additivity can make universal hypotheses decidable with arbitrarily high probability, as is claimed in proposition 13.15(a). Sextus argued that no matter how many 1s have been observed, the next observation might be a 0, so "every observation will be 1" is not verifiable with certainty. The complement $\{\bar{\zeta}\}$ of this hypothesis is the union of all fans of form $[(1, 1, 1, \ldots, 0)]$ (i.e., with handles consisting of a finite sequence of 1s followed by a single (0). Let $\mathcal{F}_i$ be the fan of this form with exactly i 1s occurring in its handle, as depicted in *Figure 13.17.*

Suppose that the prior probability of the hypothesis in question is r (i.e., that $P(\{\zeta\}) = r$). Then $P(\{\bar{\zeta}\}) = 1 - r$, by finite additivity. Since the fans are disjoint, the sum of the probabilities of the individual fans is exactly $P(\{\bar{\zeta}\}) = 1 - r$, by countable additivity. Since this infinite sum is the limit of the finite initial sums of the series, it follows that for each $u > 0$, there is a position n in the series such that the probability left in the remaining fans is less than u (*Fig. 13.18*). In other words, the bias imposed by countable additivity is in this case a bias for seeing refutations of $\{\zeta\}$ sooner rather than later.

For measure u decidability, the scientist is entitled to assume that none of the fans after position n will ever be entered. This amounts to cutting the infinite tail off of Sextus's argument, so that the demon runs out of opportunities to fool the scientist. Let $C_h = \{\zeta\}$. Let $\mathcal{K} = \{\zeta\} \cup$ the union of all fans veering off

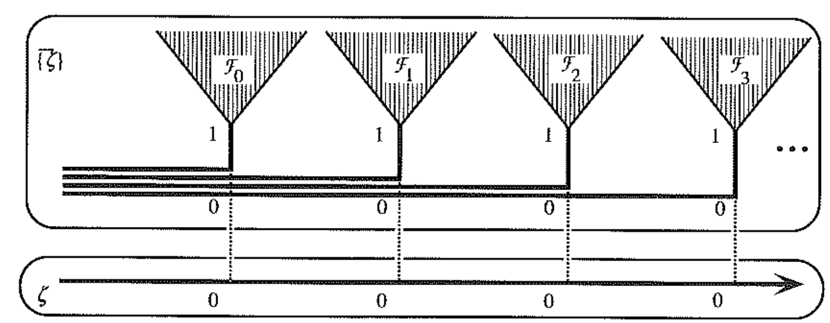

Figure 13.17

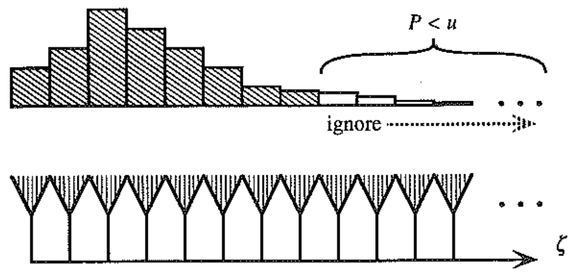

Figure 13.18

from ζ at a stage $\leq n$. Clearly, there exists an effective method α that is guaranteed to decide$_C$ h by time n given $\mathcal{K}$;

$$\alpha(h, e) = \begin{cases} 1 & \textit{if } e \subseteq \zeta \\ 0 & \textit{otherwise.} \end{cases}$$

Hence, if P is countably additive and $r > 0$, then h is decidable$_C$ by some fixed time n with probability $1 - r$ in P. Since the ancient skeptical argument is driven by concerns of error in the indefinite future, it is not surprising that an axiom of rationality imposing the attitude that a universal hypothesis will be refuted sooner rather than later should have antiskeptical consequences.[18] It might be claimed that such an attitude is justified, because finite beings have no interest in what happens in the long run. But be that as it may, such indifference ought to be reflected in the utilities we associate with outcomes in the indefinite future, rather than in general axioms governing rational beliefs about what will happen then.

In the example at hand, we have not only decidability with certainty (with high probability), but an a priori bound (calculated from P) on how long it will take to arrive at the right answer with a given probability, which is something far stronger than what is claimed in proposition 13.14. Once we discount "trifling" probabilities of error concerning a universal hypothesis, the only interesting question remaining is how long we have to wait until we are sufficiently sure we are right[19] and whether a computer could efficiently recognize counterexamples.[20]

[18] Assuming that the data are generated by independent and identically distributed (IID) trials of an experiment yields an even stronger result, namely, that the probability of refutation always decreases with time, with well-known bounds on the rate of decrease. But while no one accepts IID as a general postulate of rationality, countable additivity is often treated as such. For that reason, I focus on the convergence theorems that assume only countable additivity.

[19] This may not be true of hypotheses of higher complexity, as we shall see shortly.

[20] These questions have been drawing increased attention from computer scientists under the rubric of PAC (probably approximately correct) learning theory. At present, the results of this theory depend not just on countable additivity, but on strong independence assumptions as well.

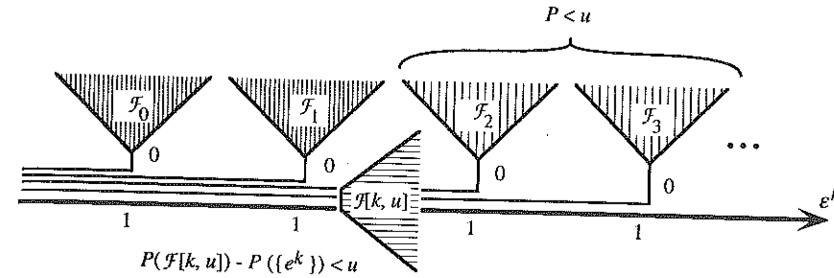

Figure 13.19

High probabilities of inductive success by a given time may evaporate depending on how the evidence comes in. For example, if the hypothesis states that heads will be observed, then the probability that the hypothesis will be verified by the tenth toss already exceeds 0.999, but the probability that it will be verified by the tenth toss given failures on the first nine tosses is just 0.5. So unlike logical reliability, which never disappears with time, a high probability of success can be an evanescent virtue.

A unit probability of success with certainty is permanent, assuming that data of probability 0 are not encountered. But universal hypotheses cannot always be verified with certainty with unit probability. For suppose that $P(\{\zeta\}) = r$ where $0 < r < 1$ and the remaining $1 - r$ probability is distributed in a countably additive way over the diverging fans, so that $P(\mathcal{F}_i) > 0$, for infinitely many i. To succeed with probability 1, the method must succeed on ζ, since $P(\{\zeta\}) = r > 0$, and must also succeed on infinitely many fans veering off from ζ, since infinitely many such fans carry positive probability. Sextus's demonic argument shows that no possible method has this property.

On the other hand, if $C_h \in Bo$, then decidability in the limit is guaranteed with unit probability when P is countably additive, by proposition 13.15. Now we consider the role of countable additivity in the argument for this claim. Let us return to the simple and familiar case of finite divisibility, introduced in chapter 3. In this example, $C_{h_{fin}}$ is the set of all infinite, Boolean-valued data streams that stabilize to 0. We know that $C_{h_{fin}} \in \Sigma_2^B - \Pi_2^B$ and that $C_{h_{fin}}$ is countable. So if h_{fin} is to be decidable$_C$ in the limit with probability 1, as proposition 13.17 says, then there must be some set $\mathcal{K}$ of probability one such that $C_{h_{fin}} \in \Pi[\mathcal{K}]_2^B$. In other words, $\bar{\mathcal{K}}$ is to be neglected.

What does $\bar{\mathcal{K}}$ look like?[21] Enumerate $C_{h_{fin}}$ as $\varepsilon^1, \varepsilon^2, \ldots, \varepsilon^n \ldots$. If we repeat the above discussion of Sextus' argument, treating $\{\varepsilon^k\}$ as a hypothesis, then we have that for each specified u, we can find a fan $\mathcal{F}[k, u]$ containing ε^k such that $P(\mathcal{F}[k, u]) - P(\{\varepsilon^k\}) < u$ (*Fig. 13.19*). This is again because countable additivity forces us to drop probability on the complement of $\{\varepsilon^k\}$ in "lumps" on the countably many disjoint fans veering off from ε^k, so that it is eventually

[21] The construction that follows is a minor adaptation of the standard proof that the [0, 1] interval can be partitioned into a meager set and a set of Lebesgue measure zero (Royden 1988: 161, no. 33b.)

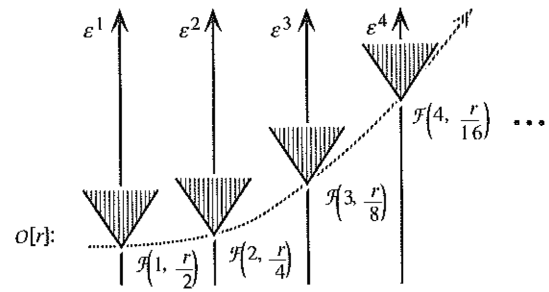

Figure 13.20

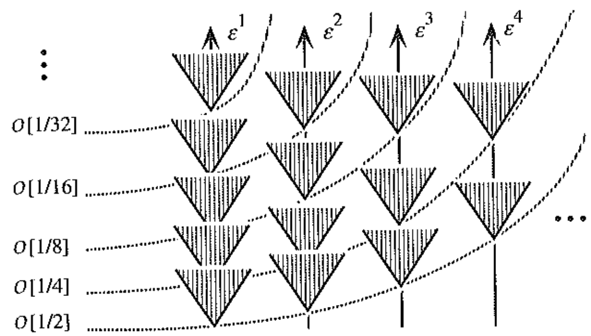

Figure 13.21

almost entirely used up (*Fig. 13.19*). For each $r \in [0, 1]$, we have:

$$\sum_{n=1}^{\infty} \frac{r}{2^n} = r.$$

Define for each $r \in [0, 1]$:[22]

$$O[r] = \bigcup_{n=1}^{\infty} \mathcal{F}\left[k, \frac{r}{2^n}\right].$$

$O[r]$ is open, $\varepsilon^k \in O[r]$ and, by construction, $P(O[r] - P(C_{h_{fin}}) \leq r$. Also, we have $O[r] \subseteq O[r']$ if $r \leq r'$ (*Fig. 13.20*). Now define:[23]

$$\mathcal{R} = \bigcap_{n=1}^{\infty} O\left[\frac{1}{2^n}\right].$$

$\mathcal{R} \in \Pi_2^B$, since $\mathcal{R}$ is a countable, downward-nested intersection of open sets, and $C_{h_{fin}} \subseteq \mathcal{R}$, since each ε^i is included in each open set in the intersection (*Fig. 13.21*). So we have a nested sequence of open set around $\mathcal{R}$ whose probabilities

[22] $\overline{O[r]}$ is an example of a nowhere dense, closed set.
[23] $\mathcal{R}$ is a meager set.

converge to $P(C_{h_{fin}})$. By continuity (cf. proposition 13.1), $P(\mathcal{R}) = P(C_{h_{fin}})$. But recall from proposition 13.1 that continuity is equivalent to countable additivity in light of the other axioms. So this construction of a probability 1 approximation of $C_{h_{fin}}$ appeals to countable additivity twice. The first application yields arbitrarily close, downward nested open approximations $O_1, O_2, \ldots, O_n, \ldots$ while the second shows that the probability of the intersection $\mathcal{R}$ of these sets is the limit of the probabilities of the O_i in the sequence.

5. Probabilistic Reliability without Countable Additivity

We have seen that countable additivity suffices for probabilistic reliability over all Borel hypotheses. It remains to see that probabilistic reliability is *impossible* to achieve for some merely finitely additive probability measures.

Let's return again to Sextus's argument. Let $C_h = \{\zeta\}$, where ζ is the everywhere 0 data stream. It makes life too easy for the scientist if $P(\{\zeta\}) = 0$, so let's set $P(\{\zeta\}) = 0.5$, so $P(\overline{\{\zeta\}}) = 0.5$ as well. Now let's turn to the probabilities of the fans veering off from ζ that form a countably infinite partition of $\overline{\{\zeta\}}$. Dropping countable additivity, we no longer have to concentrate almost all of the probability on $\overline{\{\zeta\}}$ over a finite number of fans veering off from ζ. But *finite* additivity still implies that we cannot place some uniform, positive probability on each fan, or some finite union of fans would carry more than unit probability. But we may place probability 0 on each fan, so that by finite additivity, every finite union of fans veering off from ζ carries probability 0. Nonetheless, the countable union of all these fans is $\overline{\{\zeta\}}$, which must carry probability 0.5. Every union of a cofinite set of fans veering off from ζ must carry probability 0.5 as well, since finite unions of these fans carry probability 0 by finite additivity. One might doubt whether these assumptions are coherent, in the sense that there really does exist a finitely additive probability measure that satisfies them. After all, contradictions need not be obvious. But it is shown in section 13.9 that:

Proposition 13.18

There is a finitely additive probability measure P on 2^ω satisfying the constraints just described (Fig. 13.22). ■

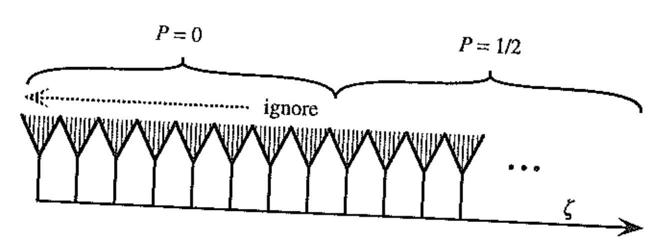

Figure 13.22

Figure 13.22 depicts an *exact reversal* of the epistemological situation guaranteed by countable additivity. Instead of *most* of the probability mass on $\overline{\{\zeta\}}$ being exhausted after a finite amount of time, we now have that *none* of the mass on $\overline{\{\zeta\}}$ is exhausted after any finite amount of time. Whereas the limit was formerly negligible, it now carries the full weight of the probability on $\overline{\{\zeta\}}$.

What happens to a conditionalizer that starts out with such a joint measure? Consider $\zeta|n$, which is a sequence of n zeros. The fan $[\zeta|n]$ includes $\{\zeta\}$. Now consider the fan $\mathcal{F}_m = [(\zeta|m)*1]$. $\mathcal{F}_m$ has a handle consisting of m 0s followed by a 1, and hence is the mth fan veering off from ζ in Figure 13.22. For each $m \geq n$, $\mathcal{F}_m \subseteq [\zeta|n]$. Since the union of all the $\mathcal{F}_m$ has probability 0.5 and $\{\zeta\}$ also has probability 0.5, finite additivity yields $P([\zeta|n]) = 1$. That is, each finite data sequence consistent with $\{\zeta\}$ has probability 1. Since $\{\zeta\} \subseteq [\zeta|n]$, we also have $P(\{\zeta\} \cap [\zeta|n]) = P(\{\zeta\})$. Hence, $P(\{\zeta\}, [\zeta|n]) = P(\{\zeta\}) = 0.5$. Conditionalization fails to approach 1 on data stream ζ, and so has a *probability* of failure of at least 0.5, despite the fact that h is refutable with certainty over all of $\mathcal{N}$ in the logical sense. Moreover, this 0.5 probability of error is never diminished through time. This shows that the measure 1 limiting success of Bayesian method reported in proposition 13.14 depends on the biasing effect of countable additivity.

We can also see from this example that the finitely additive versions of propositions 13.15(a) and 13.16(a) fail.

Proposition 13.19

Let P be the merely finitely additive measure just defined. Let $C_h = \{\zeta\}$. *Then no possible assessment method* α *can verify*$_C$ *h with certainty with a probability (according to P) greater than* 0.5.

Proof: If α never outputs 1 with certainty on ζ, then α fails with probability 0.5. So suppose α declares its certainty for h on $\zeta|n$. Then α is wrong on each data stream in the union of the fans veering off from ζ whose handles extend $\zeta|n$. The probability of this (cofinite) union of fans is 0.5, so in this case, the probability that α has committed an error is again 0.5. ■

So when we relax countable additivity, the demons of logical reliabilism can return to haunt the probabilist. The skepticism of the example can be made to appear less extreme by mixing the finitely additive measure with a more standard, countably additive measure. Then conditionalization can "learn from experience" and yet a healthy, skeptical concern about refutation in the unbounded future will always remain, if the hypothesis is true.

It remains to show that probability 1 success may be unachievable by arbitrary methods in the limiting case when countable additivity is dropped, so that the finitely additive versions of propositions 13.15(b) and 13.16(b) also fail. The following construction accomplishes this by showing that there is a Σ_2^B hypothesis and a finitely additive probability measure P such that no

possible method (conditionalization or otherwise) can refute the hypothesis in the limit with probability 1. This is the main result of the chapter.

Proposition 13.20 (with C. Juhl)

There is a merely finitely additive probability measure P on $\mathcal{N}$ *such that* h_{fin} *is not refutable*$_C$ *in the limit with probability* 1 *in P.*

Proof sketch: (a detailed proof is presented in the appendix at the end of the chapter). Let $\mathcal{S}$ denote $C_{h_{fin}} = \{\varepsilon: \varepsilon \text{ stabilizes to } 0\}$. The idea is to place discrete, nonzero probabilities summing to 0.5 on the countably many elements of $\mathcal{S}$ and to use the latitude permitted by dropping countable additivity to add a thin cloud of "virtual probability" of total mass 0.5 just outside of $\mathcal{S}$ so that any Π_2^B *superset* of $\mathcal{S}$ picks up this extra mass (and hence differs in probability from $\mathcal{S}$ by 0.5) and any Π_2^B *nonsuperset* of $\mathcal{S}$ loses the mass on some element of $\mathcal{S}$ (and hence has a probability strictly less than that of $\mathcal{S}$) (*Fig. 13.23*). Then we have that $\mathcal{S}$ has no Π_2^B approximation that differs in probability by a set of measure 0, from which the proposition follows. The bulk of the proof in section 13.9 is devoted to showing that the metaphor of the thin cloud of probability 0.5 is in fact consistent with the axioms of finitely additive probability. ■

Note that we cannot recover from this construction any *fixed* lower bound on the difference between the probability of $\mathcal{S}$ and the probability of an arbitrary Π_2^B approximation to $\mathcal{S}$ because finite additivity alone ensures that the elements of $\mathcal{S}$ carry arbitrarily small probabilities, and we cannot be sure which of them will be missed by a Π_2^B nonsuperset of $\mathcal{S}$.

The mysterious cloud of probability adhering to $\mathcal{S}$ throws a wrench into the countably additive argument (presented in section 13.4) that h_{fin} is refutable$_C$ in the limit with probability 1. It is interesting to consider where the breakdown occurs. Recall that the convergence argument in section 13.4 involves the construction of an infinite sequence of open supersets $O[1], O[1/2], O[1/4], \ldots$ of $C_{h_{fin}} = \mathcal{S}$. We observed that the probabilities of these open sets converge to the probability of $\mathcal{S}$ and then invoked continuity (proposition 13.1)

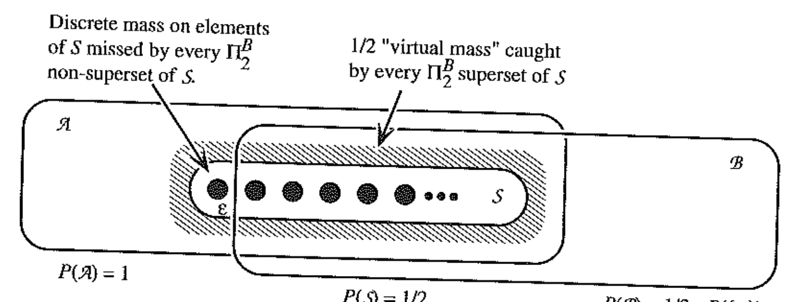

Figure 13.23

to argue that the intersection $\mathcal{R}$ of these sets has the same probability as $\mathcal{S}$. But in this case, each open superset $O[r]$ of $\mathcal{S}$ is a fortiori a Π_2^B superset of $\mathcal{S}$. Hence, each such set catches the probability cloud around $\mathcal{S}$, together with all the probability on $\mathcal{S}$, and therefore has probability 1. Thus, the limit is 1 rather than $P(\mathcal{S}) = 1/2$.

The general idea is that in countably additive measures, continuity determines the probabilities of events of higher Borel complexity as the limits of probabilities of events of lower Borel complexity. The thin clouds of probability admitted by finite additivity block this determination, and thereby block the countably additive, measure 1 approximation and convergence arguments.

6. Probabilistic Mathematics and Nonprobabilistic Science

Philosphers have long been accustomed to examine mathematical method in terms of logic and to study scientific method in terms of probability. Empiricists like Hume attempted to provide an account of this distinction. Mathematics, on their account, is the mere comparison of ideas within the mind, as though one idea could be laid on top of the other and seen directly, all at once, by the mind's eye. Empirical generalizations, on the other hand, cannot be compared against the future "all at once" so Sextus's argument applies.

Nowadays, few would assent to such a picture of mathematical method. The modern theory of computability has shown us how to view formal proof systems as computable positive tests. In chapter 6, I argued that a computer is akin to an empirical scientist in that it can scan mathematical structures only by finite chunks. Recall that the halting problem was reconstructed in chapter 6 as an instance of Sextus's skeptical argument against inductive generalization. In hindsight, it is no great step from uncomputability to Gödel's celebrated incompleteness results and beyond. One interpretation of these results is that effectively checkable mathematical proofs are limited in their power by precisely the same sorts of demonic arguments that plague inductive generalization. The difference, however, is this. In discussions of proof, we take demonic arguments seriously, while in discussions of inductive method, we discard them in sets of probability zero. To underscore this difference in attitude, let us consider what the theory of computability might look like if it were approached from the probabilistic methodologist's more lenient point of view.

Computer scientists have recognized for some time that what appears intractable from a purely logical point of view can be viewed as tractable if a measure is introduced. For example, *expected complexity theory* imposes a uniform measure over the (finite) set of inputs of a given size, as though for each input size we are uncertain which inputs of that size we will encounter. The result is that some problems that seem to be intractable in the worst case are tractable in the expected case. There also has been interest in finding tractable *random algorithms* for primality testing,[24] where

[24] E.g., Solovay and Strassen (1977) and Adleman and Manders (1977).

a random algorithm is run by a machine that flips coins as part of its operation.

But when someone turns on a computer, one evidently does so because one is uncertain about what the machine will do on the input provided. One is not uncertain about the input one will provide to the machine (the designer of the algorithm may have been uncertain about that). Nor is it that the machine is actually equipped with dice or decaying radium atoms; it is paradigmatically discrete and deterministic. The user of a deterministic machine is, in fact, uncertain about what the output will be on the input he is interested in because the operation of the program is too complicated for him to see in his mind's eye in an instant. If it weren't for this kind of uncertainty, there would be no computer industry.

One way to try to resolve the uncertainty in question is to turn on the machine, provide the input, and run it! But here we face a difficulty: how long do we wait before concluding that nothing will ever come out? In short, the halting problem becomes an empirical problem for the user of the machine. But this is just the sort of problem that can be "solved" with arbitrarily high probability along the lines discussed in section 13.3. If this is "good enough" in the empirical case, then it should be good enough for computation, and hence for formal proof.

The reader may object at this point. In the empirical case, it is no contradiction to suppose that all ravens are black, and it is no contradiction to suppose that some raven is not black. But in the computational case, if the program we know the machine to be running halts on input n in k steps, then it is a fact of elementary arithmetic that it does so. If we place probability 1 on the axioms of elementary arithmetic, the axioms of probability theory dictate that we already have probability 1 or 0 on the outcome of the computation in question.[25]

But this objection puts the computer user in a dilemma. If he does not know the outcome of the computation in advance, he is incoherent and hence irrational. And if he does know the answer in advance, going to the trouble and expense of using the computer is again irrational.[26]

Moreover, intuitively probabilistic reasoning does occur in using a computer. When undergraduates write their first computer program and it goes into an infinite loop, their degrees of belief that it will eventually halt go down rapidly, and they eventually terminate the computation and look for the bug.

The same may be said of mathematicians looking for proofs. Many mathematicians are uncertain whether P = NP, but they are inclined to believe that it is false, in the sense that they would try to prove its negation first. Such

[25] I am disregarding uncertainty about the program the machine runs, the fidelity of the machine to its program, and the input provided to the machine. The kind of uncertainty one has when one turns on a computer is clearly *not* the sum of these, or computers would never be used.

[26] Note the similarity to Plato's *Meno* paradox.

opinions suggest personal degrees of belief on an undecided proposition of arithmetic.[27]

If it is conceded, at least provisionally, that degrees of belief should be entertained over mathematical propositions, then probabilistic reliability quickly removes many skeptical concerns in the philosophy of mathematics. Let $S \subseteq \omega$ be a purely computational problem that is arithmetically definable, so that for some n, $S \in \Sigma_n^A$. Then by the universal indexing theorem,[28] there is an i such that for each x,

$$(*)\quad S(x) \Leftrightarrow \exists y_1 \, \forall y_2 \ldots T(i, \langle x, y_1, \ldots, y_n \rangle, y_{n+1}).$$

where $T(i, n, k) \Leftrightarrow$ program i halts on input x in k steps. Assuming that we are uncertain about what will happen when we feed a computer a given input x, we are in fact uncertain about the extension of the Turing predicate T. In light of our uncertainty about the extension of T, the statement (*) may be thought of as an *empirical hypothesis*, to be investigated by running different program indices on different inputs for different amounts to time, observing what happens. For simplicity, I adopt the convention that we receive a binary data stream such that a 1 or a 0 at position $\langle\langle x, y_1, \ldots, y_n\rangle y_{n+1}\rangle$ indicates whether $T(i, \langle x, y_1, \ldots, y_n\rangle, y_{n+1})$ or $\neg T(i, \langle x, y_1, \ldots, y_n\rangle, y_{n+1})$, respectively. In other words, the purely formal question "$x \in S$?" may be viewed as posing the following inductive problem:

$$C^S(\varepsilon, x) \Leftrightarrow \exists y_1 \, \forall y_2 \ldots \varepsilon_{\langle x, y_1, \ldots, y_n \rangle y_{n+1}\rangle} = 1.$$

$$\mathcal{K} = 2^\omega.$$

Suppose we have a computable method α that decides$_{C^S}$ ω with certainty. Let ε be the unique data stream corresponding to the true extension of T. It is immediate that ε is computable. Then we can decide S by running $\alpha(x, e)$ on greater and greater initial segments e of ε and passing along the first certain conjecture of α (i.e., the first one occurring after "!"). So a computable scientific decision procedure for the scientific problem generated by S yields an ordinary decision procedure for S. But if we follow the modern point of view, it should suffice to have a probabilistic solution to this inductive problem. Say that S is *probability r recursive* just in case there is a computable method α that decides$_{C^S}$ ω with probability r.

It is immediate that for each x, $C_x^S \in \Sigma_n^A$. Thus, for each x, $C_x^S \in Bo$. Let P be a probability measure on Bo. By proposition 13.14, we have that for each

[27] This discussion is closely related to ideas in Garber (1983). Garber proposes degrees of belief over entailment relations, in the spirit of the preceding discussion. He is interested in how old evidence (evidence already conditioned on) can raise the posterior probability of a new theory. His answer is that one later conditions on the formal fact that the data is entailed by the theory.

[28] Cf. proposition 7.4 and the example following proposition 7.18.

$x \in \omega$, $r > 0$ there is an $\mathcal{R} \subseteq 2^\omega$ such that $P(\mathcal{R}) = 1 - r$ and $C_x^S \cap \mathcal{R}$ is $\mathcal{R}$-clopen. Each $\mathcal{R}$-clopen subset of 2^ω is a finite union of fans (cf. exercise 4.14). Since this union can be specified by a lookup table, it follows that $C_x^S \cap \mathcal{R}$ is *$\mathcal{R}$-recursive*. So by proposition 7.8, some computable method decides ω with certainty given $\mathcal{R}$ (and hence with probability $1 - r$ given 2^ω). Thus, we see that by the probabilistic reliabilist's standards of success, Turing computability is no upper bound on what is effectively computable (up to "mere trifles" of uncertainty in the computer user).

The point is that uncomputability is about quantifiers over the natural numbers. Countably additive probability permits us to truncate the infinite universe of quantification up to "trifles," and thus to dismiss the demons of uncomputability along with those of induction. So there is some pressure on the traditional view that the stringent standards of logical reliabilism should be relaxed in the philosophy of science but not in the philosophy of mathematics. In either case, lowering standards flatters our abilities.

7. Probabilistic Theories

So far, the discussion has proceeded as though hypotheses always pick out a determinate subset of $\mathcal{N}$. But some scientific theories, including our most fundamental theory of matter, entail only probability statements concerning observable events. Moreover, even a thoroughly deterministic theory like classical mechanics can be viewed as making only probabilistic predictions if it is assumed that real-valued measurements are all subject to a normally distributed probability of error.

A logical reliabilist's approach toward such theories depends on how the probability statements are interpreted. A thoroughgoing personalist takes all probabilities to be someone's coherent degrees of belief. This is a fairly clear proposal and it causes no difficulty from a logical reliabilist's point of view. Recall that the personalist typically advocates a double reduction of probability to degree of belief and of degree of belief to betting behavior. Hence, the problem of inferring physical probabilities reduces, on this view, to the problem of reliably discovering the odds at which the subject will bet on a given event. The personalists themselves have proposed such inferential procedures in their foundational discussions of the meaning of probability. So in this sense there is no incompatibility between logical reliabilism and the purely personalistic account of physical probabilities. It is just that the data streams, to which the probabilities in quantum theory are logically connected, concern somebody's acceptance or rejection of various proferred bets.

It seems strange that the probabilities in quantum mechanics should literally refer to somebody's opinions. So it is common, even among advocates of personalism, to allow that the probabilities introduced in scientific theories refer to objective, physical realities. As we have seen in chapter 2, one such view is that physical probabilities either are or at least logically entail limiting relative frequencies of experimental outcomes. As was shown in chapters 3 and 9, this *logical* connection between probabilistic hypotheses and the data stream permits a nontrivial learning-theoretic analysis of the assessment and discovery of such hypotheses.

A careful frequentist cannot adopt all of the methods and results of Kolmogorov's probability theory, however, since many of those results are proved using countable additivity, and we have seen that countable additivity is not satisfied by limiting relative frequencies.[29] Moreover, for those who are used to the Kolmogorov axioms, frequentism appears to be an elementary fallacy. The strong law of large numbers, which is a consequence to the Kolmogorov axioms, implies that if P corresponds to independent trials with a fair coin, then $P(LRF_{\{0.5\}}(\text{heads})) = 1$. In other words, the probability that an infinite sequence of independent flips of a fair coin will have a limiting relative frequency of 0.5 is one. But a unit probability falls far short of logical necessity, as the discussion of probability one convergence in this chapter makes abundantly clear. From the point of view of a probability theorist who starts out with the Kolmogorov axioms, the frequentist simply confuses a measure one property of sequences with a logical definition of the probability of an outcome type in an infinite run of trials. The problem is compounded by the fact that there are stronger "measure one" properties than the strong law of large numbers that could have been chosen to define probability, and there is a lack of motivation for choosing one such law rather than another.[30]

Such reflections lead to the propensity account of objective probability. This account tells us just three things about propensities: they satisfy Kolmogorov's axioms; they are physical rather than mental; and whatever they are, they do not entail limiting relative frequencies of outcomes in the data stream. Evidently, this proposal permits free use of Kolmogorov's axioms and sidesteps the nettlesome project of forging a logical connection between probability and the data stream. But there seems to be a principle of conservation of mystery operating, for it becomes unintelligible on the propensity view (a) why propensities should be of any practical interest to anyone and (b) how we could ever find out about propensities with any reliability, since they are globally underdetermined in the sense of chapter 2.

The propensity theorist responds by tying personal degrees of belief to propensities by means of an axiom known as the *direct inference principle*. Let P_b be a probability measure representing personal degrees of belief. Let P be some arbitrary probability measure over the events thought to be governed by propensities. The hypothesis that propensities are distributed according to P may be abbreviated as h_P. Then the direct inference principle may be stated roughly[31]

[29] Historically, this seems to have been the decisive argument when most probability theorists dropped von Mises' theory in favor of Kolmogorov's axioms (Van Lambalgen: 1987).

[30] The *law of the iteratred logarithm* is such a theorem (Billingsly 1986). A neofrequentist could strengthen his theory by defining probability in terms of this law rather than in terms of the strong law of large numbers.

[31] A more sophisticated account is stated in Levi (1983): 255.

as follows:

$$P_b(\mathcal{E}, h_P) = P(\mathcal{E}).$$

In other words, given only that the true propensity function if P, one's degree of belief in $\mathcal{E}$ should also be the propensity of $\mathcal{E}$. Given this principle, evidence can bear on propensity hypotheses via Bayes' theorem:

$$P_b(h_P, \mathcal{S}) = \frac{P_b(h_P)P_b(\mathcal{S}, h_P)}{P_b(\mathcal{S})} \quad \text{(Bayes' theorem)}$$

$$= \frac{P_b(h_P)P(\mathcal{E})}{P_b(\mathcal{S})} \quad \text{(the direct inference principle).}$$

The direct inference principle also makes propensities figure into one's personal expected utilities for various acts, and hence makes them relevant to practical deliberation. Without it, propensities are irrelevant, inaccessible, metaphysical curiosities, whether or not they exist (*Fig. 13.24*). There are some who maintain that propensities are no worse than any other theoretical entities in this regard. But I reject the analogy. Theoretical entities can be unobservable in principle and yet figure into determinate, logical predictions about observable events. For example, the light waves in Fresnel's optical theory figured into deductions of exact positions for diffraction bands. Propensities are very different, in that a propensity theory entails only propensity consequences, which in turn bear no logical connection to the data stream, even in the limit. Without the direct inference principle, propensities are entirely insulated from experience.

But now the question is: Why accept the principle? Some defend it

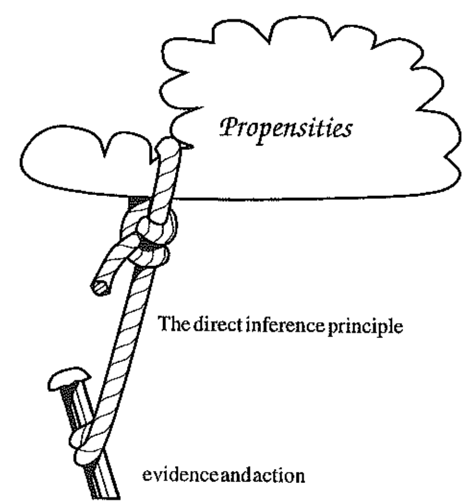

Figure 13.24

as an "incorrigible background stipulation" that requires no justification.[32] But I am not so sure. Suppose there is an omnipotent, omniscient god. Suppose this god has an infinite tape marked out with squares labeled by all the propositions expressible in some language more expressive than any we will ever use. The god also has an infinite supply of infinitely divisible mud. Then, for his own amusement, he places a mud pie on each tape square so that the masses of the pies satisfy Kolmogorov's axioms with respect to the corresponding propositions (normalizing to unit mass on the tautology). After this is done, he determines the truth values of all the sentences by zapping some universe into existence, without giving any thought to his mud pies.

For all that has been said about propensities, they could be the masses of the god's mud pies. These masses are real and they satisfy the Kolmogorov axioms. Certainly, they don't entail limiting relative frequencies for outcomes of experiments. The question is: Why should I, in accordance with the direct inference principle, contribute 2.4% of the stake on a bet that an alpha particle will be emitted from a given uranium atom because the mass of the mud pie on that proposition is 0.024? I see no reason.

In a recent text on Bayesian philosophy of science, Howson and Urbach[33] argue for the direct infererence principle by interpreting statistical probabilities as limiting relative frequencies and pointing out that those who bet in accordance with known limiting relative frequencies will minimize their limiting relative frequency of losses. This defense of the principle succeeds, but only by undercutting its importance, since the admission that propensities entail limiting relative frequencies opens propensities to empirical investigation along the lines discussed in chapter 3, with no detour through personal probabilities or direct inference. Of course, it is still possible on their approach to study, along the lines of section 2 of this chapter, whether conditionalization together with the direct inference principle is as reliable as other methods might be, both for ideal and for computable agents, but in such a study the direct inference principle is a restriction on one's choice of possible inductive strategies rather than the theoretical linchpin of statistical inference.

In summary, logical reliabilism is compatible with the foundations both of frequentism and of pure personalism, even though there might be quibbles about actual practice (e.g., with avowed frequentists who nonetheless assume countable additivity). It is harder to square the logical reliabilist approach with the increasingly popular hybrid view in which globally underdetermined propensities are connected to personal degrees of belief by the direct inference principle. That position avoids standard objections raised against its competitors, but it does so at a cost that should not be neglected. According to it, science is arbitrary (but coherent) opinion about propensities that entail nothing about what will be observed (except for more propensities, about which the same may be said), even in the limit.

[32] Levi (1983): 255.
[33] Howson and Urbach (1990).

I do not mean to suggest that hybrid probabilism is a mistake. I only wish to make it clear that the apparent advantages of the view (e.g., its correspondence with statistical practice and its ability to overcome global underdetermination) come at the expense of a clear, logical account of what science is about, of what its import is for observation, and of how scientific method leads to the truth. Logical reliabilism begins with a clear, logical account of precisely these issues. Accordingly, its strengths and weaknesses are somewhat complementary to those of today's hybrid probabilism. Logical reliability results are hard to apply in particular cases because it is hard to say in particular cases what the empirical import of a hypothesis actually is. But on the other hand, the relations between empirical content, underdetermination, computability, methodological recommendations, and the reliability of inquiry as a means for finding the truth all stand out in bold relief.

At the very least, logical reliabilism provides a vantage point from which to raise questions that do not arise naturally within the hydrid probabilist perspective; questions such as whether updating one joint probability measure would be more reliable than updating another concerning a given hypothesis. More ambitiously, logical reliabilism is a persistent remnant of the ancient idea that scientific inquiry should have a clear logical connection with finding the truth, in an age in which both philosophy and professional methodology urge us to forget about it.

8. Logic and Probability

On the face of it, the probabilistic approach to induction seems much more practical and natural than the logical one presented in this book. My purpose in this chapter has been to clarify the relationship between the two perspectives and to urge that some of the apparent advantages of the probabilistic approach may be merely apparent. The first point was that probabilistic methods may fail in the logical sense to solve problems solvable by simple, nonprobabilistic methods. Then I examined some probabilistic convergence theorems and showed how they depend on countable additivity, a powerful epistemological axiom rejected by some personalists and inconsistent with frequentism that is often advertised as a mere technical convenience. To address the objection that logical reliabilism is a cranky, skeptical viewpoint that is never taken seriously, I appealed to the analogy between induction and computation that was developed in chapter 6. According to this analogy, the demonic arguments of logical reliabilism are routinely taken seriously in the philosophy of mathematics. Finally, I reviewed the standard interpretations of stochastic theories, arguing that the logical reliabilist perspective is compatible with pure personalism, pure frequentism, and mixtures of personalism and frequentism. The view that mixes personalism and propensity theory is founded on the direct inference principle, which is questionable exactly insofar as it is required to link propensities to evidence and to practical action.

9. Proofs of Propositions 13.18 and 13.20

Proposition 13.18

Let $\mathcal{F}_n = [\zeta|n*1]$. Then there is a finitely additive probability measure P on the power set of 2^ω such that

(i) $P(\zeta) = 0.5$,

(ii) *For each cofinite union $\mathcal{S}$ of fans in $\{\mathcal{F}_n: n \in \omega\}$, $P(\mathcal{S}) = 0.5$, and*

(iii) *For each finite union $\mathcal{S}$ of fans in $\{\mathcal{F}_n: n \in \omega\}$, $P(\mathcal{S}) = 0$.*

Proof: Define

Θ is an algebra on $\mathcal{N} \Leftrightarrow \mathcal{N} \in \Theta$ and Θ is closed under finite union and complementation.

Let Θ be an algebra on $\mathcal{N}$,

P is a finitely additive probability measure on $\Theta \Leftrightarrow$

(1) *for each $\mathcal{S} \in \Theta$, $P(\mathcal{S}) \geq 0$,*

(2) *$P(\mathcal{N}) = 1$ and*

(3) *If $\mathcal{S}_0, \mathcal{S}_1, \ldots, \mathcal{S}_n$ is a finite sequence of pairwise disjoint elements of Θ, then $\bigcup_{i=1}^{n} \mathcal{S}_i \in \Theta \Rightarrow P\left(\bigcup_{i=1}^{n} \mathcal{S}_i\right) = \sum_{i=1}^{n} P(\mathcal{S}_i)$.*

Lemma 13.21

If P is a finitely additive probability measure on an algebra Θ on $\mathcal{N}$, then there is a finitely additive P′ defined on the entire power set of $\mathcal{N}$ such that for all $\mathcal{S} \in \Theta$, $P(\mathcal{S}) = P'(\mathcal{S})$. (Unlike the countably additive case, the extension is not necessarily unique.)

Proof: Consequence of the Hahn Banach theorem (Ash 1972). ■

In light of lemma 13.21, it suffices to prove that there is a finitely additive measure satisfying the intended constraints over some algebra of subsets of $\mathcal{N}$. Define:

Ξ = *the set of all cofinite unions of fans in* $\{\mathcal{F}_n: n \in \omega\}$.

Θ = *the least algebra containing* $\{\zeta\} \cup \Xi$.

Consider two different types of elements of Θ.

$Q \in A \Leftrightarrow$ *for some* $\mathcal{R} \in \Xi$, $\mathcal{R} \subseteq Q$.

$Q \in co\text{-}A \Leftrightarrow \bar{Q} \in A$.

Lemma 13.22

$\{A, co\text{-}A\}$ partitions Θ.

Proof: $A \cap co\text{-}A = \varnothing$, since if Q contains a cofinite union of fans in $\{\mathcal{F}_n: n \in \omega\}$ then $\bar{Q}$ does not, and conversely. $\{\zeta\} \in co\text{-}A$ and each $\mathcal{R} \in \Xi$ is in A. Let Q_1, $Q_2 \in \Theta$. Suppose $Q_1, Q_2 \in A$. Then $Q_1 \cup Q_2 \in A$ and $\overline{Q_1} \in co\text{-}A$. If $Q_1 \in A$ and $Q_2 \in co\text{-}A$ then $Q_1 \cup Q_2 \in A$. If Q_1, $Q_2 \in co\text{-}A$ then $Q_1 \cup Q_2 \in co\text{-}A$ and $\bar{Q}_1 \in A$. So by induction, each element of Θ is in $A \cup co\text{-}A$. ■

Define for each $Q \in \Theta$:

$$P(Q) = \begin{cases} 1 & \text{if } Q \in A \text{ and } \zeta \in Q \\ 0.5 & \text{if } Q \in A \text{ and } \zeta \notin Q \\ 0.5 & \text{if } Q \in Co\text{-}A \text{ and } \zeta \in Q \\ 0 & \text{if } Q \in Co\text{-}A \text{ and } \zeta \notin Q. \end{cases}$$

In light of lemma 13.22, P is uniquely defined and total over Θ. P also has properties (i), (ii), and (iii) of the proposition.

We now verify that P is a finitely additive probability measure on Θ. By lemma 13.21, the proposition follows. Number the conditions in the definition of P from 1 to 4. We must show that finite additivity is satisfied when Q_1, Q_2 satisfy each possible combination of such conditions. The combinations are (1, 1), (1, 2), (1, 3), (1, 4), (2, 2), (2, 3), (2, 4), (3, 3), (3, 4), (4, 4). If $Q_1, Q_2 \in A$ then $Q_1 \cap Q_2 \neq \varnothing$, so finite additivity is trivially satisfied. This accounts for (1, 1), (1, 2), (2, 2). If $\zeta \in Q_1$, $\zeta \in Q_2$ then $Q_1 \cap Q_2 \neq \varnothing$, so finite additivity is satisfied. This accounts for (1, 3), (3, 3). In case (1, 4), $P(Q_1) = 1$ and $P(Q_2) = 0$. $Q_1 \cup Q_2$ satisfies condition 1 so $P(Q_1 \cup Q_2) = 1 = P(Q_1) + P(Q_2)$. In case (2, 3), $P(Q_1) = 1/2$ and $P(Q_2) = 1/2$. $Q_1 \cup Q_2$ satisfies condition 1 so $P(Q_1 \cup Q_2) = 1 = P(Q_1) + P(Q_2)$. In case (2, 4), $P(Q_1) = 1/2$ and $P(Q_2) = 0$. $Q_1 \cup Q_2$ satisfies condition 2 so $P(Q_1 \cup Q_2) = 1/2 = P(Q_1) + P(Q_2)$. In case (3, 4), $P(Q_1) = 1/2$ and $P(Q_2) = 0$. $Q_1 \cup Q_2$ satisfies condition 3 so $P(Q_1 \cup Q_2) = 1/2 = P(Q_1) + P(Q_2)$. In case (4, 4), $P(Q_1) = 0$ and $P(Q_2) = 0$. $Q_1 \cup Q_2$ satisfies condition 4 so $P(Q_1 \cup Q_2) = 0 = P(Q_1) + P(Q_2)$. ■

Proposition 13.20 (with C. Juhl)

There is a merely finitely additive probability measure P defined on $2^{\mathcal{N}}$ such that h_{fin} is not refutable$_C$ in the limit with probability 1 in P.

Proof: Let $\mathcal{S} = C_{h_{fin}}$. Define:

$$\Theta_0 = \{\mathcal{S}\} \cup \{\mathcal{P} \in \Pi_2^B: \mathcal{S} \subseteq \mathcal{P}\} \cup \{\{\varepsilon\}: \varepsilon \in \mathcal{S}\}.$$

Let Θ be the least algebra containing all elements of Θ_0. Let $\mathcal{Q} \in \Theta$. Then define:

$$\mathcal{Q} \in A \Leftrightarrow \exists \mathcal{R} \in \Pi_2^B \text{ such that } \mathcal{S} \subseteq \mathcal{R} \text{ and } \mathcal{R} - \mathcal{S} \subseteq \mathcal{Q}.$$

$$\mathcal{Q} \in B \Leftrightarrow \exists \text{ finite } \mathcal{F} \subseteq \mathcal{S} \text{ such that } \mathcal{S} - \mathcal{F} \subseteq \mathcal{Q} \text{ and } \mathcal{F} \subseteq \bar{\mathcal{Q}}.$$

$$\mathcal{Q} \in co\text{-}B \Leftrightarrow \bar{\mathcal{Q}} \in B.$$

Now, define four subclasses of Θ:

$\Gamma_1 = A \cap B$ (*Fig. 13.25*).

$\Gamma_2 = A \cap co\text{-}B$ (*Fig. 13.26*).

$\Gamma_3 = \bar{A} \cap B$ (*Fig. 13.27*).

$\Gamma_4 = \bar{A} \cap co\text{-}B$ (*Fig. 13.28*).

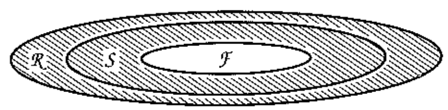

Figure 13.25

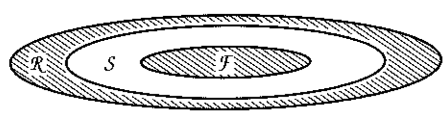

Figure 13.26

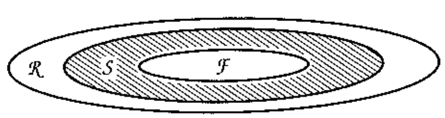

Figure 13.27

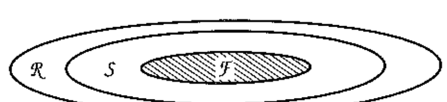

Figure 13.28

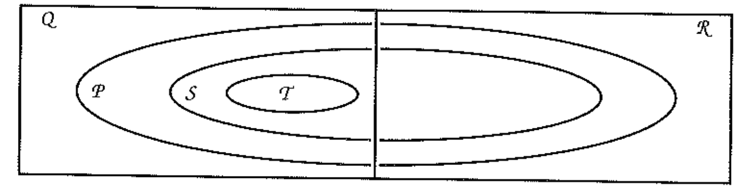

Figure 13.29

Lemma 13.23 (Closure laws)

For each $\mathcal{Q}, \mathcal{R} \in \Theta$:

(a) *If* $\mathcal{Q} \in A$ *and* $\mathcal{R} \in A$ *then* $\mathcal{Q} \cup \mathcal{R} \in A$.

(b) *If* $\mathcal{Q} \in \bar{A}$ *and* $\mathcal{R} \in \bar{A}$ *then* $\mathcal{Q} \cup \mathcal{R} \in \bar{A}$.

(c) *If* $\mathcal{Q} \in A$ *and* $\mathcal{R} \in \bar{A}$ *then* $\mathcal{Q} \cup \mathcal{R} \in A$.

(d) *If* $\mathcal{Q} \in B$ *and* $\mathcal{R} \in B$ *then* $\mathcal{Q} \cup \mathcal{R} \in B$.

(e) *If* $\mathcal{Q} \in co\text{-}B$ *and* $\mathcal{R} \in co\text{-}B$ *then* $\mathcal{Q} \cup \mathcal{R} \in co\text{-}B$.

(f) *If* $\mathcal{Q} \in B$ *and* $\mathcal{R} \in co\text{-}B$ *then* $\mathcal{Q} \cup \mathcal{R} \in B$.

Proof: All these results are immediate except for (b). Let $co\text{-}\Theta_0$ denote the set of complements of elements of Θ_0. A *term* is a finite intersection of sets in $\Theta_0 \cup co\text{-}\Theta_0$. Each $\mathcal{Q} \in \Theta$ can be expressed as a finite union of terms.

Suppose $\mathcal{Q} \in \bar{A}$ and $\mathcal{R} \in \bar{A}$. Write $\mathcal{Q} = \bigcup X$ and $\mathcal{R} = \bigcup Y$, where X and Y are finite sets of terms. In each term, we can intersect all the finite subsets[34] of $\mathcal{S}$ into one finite subset of $\mathcal{S}$, all the complemented, finite subsets of $\mathcal{S}$ into one complement of a finite subset of $\mathcal{S}$, all the Π_2^B supersets of $\mathcal{S}$ into one Π_2^B superset of $\mathcal{S}$, and all the complements of Π_2^B supersets of $\mathcal{S}$ into complements of Π_2^B supersets of $\mathcal{S}$, since the finite sets, the cofinite sets, the Π_2^B supersets of $\mathcal{S}$ and the complements of Π_2^B supersets of $\mathcal{S}$ are all closed under finite intersection. We may now assume that each term in $(X \cup Y)$ is expressed in this normal form.

Suppose for reductio that $\mathcal{Q} \cup \mathcal{R} \in A$, so that $\bigcup (X \cup Y) \in A$ (*Fig. 13.29*). Suppose some term $\mathcal{T}$ in $X \cup Y$ involves an uncomplemented subset of $\mathcal{S}$. Since $\bigcup (X \cup Y) \in A$, there is a $\mathcal{P} \in \Pi_2^B$ such that $\mathcal{P} - \mathcal{S} \subseteq \bigcup (X \cup Y)$. But no element of $\mathcal{T}$ is in $\mathcal{P} - \mathcal{S}$, so $\mathcal{P} - \mathcal{S} \subseteq \bigcup ((X \cup Y) - \{\mathcal{T}\})$ and hence $\bigcup ((X \cup Y) - \{\mathcal{T}\}) \in A$.

Let K be the result of removing from $X \cup Y$ all terms in which uncomplemented subsets of $\mathcal{S}$ occur. Hence, $\bigcup K \in A$. The remaining terms in K are of one of two types: type (i) terms are of form $\mathcal{A} \cap \mathcal{B}$ and type (ii) terms are of form $\bar{\mathcal{A}} \cap \mathcal{B}$, where $\mathcal{A}$ is a Π_2^B superset of $\mathcal{S}$ and $\mathcal{B}$ is the complement of a subset of $\mathcal{S}$. Since $\bar{\mathcal{A}} \subseteq \mathcal{B}$, $\bar{\mathcal{A}} \cap \mathcal{B} = \bar{\mathcal{A}}$. So type (ii) terms reduce to the form $\bar{\mathcal{A}}$, where $\mathcal{A}$ is a Π_2^B superset of $\mathcal{S}$.

[34] These are all singletons.

Suppose for reductio that K consists of only type (ii) terms. Then $\bigcup K$ is a finite union of complements of Π_2^B supersets of $\mathcal{S}$. So there are $\mathcal{G}_0, \ldots, \mathcal{G}_n \in \Pi_2^B$ such that for each $i \leq n$, $\mathcal{S} \subseteq \mathcal{G}_i$, and

$$\bigcup K = \bigcup_{i=0}^{n} \overline{\mathcal{G}_i}.$$

Also let $\mathcal{Y}$ be an arbitrary Π_2^B superset of $\mathcal{S}$. Since $\mathcal{S} \notin \Pi_2^B$ and Π_2^B is closed under countable intersection, we have that

$$(\mathcal{Y} - \mathcal{S}) - \bigcup K = (\mathcal{Y} - \mathcal{S}) - \bigcup_{i=0}^{n} \overline{\mathcal{G}_i} = \left(\mathcal{Y} \cap \bigcap_{i=0}^{n} \mathcal{G}_i\right) - \mathcal{S} \neq \varnothing.$$

So $(\mathcal{Y} - \mathcal{S})$ is not a subset of $\bigcup K$. Since $\mathcal{Y}$ is an arbitrary Π_2^B superset of $\mathcal{S}$, $\bigcup K \notin A$. Hence, $\bigcup (X \cup Y) \notin A$. Contradiction. So K involves at least one type (i) term $\mathcal{A} \cap \mathcal{B}$. $\mathcal{A} \cap \mathcal{B} \in A$ since $\mathcal{A}$ is a Π_2^B superset of $\mathcal{S}$ and $\mathcal{A} - \mathcal{S} \subseteq \mathcal{A} \cap \mathcal{B}$. But the term $\mathcal{A} \cap \mathcal{B}$ is either a subset of $\mathcal{Q}$ or of $\mathcal{R}$, so either $\mathcal{Q}$ or $\mathcal{R}$ is in A. Contradiction. ■

Lemma 13.24

$\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4\}$ *partitions* Θ.

Proof: First it is established by exhaustion that the Γs are pairwise disjoint. Suppose $\mathcal{Q} \in B \cap co\text{-}B$. Then $\exists$ finite $\mathcal{F}_1 \subseteq \mathcal{S}$ such that $\mathcal{S} - \mathcal{F}_1 \subseteq \mathcal{Q}$ and $\mathcal{F}_1 \subseteq \overline{\mathcal{Q}}$ and there is an $\mathcal{F}_2$ such that $\mathcal{S} - \mathcal{F}_2 \subseteq \overline{\mathcal{Q}}$ and $\mathcal{F}_2 \subseteq \mathcal{Q}$. But since $\mathcal{S}$ is infinite and $\mathcal{F}_1, \mathcal{F}_2$ are finite, there is some $\varepsilon \in \mathcal{S} - (\mathcal{F}_1 \cup \mathcal{F}_2)$ such that $\varepsilon \in \mathcal{Q}$ and $\varepsilon \in \overline{\mathcal{Q}}$, which is a contradiction. Hence $B \cap co\text{-}B = \varnothing$. Evidently, $A \cap \bar{A} = \varnothing$. Thus, the Γs are pairwise disjoint.

Now it is shown by induction that $\Theta \subseteq \Gamma_1 \cup \Gamma_2 \cup \Gamma_3 \cup \Gamma_4$. Evidently, $\Theta \subseteq A \cup \bar{A}$, so it suffices to show that $\Theta \subseteq B \cup co\text{-}B$.

Base case: For each $\varepsilon \in \mathcal{S}$, $\{\varepsilon\} \in co\text{-}B$, *choosing* $\mathcal{F} = \{\varepsilon\}$. $\mathcal{S} \in B$, choosing $\mathcal{F} = \varnothing$. Let $\mathcal{R}$ be a Π_2^B superset of $\mathcal{S}$. Then $\mathcal{R} \in B$, choosing $\mathcal{F} = \varnothing$.

Induction: By definition, $\mathcal{Q} \in B \Leftrightarrow \overline{\mathcal{Q}} \in co\text{-}B$. Union is covered by the preceding lemma. ■

Lemma 13.25

For each $\mathcal{Q} \in \Theta$, *there is a unique, finite* $\mathcal{F} \subseteq \mathcal{S}$ *such that* $\mathcal{S} - \mathcal{F} \subseteq \mathcal{Q}$ *and* $\mathcal{F} \subseteq \overline{\mathcal{Q}}$ *or* $\mathcal{S} - \mathcal{F} \subseteq \overline{\mathcal{Q}}$ *and* $\mathcal{F} \subseteq \mathcal{Q}$.

Proof: By the preceding lemma, $\{B, co\text{-}B\}$ partitions Θ. Suppose $\mathcal{Q} \in B$. *Then* (*) $\exists \mathcal{F}$ such that $\mathcal{S} - \mathcal{F} \subseteq \mathcal{Q}$ and $\mathcal{F} \subseteq \overline{\mathcal{Q}}$. Suppose distinct $\mathcal{F}, \mathcal{G}$ satisfy (*). Without loss of generality, suppose $\varepsilon \in \mathcal{F} - \mathcal{G}$. Since $\mathcal{F} \subseteq \overline{\mathcal{Q}}$, $\varepsilon \in \overline{\mathcal{Q}}$ and $\varepsilon \in \mathcal{S}$. Since $\mathcal{S} - \mathcal{G} \subseteq \mathcal{Q}$ and $\varepsilon \in \mathcal{S} - \mathcal{G}$, $\varepsilon \in \mathcal{Q}$. Contradiction. The same argument applied to $\overline{\mathcal{Q}}$ works when $\mathcal{Q} \in co\text{-}B$. ■

If $\mathcal{Q} \in \Theta$, then let $\mathcal{F}_{\mathcal{Q}}$ denote the unique $\mathcal{F}$ guaranteed by lemma 13.25. Now I define a finitely additive probability measure P on Θ. Enumerate $\mathcal{S}$ as $\varepsilon[1]$, $\varepsilon[2], \ldots, \varepsilon[n], \ldots$. Let

$$P'(\{\varepsilon[i]\}) = 1/2^{i+1}.$$

For each finite $\mathcal{F} \subseteq \mathcal{S}$ define:

$$P'(\mathcal{F}) = \sum_{\varepsilon \in \mathcal{F}} P'(\{\varepsilon\}).$$

Now for each $\mathcal{Q} \in \Theta$, define:

$$P(\mathcal{Q}) = \begin{cases} 1 - P'(\mathcal{F}_{\mathcal{Q}}) & \text{if } \mathcal{Q} \in \Gamma_1 \\ 0.5 + P'(\mathcal{F}_{\mathcal{Q}}) & \text{if } \mathcal{Q} \in \Gamma_2 \\ 0.5 - P'(\mathcal{F}_{\mathcal{Q}}) & \text{if } \mathcal{Q} \in \Gamma_3 \\ P'(\mathcal{F}_{\mathcal{Q}}) & \text{if } \mathcal{Q} \in \Gamma_4. \end{cases}$$

P is uniquely defined and total on Θ since the Γs partition Θ (lemma 13.24), and $\mathcal{F}_{\mathcal{Q}}$ is uniquely defined on Θ (lemma 13.25). We need to show coherence over Θ.

Lemma 13.26

P is a finitely additive probability measure on Θ.

Proof: Let $\mathcal{Q}, \mathcal{R} \in \Theta$. If both $\mathcal{Q}$ and $\mathcal{R}$ are in A or if both $\mathcal{Q}$ and $\mathcal{R}$ are in B, then $\mathcal{Q} \cap \mathcal{R} \neq \varnothing$, so finite additivity is trivailly satisfied. This leaves just the cases (a) $\mathcal{Q} \in \Gamma_1$ and $\mathcal{R} \in \Gamma_4$; (b) $\mathcal{Q} \in \Gamma_2$ and $\mathcal{R} \in \Gamma_3$; (c) $\mathcal{Q} \in \Gamma_3$ and $\mathcal{R} \in \Gamma_4$; and finally (d) $\mathcal{Q} \in \Gamma_4$ and $\mathcal{R} \in \Gamma_4$.

In case (a) (*Fig. 13.30*), $P(\mathcal{Q}) = 1 - P'(\mathcal{F}_{\mathcal{Q}})$ and $P(\mathcal{R}) = P'(\mathcal{F}_{\mathcal{R}})$. By lemma 13.23, $\mathcal{Q} \cup \mathcal{R} \in \Gamma_1$ so $P(\mathcal{Q} \cup \mathcal{R}) = 1 - P'(\mathcal{F}_{\mathcal{Q} \cup \mathcal{R}})$. Suppose $\mathcal{Q} \cap \mathcal{R} = \varnothing$. (Note:

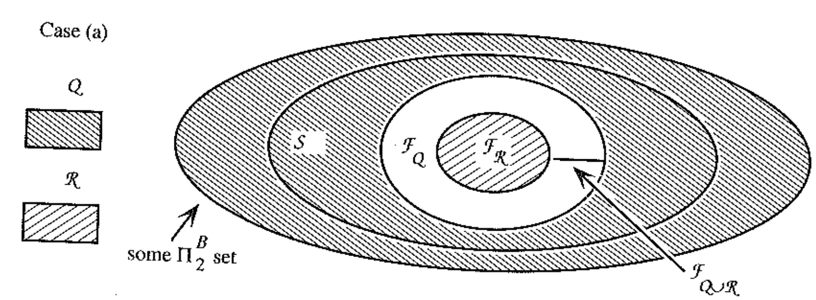

Figure 13.30

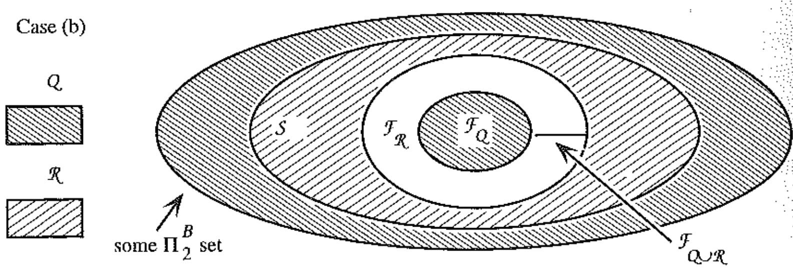

Figure 13.31

therefore $\mathcal{Q}$ and $\mathcal{R}$ cannot overlap in Figure 13.30.) Then $\mathcal{F}_\mathcal{R} \subseteq \mathcal{F}_\mathcal{Q}$. So $\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}} = \mathcal{F}_\mathcal{Q} - \mathcal{F}_\mathcal{R}$ and hence $P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = P'(\mathcal{F}_\mathcal{Q}) - P'(\mathcal{F}_\mathcal{R})$. So $P(\mathcal{Q}\cup\mathcal{R}) = 1 - P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = 1 - [P'(\mathcal{F}_\mathcal{Q}) - P'(\mathcal{F}_\mathcal{R})] = P(\mathcal{Q}) + P(\mathcal{R})$.

In case (b) (*Fig. 13.31*), $P(\mathcal{Q}) = 1/2 + P'(\mathcal{F}_\mathcal{Q})$ and $P(\mathcal{R}) = 1/2 - P'(\mathcal{F}_\mathcal{R})$. By lemma 13.23, $\mathcal{Q}\cup\mathcal{R} \in \Gamma_1$, so $P(\mathcal{Q}\cup\mathcal{R}) = 1 - P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}})$. Suppose $\mathcal{Q}\cap\mathcal{R} = \varnothing$. Then $\mathcal{F}_\mathcal{Q} \subseteq \mathcal{F}_\mathcal{R}$. So $\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}} = \mathcal{F}_\mathcal{R} - \mathcal{F}_\mathcal{Q}$ and hence $P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = P'(\mathcal{F}_\mathcal{R}) - P'(\mathcal{F}_\mathcal{Q})$. So $P(\mathcal{Q}\cup\mathcal{R}) = 1 - P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = 1 - [P'(\mathcal{F}_\mathcal{R}) - P'(\mathcal{F}_\mathcal{Q})] = 1/2 + 1/2 + P'(\mathcal{F}_\mathcal{Q}) - P'(\mathcal{F}_\mathcal{R}) = P(\mathcal{Q}) + P(\mathcal{R})$.

In case (c) (*Fig. 13.32*), $P(\mathcal{Q}) = 1/2 - P'(\mathcal{F}_\mathcal{Q})$ and $P(\mathcal{R}) = P'(\mathcal{F}_\mathcal{R})$. By lemma 13.23, $\mathcal{Q}\cup\mathcal{R} \in \Gamma_3$ so $P(\mathcal{Q}\cup\mathcal{R}) = 1/2 - P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}})$. Suppose $\mathcal{Q}\cap\mathcal{R} = \varnothing$. Then $\mathcal{F}_\mathcal{R} \subseteq \mathcal{F}_\mathcal{Q}$. So $\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}} = \mathcal{F}_\mathcal{Q} - \mathcal{F}_\mathcal{R}$ and hence $P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = P'(\mathcal{F}_\mathcal{Q}) - P'(\mathcal{F}_\mathcal{R})$. Hence, $P(\mathcal{Q}\cup\mathcal{R}) = 1/2 - P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = 1/2 - [P'(\mathcal{F}_\mathcal{Q}) - P'(\mathcal{F}_\mathcal{R})] = 1/2 - P'(\mathcal{F}_\mathcal{Q}) + P'(\mathcal{F}_\mathcal{R}) = P(\mathcal{Q}) + P(\mathcal{R})$.

In case (d) (*Fig. 13.33*), $P(\mathcal{Q}) = P'(\mathcal{F}_\mathcal{Q})$ and $P(\mathcal{R}) = P'(\mathcal{F}_\mathcal{R})$. By lemma 13.23, $\mathcal{Q}\cup\mathcal{R} \in \Gamma_4$ so $P(\mathcal{Q}\cup\mathcal{R}) = P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}})$. Suppose $\mathcal{Q}\cap\mathcal{R} = \varnothing$. Then since $\mathcal{Q} = \mathcal{F}_\mathcal{Q}$ and $\mathcal{R} = \mathcal{F}_\mathcal{R}$, $\mathcal{F}_\mathcal{Q} \cap \mathcal{F}_\mathcal{R} = \varnothing$ and $\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}} = \mathcal{F}_\mathcal{Q} \cup \mathcal{F}_\mathcal{R}$. Hence, $P'(\mathcal{F}_{\mathcal{Q}\cup\mathcal{R}}) = P'(\mathcal{F}_\mathcal{Q}) + P'(\mathcal{F}_\mathcal{R}) = P(\mathcal{Q}) + P(\mathcal{R})$. ■

By lemma 13.26 and lemma 13.21, there exists a finitely additive probability measure P'' extending P to the power set of $\mathcal{N}$. Now, let α be an arbitrary assessment method. Let $\mathcal{S}tab_0(\alpha) = \{\varepsilon: \alpha$ stabilizes to 0 on $\varepsilon\}$. Let $\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) = \overline{\mathcal{S}tab_0(\alpha)} - \mathcal{S}$, let $\mathcal{F}ail_\mathcal{S}(\alpha) = \mathcal{S}tab_0(\alpha) \cap \mathcal{S}$, and let $\mathcal{F}ail(\alpha) = \mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) \cup \mathcal{F}ail_\mathcal{S}(\alpha)$. Since $C_h = \mathcal{S}$, $\mathcal{F}ail(\alpha)$ is the set of all data streams on which α fails to refute$_C$ h in the limit, either by stabilizing to 0 on an element

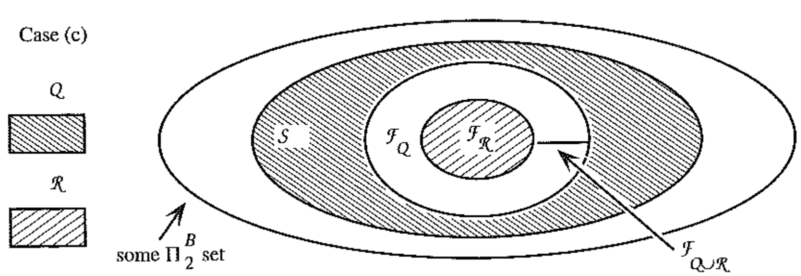

Figure 13.32

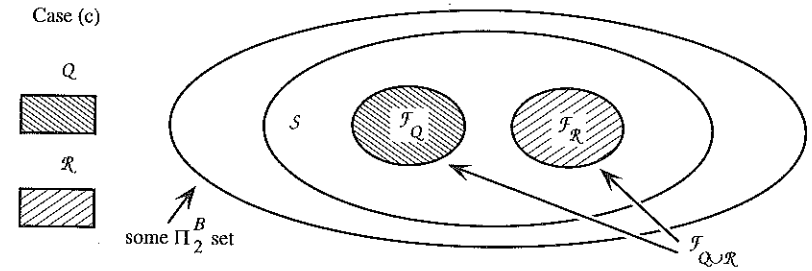

Figure 13.33

of C_h or by failing to stabilize to zero on an element of $\overline{C_h}$. Since h_{fin} is not refutable$_C$ in the limit (cf. chapter 3) we have that $\mathcal{F}ail(\alpha) \neq \varnothing$. We need to show that $P''(\mathcal{F}ail(\alpha)) > 0$. Case I: $\mathcal{F}ail_\mathcal{S}(\alpha) \neq \varnothing$. Let $\varepsilon \in \mathcal{F}ail_\mathcal{S}(\alpha)$. Then $\varepsilon \in \mathcal{S}$, so for some i, $\varepsilon = \varepsilon[i]$. So $P''(\{\varepsilon\}) = P'(\{\varepsilon[i]\}) = 1/2^{i+1} > 0$. Case II: $\mathcal{F}ail_\mathcal{S}(\alpha) = \varnothing$. Then since $\mathcal{F}ail(\alpha) \neq \varnothing$, we have $\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) \neq \varnothing$. Since $\mathcal{F}ail_\mathcal{S}(\alpha) = \varnothing$, (a) $\mathcal{S} \subseteq \overline{\mathcal{S}tab_0(\alpha)}$. Also, (b) $\overline{\mathcal{S}tab_0(\alpha)} \in \Pi_2^B$, since we can define:

$$\varepsilon \in \overline{\mathcal{S}tab_0(\alpha)} \Leftrightarrow \forall n\, \exists m > n,\ \alpha(h_{fin}, \varepsilon|m) \neq 0.$$

Then since $\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) = \overline{\mathcal{S}tab_0(\alpha)} - \mathcal{S}$, we have that $\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) \in A$. Since $\mathcal{S} \subseteq \mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha)$, $\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) \in co\text{-}B$. So $\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha) \in \Gamma_2$. So by the definition of P'', $P''(\mathcal{F}ail_{\bar{\mathcal{S}}}(\alpha)) = 0.5 > 0$. So no method α can refute$_C$ h_{fin} in the limit given 2^ω with probability 1 in P''. ■

Exercises

13.1. Prove proposition 13.2.

13.2 Prove proposition 13.1. (Hint: let $Q_0, Q_1, \ldots, Q_n, \ldots$ be an ω-sequence of mutually disjoint Borel sets sets.) Use the fact that

$$\bigcup_{i=0}^{\infty} Q_i = \overline{\bigcap_{n=0}^{\infty}\left(\overline{\bigcup_{i=0}^{n} Q_i}\right)}.$$

Then observe that the intersection is over a downward nested sequence of Borel sets and apply continuity to the probability of the intersection. On the other side, let $Q_0, Q_1, \ldots, Q_n, \ldots$ be a downward nested sequence of Borel sets. Use the fact that

$$\bigcap_{i=0}^{\infty} Q_i = Q_0 - \bigcup_{i=0}^{\infty} (Q_i - Q_{i+1}).$$

Then apply countable additivity to express the probability of the union as a limit of finite sums.)

13.3. For each $r \in (0, 1)$, let P_r be the probability measure on 2^ω corresponding to independent flips of a coin that comes up 1 on each trial with probability r. Suppose that Fred's initial joint probability measure P is a mixture of finitely many P_rs. Is Fred reliable concerning the hypothesis constructed in proposition 13.5?

13.4 Show that gradual identification can be accomplished with probability 1. (Hint: apply proposition 13.17 to the degrees of approximation of each hypothesis.)

13.5. Define

P is a countably additive probability measure on $\Delta_1^B \Leftrightarrow$

(1′) $\forall \mathcal{S} \in \Delta_1^B,\ P(\mathcal{S}) \geq 1$,

(2′) $P(\mathcal{N}) = 1$, *and*

(3′) $\forall$ *sequence* $\mathcal{S}_0, \mathcal{S}_1, \ldots, \mathcal{S}_n, \ldots$ *of pairwise disjoint* Δ_1^B *sets,*

$$\bigcup_{n=0}^{\infty} \mathcal{S}_n \in \Delta_1^B \Rightarrow P\left(\bigcup_{n=0}^{\infty} \mathcal{S}_n\right) = \sum_{n=0}^{\infty} P(\mathcal{S}_n).$$

Let P be a countably additive probability measure on Δ_1^B. Let $\mathcal{S} \in Bo$. Define:

$$\Gamma(\mathcal{S}) = \left\{\xi \in (\Delta_1^B)^\omega : \mathcal{S} \subseteq \bigcup_{i \in \omega} \xi_i\right\}.$$

$$P^*(\mathcal{S}) = \inf_{\xi \in \Gamma(\mathcal{S})} \sum_{i \in \omega} P(\xi_i).$$

P^* is called the *outer measure* generated by P. It can be shown that[35]

If P is a countably additive probability measure on Δ_1^B, then P^* restricted to Bo is the unique,
countably additive probability measure on Bo such that $P \subseteq P^*$.

Using this result, prove proposition 13.16. (Hint: use exercise 13.1.)

[35] (Royden 1988): 295, theorem 8.

14

Experiment and Causal Inference

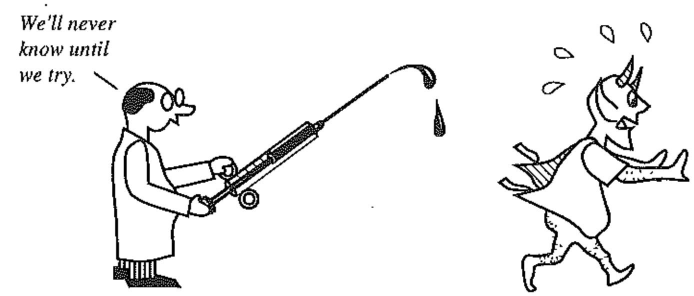

1. Introduction

Until now, I have represented the scientist as a passive observer who watches the world pass before his fixed and undirectable gaze. But in fact, an experiment is not just a question put to nature, it is an *act* that changes the state of the world under study in some way. This chapter examines active, experimental science from a logical reliabilist point of view. Since the issues that arise are more complex, the development will be correspondingly more heuristic and suggestive than in the preceding chapters. What follows may be thought of more as an outline than as a fully articulated theory.

Since Aristotle's time, empirical science has aimed at discovering necessary laws and causal mechanisms rather than merely contingent generalizations. Does scientific method measure up to this aim? Scholastic realists held the view that essences may be observed directly in virtue of a special power of the mind, and that this ability accounts for the reliable discovery of necessary truths. Nominalistic scholastics like William of Ockham undercut this theory by denying that essences could function in causal accounts of perception. Methodology was then faced with explaining how *necessary* conclusions can be drawn from *contingent* evidence if the natures grounding the necessity are not given in perception. For example, even if all ravens ever observed are black, the essence or DNA of ravens may indeed admit of other colors that *happen* not to have been realized for accidental reasons.

Francis Bacon thought that experimentation could bridge the gap between contingent data and modal conclusions. The manipulation of nature was supposed by him to somehow provide access to essences rather than to mere accidents. David Hume held, to the contrary, that necessary connections