Recall that $\mathfrak{N}$ is the Baire space $(\mathcal{N}, B)$. Let $\mathfrak{N} + \mathfrak{S}$ be the result of adding all singletons to $\mathfrak{N}$. Let $\mathfrak{D} = (\mathcal{N}, 2^{\mathcal{N}})$ be the discrete topology on $\mathcal{N}$. Let $\mathfrak{T} = (\mathcal{N}, \{\varnothing, \mathcal{N}\})$ be the trivial topology on $\mathcal{N}$. Let $\mathfrak{W} = (\mathcal{N}, \Delta)$, where Δ is the set of all closed balls under the distance function $\rho(\varepsilon, \tau) = \sup_n |\varepsilon_n - \tau_n|$, where the distance is ∞ if the sup does not exist.

(i) Show that $\forall \varepsilon$, h is verifiable$_C$ in the limit given $\mathfrak{T}$ on $\varepsilon \Leftrightarrow h$ is verifiable$_C$ in the limit given $\mathcal{N}$.

(ii) Let $C_h = \{\zeta\}$. Show that h is not verifiable$_C$ with certainty given $\mathfrak{N}$, $\mathfrak{T}$, or $\mathfrak{W}$, on ζ, whereas h is decidable$_C$ with certainty given $\mathfrak{D}$ on ζ. Show that a structure $\mathfrak{K}$ satisfies "h is decidable$_C$ with certainty given $\mathfrak{K}$ on ζ" just in case $\exists n$ such that $\zeta \in (\overline{\{\zeta\}} \mapsto \{\zeta\}$ is refuted by stage $n)_{\mathfrak{K}}$. Do you believe that if a universal hypothesis is actually true then there is some fixed time such that if the hypothesis were false it would be refuted by that time? What does this say about the prospects for Nozick's program (cf. section 2.7) as a response to the problem of local underdetermination? Show that for every $\mathfrak{K}$, $\zeta \in (\overline{\{\zeta\}} \mapsto \exists n$ such that $\{\zeta\}$ is refuted by stage $n)_{\mathfrak{K}}$. Why does the placement of the existential quantifier matter so much? (The former placement is said to be *de re*, or *of things*, while the latter placement is said to be *de dicto*, or *of words*.)

(iii) Show that for each ε, h_{fin} is not refutable in the limit given $\mathfrak{N}$, $\mathfrak{T}$, or $\mathfrak{W}$ on ε.

(iv) Show that:

(a) $\mathcal{DC}_C[h]_{\mathfrak{N}+\mathfrak{S}} = \mathcal{DC}_C[h]_{\mathfrak{N}} = C_h - bdry(C_h)$.
(b) $\mathcal{VC}_C[h]_{\mathfrak{N}+\mathfrak{S}} - \mathcal{DC}_C[h]_{\mathfrak{N}+\mathfrak{S}} = bdry(C_h) - C_h$.
(c) $\mathcal{RC}_C[h]_{\mathfrak{N}+\mathfrak{S}} - \mathcal{DC}_C[h]_{\mathfrak{N}+\mathfrak{S}} = bdry(C_h) - \overline{C_h}$.
(d) $int(\overline{C_h}) = \varnothing \Rightarrow \mathcal{DC}_C[h]_{\mathfrak{N}+\mathfrak{S}} = C_h$ (i.e., "*h is decidable with certainty*" is the *same proposition* as *h*).

(v) What happens to (b) and (c) in $\mathfrak{N}$?

(vi) Relate results (iv a–c) to the discussion in this chapter in which it was claimed that the problem of induction arises in the boundaries of hypotheses.

(vii) What is the significance of (d) for naturalistic epistemology, the proposal that we should use empirical inquiry to find out if we are or can be reliable?

5

Reducibility and the Game of Science

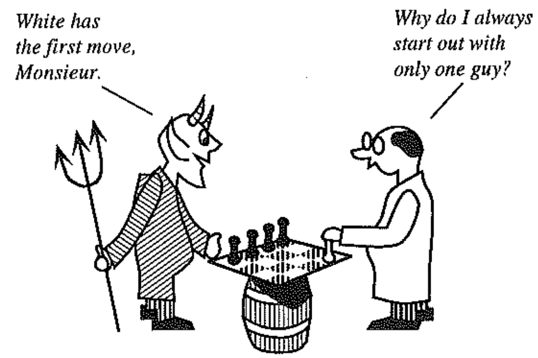

1. Introduction

Many of the negative arguments considered so far have involved demons who attempt to feed misleading evidence to the inductive method α, as though science were an ongoing game between the demon and the scientist, where winning is determined by what happens in the limit. This game-theoretic construal of inquiry is familiar in skeptical arguments from ancient times, and it is the purpose of this chapter to explore it more explicitly than we have done so far. One reason it is interesting to do so is that the theory of infinite games is central to contemporary work in the foundations of mathematics, so that the existence of winning strategies for inductive demons hinges on questions unanswered by set theory.[1] In set theory, infinite games are known to be intimately related to the notion of continuous reducibility among problems and inductive methods turn out to determine continuous reductions. It is therefore natural to consider games and continuous reducibility all at once. The discussion of continuous reducibility in this chapter will help to underscore the analogy between reliable inductive inference on the one hand and ordinary computation on the other, which will be developed in the next chapter.

[1] Material presented in this chapter assumes some elementary issues in set theory, but may be skipped without loss of comprehension in the chapters that follow.

2. Ideal Inductive Methods as Continuous Operators on the Baire Space

Consider a hypothesis assessment method α. α takes a hypothesis and a finite data sequence as arguments and returns a rational number in the interval [0, 1]. Let ξ be a fixed 1–1 map from the rational numbers in [0, 1] to the natural numbers. This way, we may think of α as conjecturing natural code numbers for rationals. Then for fixed h, we may think of α as determining a function $\Phi_{\alpha,h}: \mathcal{N} \to \mathcal{N}$ from data streams to conjecture streams such that the code number of the nth conjecture of α on h is the nth entry in the infinite conjecture stream that is the value of $\Phi_{\alpha,h}(\varepsilon)$ (*Fig. 5.1*). In other words:

$$\Phi_{\alpha,h}(\varepsilon)_n = \xi(\alpha(h, \varepsilon|n)).$$

A function Φ from $\mathcal{N}$ to $\mathcal{N}$ is called an *operator*. It turns out that the operator $\Phi_{\alpha,h}$ is continuous with respect to the Baire space. Recall that Φ is continuous just in case $\Phi^{-1}(\mathcal{S})$ is open if $\mathcal{S}$ is. In the Baire space, this amounts to the requirement that successive conjectures in $\Phi(\varepsilon)$ depend only on finite initial segments of ε, so continuity captures the bounded perspective of the scientist. The following proposition makes the connection precise (*Fig. 5.2*).

Proposition 5.1

$\Phi: \mathcal{N} \to \mathcal{N}$ is continuous $\Leftrightarrow \forall\varepsilon, \forall n, \exists e \subseteq \varepsilon$ such that $\forall\varepsilon'$ extending e, $\Phi(\varepsilon')_n = \Phi(\varepsilon)_n$.

Proof: ($\Leftarrow$) Assume the right-hand side of the fact. Let O be an open subset of $\mathcal{N}$. Then for some $G \subseteq \omega^*$, O is the union of all $[e]$ such that $e \in G$. Let $\mathcal{R}(k, n)$ be the union of all $[e']$ such that for each ε extending e', $\Phi(\varepsilon)_n = k$. $\mathcal{R}(k, n)$ is open since each $[e']$ is (*Fig. 5.3*).

Hence, the finite intersection $\mathcal{R}(e) = \mathcal{R}(e_0, 0) \cap \cdots \cap \mathcal{R}(e_{lh(e)-1}, lh(e))$ is open (*Fig. 5.4*).

The union $\mathcal{R}$ of the $\mathcal{R}(e)$ over all $e \in G$ is therefore open. Let $\varepsilon \in \mathcal{R}$. Then $\varepsilon \in \mathcal{R}(e)$, for some $e \in G$. Thus, $e \subseteq \Phi(\varepsilon)$, by the definition of $\mathcal{R}(e)$, so $\Phi(\varepsilon) \in O$ (*Fig. 5.5*). Thus $\mathcal{R} \subseteq \Phi^{-1}(O)$.

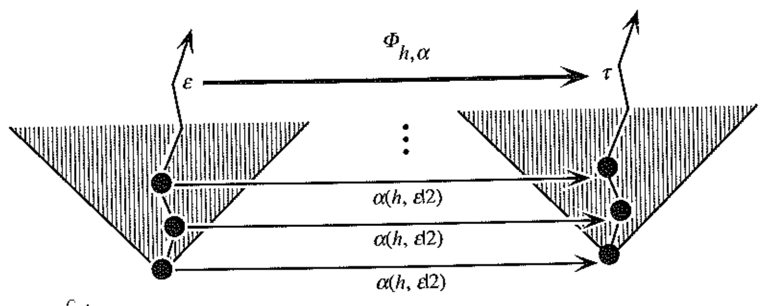

Figure 5.1

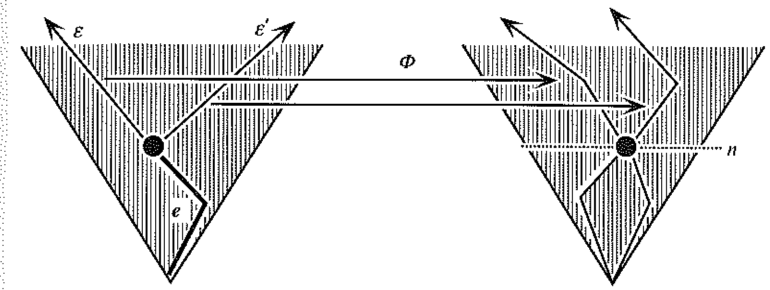

Figure 5.2

Let $\varepsilon \in \Phi^{-1}(O)$, so $\Phi(\varepsilon) \in O$. Since $\Phi(\varepsilon) \in O$, there is some $e \in G$ such that $e \subseteq \Phi(\varepsilon)$. But by hypothesis, for each $n \le lh(\varepsilon)$ there is an $e' \subseteq \varepsilon$ such that for each ε' extending e' $\Phi(\varepsilon')_n = \Phi(\varepsilon)_n$. Thus, $\varepsilon \in \mathcal{R}(e') \subseteq \mathcal{R}$. So $\Phi^{-1}(O) = \mathcal{R}$, and hence $\Phi^{-1}(O)$ is open.

($\Rightarrow$) Suppose that Φ is continuous. Then for each open set O, $\Phi^{-1}(O)$ is open. Let $\varepsilon \in \mathcal{N}$, $n \in \omega$ be given. Suppose $\Phi(\varepsilon)_n = k$. Let $\mathcal{S}(k, n)$ be the set of all infinite sequences ε' such that $\varepsilon'_n = k$. $\mathcal{S}(k, n)$ is just the union of all fans $[e]$ such that $lh(e) \ge n$ and $e_n = k$, and hence is open. By the continuity of Φ, $\Phi^{-1}(\mathcal{S}(k, n))$ is open. Hence, there is some $G \subseteq \omega^*$ such that $\Phi^{-1}(\mathcal{S}(k, n))$ is the union of all $[e'']$ such that $e'' \in G$ (*Fig. 5.6*). Since $\Phi(\varepsilon)_n = k$, some initial segment e^* of ε is in G. But then for each extension ε'' of e^*, $\varepsilon'' \in \Phi^{-1}(\mathcal{S}(k, n))$, so $\Phi(\varepsilon'') \in \mathcal{S}(k, n)$ and hence $\Phi(\varepsilon'')_n = k$. ■

Since $\Phi_{\alpha,h}$ clearly satisfies the right-hand side of proposition 5.1, the continuity of $\Phi_{\alpha,h}$ follows immediately. It is not the case, however, that each continuous Φ is identical to some $\Phi_{\alpha,h}$. Continuity leaves open the possibility of waiting for relevant data to arrive. For example, define:

$$\Phi(\varepsilon) = \begin{cases} \tau & \text{if } \varepsilon_{100} = 0 \\ \tau' & \text{otherwise.} \end{cases}$$

τ and τ' are any conjecture streams that differ up to time 100. Φ is continuous, but no α that is forced to produce a conjecture on each finite initial segment of the data can induce Φ. Nevertheless, it turns out that this extra generality

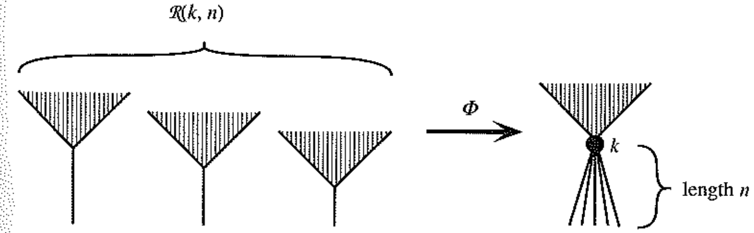

Figure 5.3

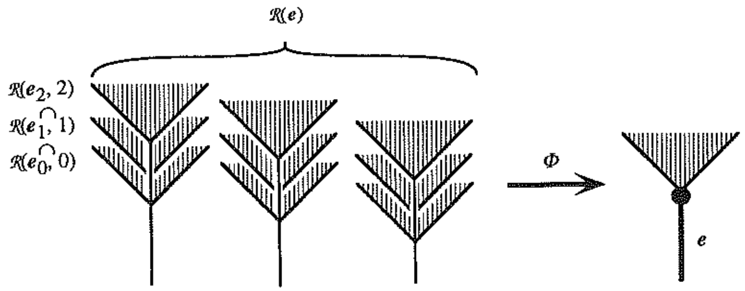

Figure 5.4

doesn't matter for our applications, since we are interested in methods that converge, and convergence is undisturbed if a method stalls by repeating its previous conjecture until some relevant datum arrives in the data stream, as we shall see in the next section.

3. Assessment as Reduction

Many results in topology, descriptive set theory, logic, and the theories of computability and computational complexity are about *reducibility*. The following approach is common to each of these fields. Let $\boldsymbol{G}$ be a collection of operators on $\mathcal{N}$. We may think of $\boldsymbol{G}$ as characterizing all operators corresponding to agents of a given cognitive power. For example, arbitrary operators correspond to the abilities of the Judeo-Christian deities who can see the entire future at a timeless glance. Continuous operators correspond to the abilities of "ideal agents" who cannot see the future but who can intuit all mathematical relations in an instant. Computable operators (cf. chapter 7) correspond to the abilities of a scientist who follows an algorithmic inductive method (or who lets a digital computer do his work for him). Finite state operators (cf. chapter 8) correspond to the abilities of a computer with a fixed memory store. Let $\mathcal{A}, \mathcal{B} \subseteq \mathcal{N}$. Then define (*Fig. 5.7*):

$$\mathcal{A} \leq_G \mathcal{B} \Leftrightarrow \textit{there is a } \Phi \in \boldsymbol{G} \textit{ such that for each } \varepsilon \in \mathcal{N},\ \varepsilon \in \mathcal{A} \Leftrightarrow \Phi(\varepsilon) \in \mathcal{B}.$$

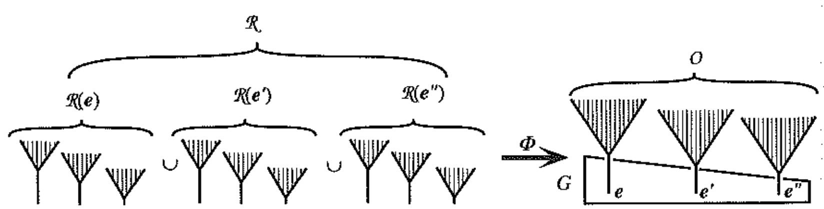

Figure 5.5

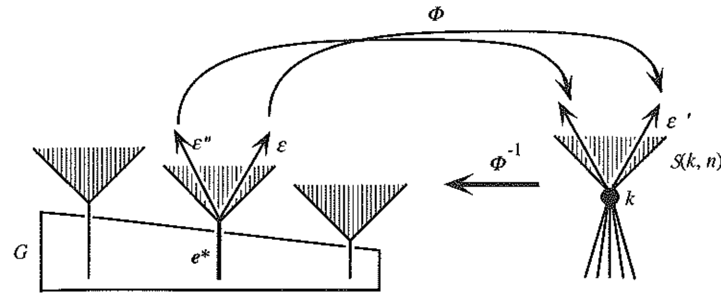

Figure 5.6

When $\mathcal{A} \leq_G \mathcal{B}$, we may say that $\mathcal{A}$ is ***G*-reducible** to $\mathcal{B}$. ***G*-equivalence** is defined in the obvious manner:

$$\mathcal{A} \equiv_G \mathcal{B} \Leftrightarrow \mathcal{A} \leq_G \mathcal{B} \textit{ and } \mathcal{B} \leq_G \mathcal{A}.$$

We may think of $\mathcal{A} \leq_G \mathcal{B}$ as saying that $\mathcal{A}$ is no harder than $\mathcal{B}$, so far as operators in $\boldsymbol{G}$ are concerned. Let ***Cnt*** denote the continuous operators on $\mathcal{N}$. Then the relation $\leq_{Cnt}$ of continuous reducibility satisfies the following properties.[2]

Proposition 5.2

(a) $\leq_{Cnt}$ *is reflexive and transitive over the subsets of* $\mathcal{N}$, so $\equiv_{Cnt}$ *is an equivalence relation over these sets.*

(b) *If* $\mathcal{A} \leq_{Cnt} \mathcal{B}$ *then* $\bar{\mathcal{A}} \leq_{Cnt} \bar{\mathcal{B}}$.

(c) *If* $\mathcal{A} \in \begin{bmatrix} \Sigma_n^B \\ \Pi_n^B \\ \Delta_n^B \end{bmatrix}$ *and* $\mathcal{B} \leq_{Cnt} \mathcal{A}$ *then* $\mathcal{B} \in \begin{bmatrix} \Sigma_n^B \\ \Pi_n^B \\ \Delta_n^B \end{bmatrix}$.

(d) *We may substitute D for B in* (c).

Proof: Exercise 5.1. ■

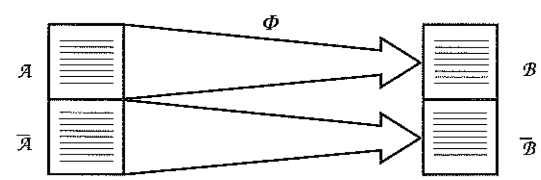

Figure 5.7

[2] Continuous reducibility is called *Wadge reducibility* in the logical literature and is denoted by $\leq_w$ (Moschavakis 1980).

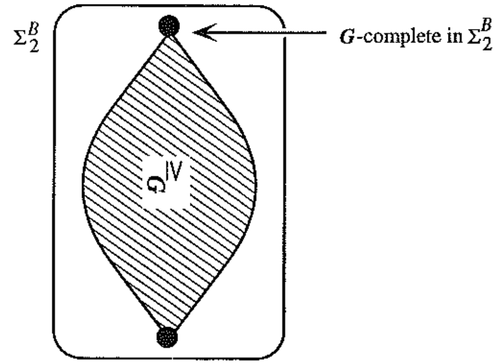

Figure 5.8

Let Θ be a complexity class of subsets of $\mathcal{N}$ (e.g., Σ_n^B). Then a ***G*-complete** set in Θ is a maximally complex member of Θ (*Fig. 5.8*).

$$\mathcal{A} \text{ is } \textbf{\textit{G}-complete in } \Theta \Leftrightarrow \text{(i) } \mathcal{A} \in \Theta \text{ and (ii) for each } \mathcal{B} \in \Theta, \mathcal{B} \leq_G \mathcal{A}.$$

It turns out that if we fix $\mathcal{K} = \mathcal{N}$, then reliability is equivalent to ***Cnt*-reducibility** to one of the following sets, which correspond to the various notions of convergence. Recall that ξ is a fixed, 1–1 map from the rationals in [0, 1] to the natural numbers.

$$\mathcal{S}_1 = \{\tau \in \mathcal{N}: \text{there is an } n \text{ such that } \tau_n = 1\}.$$

$$\mathcal{S}_2 = \{\tau \in \mathcal{N}: \text{there is an } n \text{ such that for all } m \geq n, \tau_m = 1\}.$$

$$\mathcal{S}_2(k) = \{\tau \in \mathcal{N}: \tau \in \mathcal{S}_2 \ \& \text{ there are at most } k \text{ positions } i \text{ such that } \tau_i \neq \tau_{i+1}\}.$$

$$\mathcal{S}_3 = \{\tau \in \mathcal{N}: \text{for each rational } s > 0 \text{ there is an } n \text{ such that for all } m \geq n, 1 - \xi^{-1}(\tau_m) \leq s\}.$$

The correspondence is made precise in the following proposition.

Proposition 5.3

(a) H is verifiable$_C \begin{bmatrix} \text{with certainty} \\ \text{in the limit} \\ \text{gradually} \end{bmatrix} \Leftrightarrow$ for each $h \in H$, $C_h \leq_{Cnt} \begin{bmatrix} \mathcal{S}_1 \\ \mathcal{S}_2 \\ \mathcal{S}_3 \end{bmatrix}$.

(b) H is refutable$_C \begin{bmatrix} \text{with certainty} \\ \text{in the limit} \\ \text{gradually} \end{bmatrix} \Leftrightarrow$ for each $h \in H$, $\overline{C_h} \leq_{Cnt} \begin{bmatrix} \mathcal{S}_1 \\ \mathcal{S}_2 \\ \mathcal{S}_3 \end{bmatrix}$.

(c) H is decidable$_C \begin{bmatrix} \text{with certainty} \\ \text{in the limit} \end{bmatrix}$

$\Leftrightarrow$ for each $h \in H$, $\overline{C_h}$, $C_h \leq_{Cnt} \begin{bmatrix} \mathcal{S}_1 \\ \mathcal{S}_2 \end{bmatrix}$.

(d) If $0 < r < 1$, then H is decidable$_C$ with n mind changes starting with

$$\begin{bmatrix} 0 \\ 1 \\ r \end{bmatrix} \Leftrightarrow \text{for each } h \in H, \begin{bmatrix} C_h \\ \overline{C_h} \\ C_h, \overline{C_h} \end{bmatrix} \leq_{Cnt} \mathcal{S}_2(n).$$

Proof: I present only the certainty case of (a), the other cases being similar. ($\Rightarrow$) Let α decide$_C$ h with 1 mind change starting with 0. $\Phi_{\alpha,h}$ is the required reduction. ($\Leftarrow$) Suppose that for each $h \in H$, $C_h \leq_{Cnt} \mathcal{S}_1$. Let continuous Φ_h be such that for each $\varepsilon \in \mathcal{N}$, $\varepsilon \in C_h \Leftrightarrow \Phi_h(\varepsilon) \in \mathcal{S}_1$. Define the method α as follows. (Remember that conjectures following the declaration of certainty are irrelevant.)

$$\alpha(h, \mathbf{e}) = \begin{cases} 1 & \text{if } \mathbf{e} \neq 0 \text{ and } \alpha(h, \mathbf{e}-) = \text{'!'}, \text{ else} \\ \text{'!'} & \text{if there is an } n \text{ such that for each } \varepsilon \in [\mathbf{e}], \Phi_h(\varepsilon)_n = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Let $\varepsilon \in \mathcal{N}$. Suppose $\varepsilon \in C_h$. Then $\Phi_h(\varepsilon) \in \mathcal{S}_1$. So for some n, $\Phi_h(\varepsilon)_n = 1$. So by proposition 5.1, there is some finite $\mathbf{e} \subseteq \varepsilon$ such that for each ε' extending $\mathbf{e}$, $\Phi_h(\varepsilon')_n = 1$. Hence, $\alpha(h, \mathbf{e})$ produces a certainty mark '!' followed by 1. Suppose $\varepsilon \notin C_h$. Then by a similar argument, $\alpha(h, \mathbf{e})$ never produces the certainty mark. ■

Proposition 5.3 shows that the ability of a continuous map to wait is of no consequence when the question concerns reduction to a convergence set, since a method α that is always forced to produce a conjecture at every stage can simply repeat its previous conjecture until the information the continuous map is waiting for appears.

4. Ideal Transcendental Deductions as *Cnt*-Completeness Theorems

We have just seen that a reliable method is a continuous reduction to some convergence set. It also turns out that when $\mathcal{K} = \mathcal{N}$ the characterization theorems of the preceding chapter may be thought of as proofs that the various notions of convergence are complete for their respective Borel complexity classes.

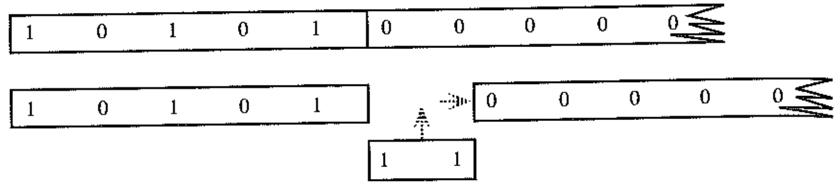

Figure 5.9

Proposition 5.4

$$\begin{bmatrix} \mathcal{S}_1 \\ \mathcal{S}_2(n) \\ \mathcal{S}_2 \\ \mathcal{S}_3 \end{bmatrix} \textit{ is } \textbf{\textit{Cnt}}\textit{-complete in} \begin{bmatrix} \Sigma_1^B \\ \Sigma_n^D \\ \Sigma_2^B \\ \Pi_3^B \end{bmatrix}.$$

Proof of Σ_2^B *case*: Let C_h be an arbitrary Σ_2^B set. Then by proposition 4.10, some α verifies$_C$ h in the limit given $\mathcal{K}$. So $\Phi_{\alpha,h}$ is a continuous reduction of C_h to $\mathcal{S}_2$. So for each $\mathcal{P} \in \Sigma_2^B$, $\mathcal{P} \leq_{Cnt} \mathcal{S}_2$. Moreover, it is immediate by definition that $\mathcal{S}_2 \in \Sigma_2^B$. So $\mathcal{S}_2$ is ***Cnt***-complete in Σ_2^B. The other cases are similar. ■

The idea behind proposition 5.4 is more general than the particular notions of convergence considered. Say that $\mathcal{S}$ is *closed under finite repetition* just in case for each $\varepsilon \in \mathcal{S}$, the result of inserting an arbitrary, finite number of xs after each occurrence of x in ε is still in $\mathcal{S}$ (*Fig. 5.9*).

It is easy to see that $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{S}_2(n)$, and $\mathcal{S}_3$ are all closed under finite repetition. For each convergence set closed under finite repetition, the existence of a reliable method is equivalent to the existence of a continuous reduction, as in proposition 5.3. To extend inquiry to Σ_n^B, all we have to do is to invent a new notion of convergence that yields a convergence set closed under finite repetition that is ***Cnt***-complete in Σ_n^B (cf. exercise 5.4).

5. Inductive Demons as Continuous Counterreductions

In each of the demonic arguments employed in the preceding chapters, the demon looks at the conjectures produced by the scientist so far and generates the next datum to be fed to the scientist. So we may view the demon explicitly as a map δ from finite sequences of conjectures to data points. In this sense, the demon is a mirror image of the scientist, who maps finite sequences of data points to conjectures (*Fig. 5.10*).

The demon's goal is also a mirror image of the scientist's goal. The scientist's method is supposed to reduce C_h to $\mathcal{S}$, where $\mathcal{S}$ is the set of all sequences that converge in the relevant sense. The demon's job is to produce a data stream in

Figure 5.10

C_h if and only if the given conjecture stream is not in $\mathcal{S}$. Just like a scientific method, a demonic strategy δ induces an operator Φ_δ on $\mathcal{N}$. A successful demon therefore witnesses the relation $\bar{\mathcal{S}} \leq_{Cnt} C_h$, just as a successful scientific method witnesses the relation $C_h \leq_{Cnt} \mathcal{S}$. How does $\bar{\mathcal{S}} \leq_{Cnt} C_h$ conflict with $C_h \leq_{Cnt} \mathcal{S}$? The answer is that demonic arguments, like characterization theorems, may be viewed as completeness proofs (*Fig. 5.11*).

Proposition 5.5

$$\textit{If } C_h \in \begin{bmatrix} \Sigma_1^B \\ \Sigma_n^D \\ \Sigma_2^B \\ \Pi_3^B \end{bmatrix} \textit{ and } \begin{bmatrix} \overline{\mathcal{S}_1} \\ \overline{\mathcal{S}_2(n)} \\ \overline{\mathcal{S}_2} \\ \overline{\mathcal{S}_3} \end{bmatrix} \leq_{Cnt} C_h,$$

$$\textit{then } C_h \textit{ is } \textbf{\textit{Cnt}}\textit{-complete in} \begin{bmatrix} \Pi_1^B \\ \Pi_n^D \\ \Pi_2^B \\ \Sigma_3^B \end{bmatrix}.$$

Proof: Immediate consequence of proposition 5.4. ■

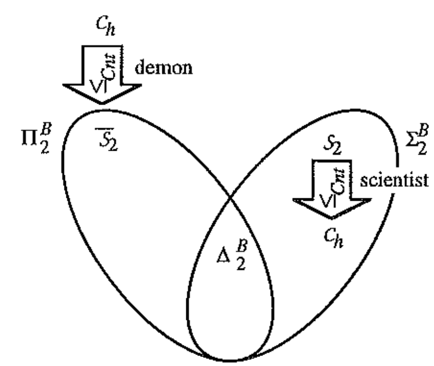

Figure 5.11

For example, $C_{h_{inf}}$ is Π_2^B complete by proposition 5.5 and the demonic argument given in chapter 3. By the Borel hierarchy theorem (proposition 4.9), there is a non-Σ_2^B set $\mathcal{R} \in \Pi_2^B$. So $C_{h_{inf}}$, which is Π_2^B-complete, reduces $\mathcal{R}$. By the contrapositive of proposition 5.2(c), any set that continuously reduces a non-Σ_2^B set is a non-Σ_2^B set, so $C_{h_{inf}}$ is not Σ_2^B. By the characterization theorem for limiting verifiability (proposition 4.10), h is not verifiable$_C$ in the limit. Similar arguments apply in the other cases. So a demonic argument shows not only that C_h is not in the Borel complexity class characterizing inductive success of the required sort, but also that C_h is at least as impossible to solve for ideal agents as the worst problem in the dual complexity class.

The parallelism between the demon and the scientist can be extended to background knowledge. $\mathcal{K}$ tells the scientist what sort of data stream the demon may produce in the limit. Suppose we were to restrict the scientist to producing conjecture streams in some set $\mathcal{D}$, so that $\mathcal{D}$ is the demon's background knowledge about what the scientist will do in the limit. The demon's task would become easier, just the way the scientist's task becomes easier when $\mathcal{K}$ is strengthened. For example, consider what happens when the hypothesis investigated by the scientist is h_{fin}, but the demon knows that the scientist will never say 1 again once he stops conjecturing 1.

6. Science as a Limiting Game

We have seen that a continuous demon suffices to show that an inductive problem is unsolvable and that when the problem is solvable, it is as though the scientist has turned tables on the demon, exploiting the latter's restricted view of the scientist's conjectures. But this raises a question, namely, whether there are inductive problems for which neither the scientist nor the demon has a winning strategy. In other words, is the game of inquiry always *determined*?

The question belongs to the theory of infinite games of perfect information, which has found important applications in logic and set theory.[3] We assume that there are two players, which I will provide with the suggestive names S and D, for the scientist and the demon, respectively. A *strategy* is just a function $\phi: \omega^* \to \omega$ from finite sequences of natural numbers to natural numbers. Let α, δ be strategies for S and for D, respectively. The *play sequence* $P_{\alpha,\delta}$ of α and δ is built up as follows. The strategy α of S is applied to the empty sequence, yielding the first play $\alpha(\mathbf{0}) = s_0$. Then the strategy δ of D is applied to the sequence (s_0), yielding the first play by D, namely $\delta((s_0)) = d_0$. In general, the nth play of S is the value of α on all previous plays by δ, and the nth play of D is the value of δ on all previous plays by α. In the limit, the interaction of α and δ generates the infinite play sequence:

$$P_{\alpha,\delta} = (s_0, d_0, s_1, d_1, s_2, d_2, \ldots)$$

[3] For a review, see Moschavakis (1980).

where s_i is the ith play of the scientist α and d_i is the ith play of the demon δ. If $\varepsilon \in \mathcal{N}$, then define

$$P_{\alpha,\varepsilon} = (\alpha(\varepsilon|0), \alpha(\varepsilon|1), \alpha(\varepsilon|2), \ldots) \quad \text{and} \quad P_{\varepsilon,\delta} = (\delta(\varepsilon|0), \delta(\varepsilon|1), \delta(\varepsilon|2), \ldots).$$

In other words, a sequence of numbers may be viewed as a fixed strategy that ignores past plays by the opponent. The subsequence

$$D_{\alpha,\delta} = (d_0, d_1, d_2, \ldots)$$

is the infinite sequence of plays of δ in response to α, and the infinite play sequence

$$S_{\alpha,\delta} = (s_0, s_1, s_2, \ldots)$$

is the infinite sequence of conjectures produced by α in response to δ.

An *infinite game of perfect information* is now just a set $\mathcal{W} \subseteq \mathcal{N}$. Winning is defined with respect to $\mathcal{W}$ in the following way:

α wins for S against δ in $\mathcal{W} \Leftrightarrow P_{\alpha,\delta} \in \mathcal{W}$

δ wins for D against α in $\mathcal{W} \Leftrightarrow \alpha$ does not win for S against δ in $\mathcal{W}$.

Once the players select their respective strategies, one and only one of them is destined to win the game $\mathcal{W}$, by the definition of winning. But we are not interested in mere winning. We are interested in strategies logically guaranteed to win, no matter which strategy the opponent selects.

α is a winning strategy for S in $\mathcal{W}$
$\Leftrightarrow \forall$ strategy δ, α wins for S against δ in $\mathcal{W}$.

δ is a winning strategy for D in $\mathcal{W}$
$\Leftrightarrow \forall$ strategy α, δ wins for D against α in $\mathcal{W}$.

S has a winning strategy in $\mathcal{W}$
$\Leftrightarrow \exists\alpha$ such that α is a winning strategy for S in $\mathcal{W}$.

D has a winning strategy in $\mathcal{W}$
$\Leftrightarrow \exists\alpha$ such that δ is a winning strategy for D in $\mathcal{W}$.

When either S or D has a winning strategy in $\mathcal{W}$, we say that $\mathcal{W}$ is *determined*. It is not obvious from the definition of determinacy that all games $\mathcal{W}$ are determined. In fact, the answer to this question hinges on debates concerning the foundations of mathematics so that the foundations of inductive methodology and the foundations of mathematics are intertwined.

In standard, Zermelo-Fraenkel set theory[4] with the axiom of choice added (ZFC), there is a powerful positive result for determinacy.

Proposition 5.6 (Martin 1975)

If $\mathcal{W}$ is a Borel set then $\mathcal{W}$ is determined. ■

As a corollary to this result, whenever C_h is a Borel set, underdetermination (in any of the senses introduced in chapter 3) implies a winning strategy for the demon, because the definition of the winning set in such a game adds just a bit of complexity to the Borel complexity of C_h.

Since conjectures less than 1 can be converted to 0 conjectures, we may assume without loss of generality that scientific strategies for limiting verification produce only 0–1 conjectures. This permits us to dispense with the fussy encoding of rationals into natural numbers. Let ρ be an infinite play sequence and let ρ_D and ρ_S be the demon's and the scientist's plays in ρ, respectively. Since α produces only 0–1 conjectures, $\rho_S \in \mathcal{N}$. Then the game of limiting verification can be defined as follows:

$$Limver_C(h) = \{\rho\colon \rho_D \in C_h \Leftrightarrow \rho_S \in \mathcal{S}_2\}.$$

It is clear that $Limver_C(h)$ is Borel if C_h is, and the same is true for the other criteria of success.

Outside of the Borel hierarchy, however, ZFC yields spotty results. Above the Borel hierarchy is the *projective hierarchy*, and the sets in this hierarchy are known as the *projective sets*. The precise definition of the projective hierarchy is not essential for our purposes here. Suffice it to say that they are formed by existential quantification over functions on the natural numbers rather than by existential quantification over natural numbers (i.e., countable union).

Proposition 5.7[5]

There are projective sets whose determinacy is independent of ZFC. ■

Projective determinacy is the hypothesis that all games based on projective sets are determined. Projective determinacy has been shown to be independent of ZFC, and hence has been entertained as an additional axiom of set theory. Beyond the projective sets, ZFC entails the existence of nondetermined games:

Proposition 5.8 (Gale and Stewart 1953)

ZFC implies the existence of a nondetermined game. ■

[4] Cf. Kunen (1980) for an introduction to ZFC.
[5] Moschovakis (1980): 297.

The Gale-Stewart theorem does not settle our question in the negative, however, since scientific games have a special form and the game produced by the Gale-Stewart construction is not guaranteed to have this form. If we grant the continuum hypothesis, then a nondetermined limiting verification game exists. The *continuum hypothesis* (CH) says that there are no uncountable cardinals between $|\omega|$ and $|2^\omega| = |\mathcal{N}|$. The continuum hypothesis is known to be independent of ZFC (i.e., neither it nor its negation is entailed by ZFC) *if* ZFC is itself consistent.[6] I leave open the question whether a nondetermined limiting verification game exists when the continuum hypothesis is dropped.

Proposition 5.9 (Juhl 1991)

ZFC + CH implies that there exists a C_h such that $Limver_C(h)$ is not determined.

Proof: Using the axiom of choice, well-order the set of all scientific strategies as $\{\alpha_\mu\colon \mu < \omega_1\}$ and the set of all demonic strategies as $\{\delta_\mu\colon \mu < \omega_1\}$, where ω_1 is the first uncountable ordinal. By the continuum hypothesis, each $\mu < \omega_1$ is countable. Now we define an ordinal sequence $\mathcal{A}_\mu$, $\mathcal{B}_\mu$ of sets of data streams such that at each stage $\mu < \omega_1$, $\mathcal{A}_\mu \cap \mathcal{B}_\mu = \varnothing$. We will let $C_h = \bigcup_{\mu<\omega_1} \mathcal{A}_\mu$. The idea is that at each stage $\mu < \omega_1$, we add a data stream to $\mathcal{A}_\mu$ or to $\mathcal{B}_\mu$ that ensures that neither the scientific strategy α_μ nor the demonic strategy δ_μ wins $Limver_C(h)$. The idea is complicated by the fact that we may not be able to find the right sort of data stream to add to ensure that δ_μ loses, but it will turn out in such cases that δ_μ is not a winning strategy anyway.

Stage 0: $\mathcal{A}_0 = \mathcal{B}_0 = \varnothing$.

Stage μ: Let $\mathcal{A} = \bigcup_{\mu'<\mu} \mathcal{A}_{\mu'}$ and $\mathcal{B} = \bigcup_{\mu'<\mu} \mathcal{B}_{\mu'}$.

To ensure that α_μ is not a winning strategy: Find the first δ such that $D_{\alpha_\mu,\delta} \notin \mathcal{A} \cup \mathcal{B}$. Such a δ exists because $\forall \mu' \leq \mu$, μ' is countable and $\mathcal{A} \cup \mathcal{B}$ is therefore countable since at most two items are added at each stage μ' (as will be seen in the construction that follows). If $S_{\alpha_\mu,\delta}$ stabilizes to 1, then set $\mathcal{A}'_\mu = \mathcal{A}$ and set $\mathcal{B}'_\mu = \mathcal{B} \cup \{D_{\alpha_\mu,\delta}\}$. Otherwise, set $\mathcal{A}'_\mu = \mathcal{A} \cup \{D_{\alpha\mu,\delta}\}$ and set $\mathcal{B}'_\mu = \mathcal{B}$. In other words, we put $D_{\alpha_\mu,\delta}$ wherever it would make α_μ fail.

To ensure that δ_μ is not a winning strategy: Find the first α such that $D_{\alpha,\delta_\mu} \notin \mathcal{A}'_\mu \cup \mathcal{B}'_\mu$, if there is one. If $S_{\alpha,\delta\mu}$ stabilizes to 1, then set $\mathcal{A}_\mu = \mathcal{A}'_\mu \cup \{D_{\alpha,\delta_\mu}\}$ and set $\mathcal{B}_\mu = \mathcal{B}'_\mu$. Otherwise, set $\mathcal{A}_\mu = \mathcal{A}'_\mu$ and set $\mathcal{B}_\mu = \mathcal{B}'_\mu \cup \{D_{\alpha_\mu,\delta}\}$. In other words, we put D_{α,δ_μ} wherever it would make δ_μ fail.

Now let $C_h = \bigcup_{\mu<\omega_1} \mathcal{A}_\mu$. It is clear that α_μ fails to be a winning strategy for

[6] Cf. Kunen (1980): 209.

the scientist because it loses against the δ chosen at stage μ. It is also clear that δ_μ fails to be a winning strategy for the demon whenever the α sought at stage μ exists. It remains only to show that δ_μ fails to be a winning strategy for the demon even when the α sought at stage μ does *not* exist. In that case,

$$(*) \quad \forall \alpha, D_{\alpha, \delta_\mu} \in \mathcal{A}'_\mu \cup \mathcal{B}'_\mu.$$

(*) says that the demon δ_μ responds to all possible scientific strategies with just a countable set of data streams. Since the demon's response is a function only of what the scientist actually plays, the (countable) set of all of δ_μ's responses is

$$\mathcal{P} = \{D_{\tau, \delta_\mu}: \tau \in \mathcal{N}\}.$$

Recall that $\mathcal{S}_2$ denotes the set of all ω-sequences that stabilize to 1. For reductio, assume that

$$(**) \quad \forall \tau, \tau' \in \mathcal{N}, D_{\tau, \delta_\mu} = D_{\tau', \delta_\mu} \Rightarrow [S_{\tau, \delta_\mu} \in \mathcal{S}_2 \Leftrightarrow S_{\tau', \delta\mu} \in \mathcal{S}_2] \text{ (Fig. 5.12).}$$

(**) implies that if we look at how δ_μ responds to fixed conjecture sequences by the scientist, any two fixed conjecture sequences $\tau, \tau' \in \mathcal{N}$ that lead δ_μ to produce the same data stream either both converge to 1 or both fail to converge to 1. We may therefore define $\overline{\mathcal{S}_2}$ as the set of all fixed scientific strategies $\tau \in \mathcal{N}$ against which δ_μ produces a data stream in the countable set $\mathcal{A}'_\mu \cup \mathcal{B}'_\mu$ that is (in light of **) generated only against nonconvergent conjecture sequences. More precisely, define the set of all data streams produced by δ_μ against nonconvergent conjecture streams as follows:

$$\mathcal{R} = \{\varepsilon \in \mathcal{P}: \forall \sigma \in \mathcal{N}, D_{\sigma, \delta_\mu} = \varepsilon \Rightarrow \sigma \notin \mathcal{S}_2\}.$$

Now I claim:

$$(***) \quad \tau \notin \mathcal{S}_2 \Leftrightarrow D_{\tau, \delta_\mu} \in \mathcal{R}.$$

For suppose $D_{\tau, \delta_\mu} \in \mathcal{R}$. By the definition of $\mathcal{R}$, $D_{\tau, \delta_\mu} \in \mathcal{P}$ and $\forall \sigma \in \mathcal{N}, D_{\sigma, \delta_\mu} = D_{\tau, \delta_\mu} \Rightarrow \sigma \notin \mathcal{S}_2$. Choosing σ as τ, we obtain $\tau \notin \mathcal{S}_2$. Now suppose $D_{\tau, \delta_\mu} \notin \mathcal{R}$. Then by the definition of $\mathcal{R}$, $\exists \sigma \in \mathcal{N}$ such that $D_{\sigma, \delta_\mu} = D_{\tau, \delta_\mu}$ and $\sigma \in \mathcal{S}_2$. By (**), $\tau \in \mathcal{S}_2$.

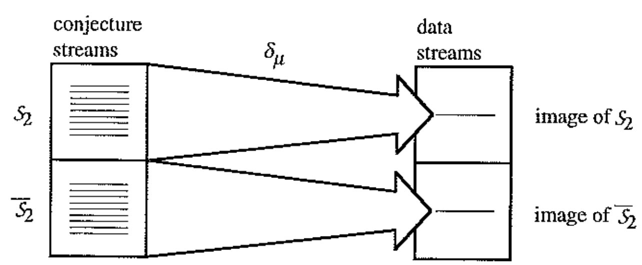

Figure 5.12

By (***), we have:

$$\tau \notin \mathcal{S}_2 \Leftrightarrow \exists \varepsilon \in \mathcal{R} \text{ such that } \forall n, \varepsilon_n = D_{\tau | n, \delta_\mu}.$$

Since $\mathcal{R}$ is countable and $\{\tau: \varepsilon_n = D_{\tau | n, \delta_\mu}\}$ is clopen, $\overline{\mathcal{S}_2} \in \Sigma_2^B$, which is absurd since $\mathcal{S}_2$ is complete in Σ_2^B. So we may conclude that (**) is false. Hence:

$$\exists \tau, \tau' \in \mathcal{N} \text{ such that } D_{\tau, \delta_\mu} = D_{\tau', \delta_\mu} \text{ and } S_{\tau, \delta_\mu} \in \mathcal{S}_2 \text{ but } S_{\tau', \delta_\mu} \notin \mathcal{S}_2.$$

So whether or not $D_{\tau, \delta_\mu} \in C_h$, δ_μ loses against either the fixed strategy τ or the fixed strategy τ'. Hence, $Limver_C(h)$ is undetermined. ■

Recall that the characterization of n-mind change decidability (proposition 4.18) was proved in terms of a generalized demonic construction ($n + 1$-feathers) whereas no such construction was given for the verification in the limit case (proposition 4.10). It seems very intuitive that there should be such a construction, since the limiting demonic arguments given in particular cases have been based on pictures that look like infinite dimensional feathers (e.g., Fig. 3.16). But even though it appears to be a small step from n-feathers to ω-feathers, no such proof of proposition 4.10 should be expected, and the preceding proposition shows why. Let GD be the claim that a generalized demonic construction for proving proposition 4.10 exists (e.g., ω-feathers). Let LVD denote the claim that all limiting verification games are determined. Now suppose that ZFC entails GD. Then (a) ZFC entails LVD, since GD implies that the demon has a winning strategy whenever the scientist does not. By proposition 5.4, ZFC + CH entails ¬LVD. So by (a), we have that ZFC entails ¬CH. But by the independence of the continuum hypothesis, if ZFC is consistent, then ZFC does not entail ¬CH. Hence, ZFC is inconsistent. We have just shown (by contraposition) that:

Corollary 5.10

If ZFC is consistent, then ZFC does not entail GD. ■

Since ZFC is the mathematical framework we have tacitly assumed throughout this study, we should not expect to show that the demon has a winning strategy whenever the scientist does not. It may then be wondered why demonic arguments were found in each particular example considered in chapter 3. The answer is that each such example involved a Borel hypothesis, and such hypotheses are covered by proposition 5.6.

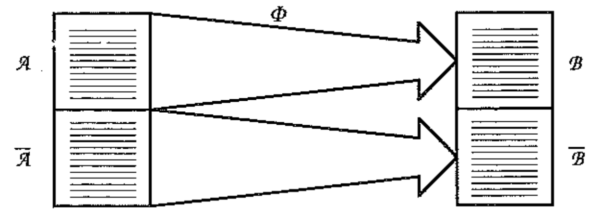

Figure 5.13

Exercises

5.1. Prove proposition 5.2.

5.2. Finish the proof of proposition 5.3.

5.3. Show that for each n, there are no ***Cnt***-incomplete sets in $\Sigma_n^D - \Pi_n^D$ or in $\Pi_n^D - \Sigma_n^D$, and for each $n \leq 3$, there are no ***Cnt***-incomplete sets in $\Sigma_n^B - \Pi_n^B$ or in $\Pi_n^B - \Sigma_n^B$. (Hint: use propositions 5.5 and 5.6.) Extend the latter result to the case in which $n = 4$. (Hint: cf. exercise 5.4.)

5.4. Invent a convergence criterion that is Σ_4^B-complete.

*5.5. This chapter has illustrated how the transcendental deductions and demonic arguments of logical reliabilism correspond to the structures of reduction and completeness, which are familiar to logicians and computability theorists. Recall, however, that the entire discussion assumes that $\mathcal{K} = \mathcal{N}$. When there is background knowledge, the tight relationships between reducibility, methods, and demons break down unless we modify the standard concept of reduction. Recall that when $\mathcal{A} \leq_G \mathcal{B}$, the picture in (*Fig. 5.13*) obtains.

When there is background knowledge, a successful scientist is free to do anything on a data stream violating that knowledge (*Fig. 5.14*):

This could be handled simply by defining reduction modulo a set, so that $\mathcal{A} \leq_G \mathcal{B}$ mod $\mathcal{K}$ if and only if for some $\Phi \in G$, for each $\varepsilon \in \mathcal{K}$, $\varepsilon \in \mathcal{A} \Leftrightarrow \Phi(\varepsilon) \in \mathcal{B}$. So far so good. But that won't help with the demon. While the scientist is free to ignore all sequences outside of $\mathcal{K}$, the demon is required to produce a sequence in $\mathcal{K}$ (*Fig. 5.15*).

What is really going on is that we are interested in reductions and counterreductions between binary partitions. The scientist is supposed to reduce the partition $(\mathcal{A} \cap \mathcal{K}, \bar{\mathcal{A}} \cap \mathcal{K})$ to the partition $(\mathcal{B}, \bar{\mathcal{B}})$, and the demon is supposed to reduce the partition $(\bar{\mathcal{B}}, \mathcal{B})$ to the partition $(\mathcal{A} \cap \mathcal{K}, \bar{\mathcal{A}} \cap \mathcal{K})$. The order of the cells matters. In general, let

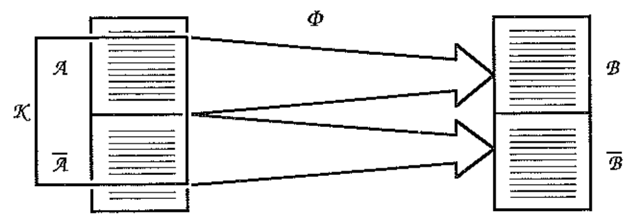

Figure 5.14

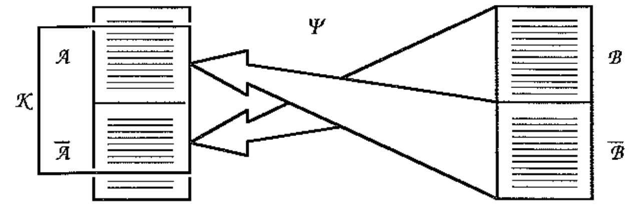

Figure 5.15

$\Pi = \mathcal{S}_1, \ldots, \mathcal{S}_k, \ldots)$ and let, $\Pi' = (\mathcal{R}_0, \ldots, \mathcal{R}_k, \ldots)$ be sequences of subsets of $\mathcal{N}$ such that for each distinct i, j, $\mathcal{S}_i \cap \mathcal{S}_j = \varnothing$ and $\mathcal{R}_i \cap \mathcal{R}_j = \varnothing$. Then define:

$\Pi \leq_G \Pi' \Leftrightarrow$ *there is a* $\Phi \in G$ *such that for each* $\varepsilon \in \mathcal{N}$, *for each* $i \in \omega$, $\varepsilon \in \mathcal{S}_i$
$\Leftrightarrow \Phi(\varepsilon) \in \mathcal{R}_i$ (*Fig.* 5.16).

This generalized notion of reduction properly handles the background assumptions $\mathcal{K}$. How do the results of this chapter work out when we add nontrivial background knowledge and move to partition reductions? Give an account of empirical verifiability and decidability in the limit given $\mathcal{K}$ in terms of partition reductions.

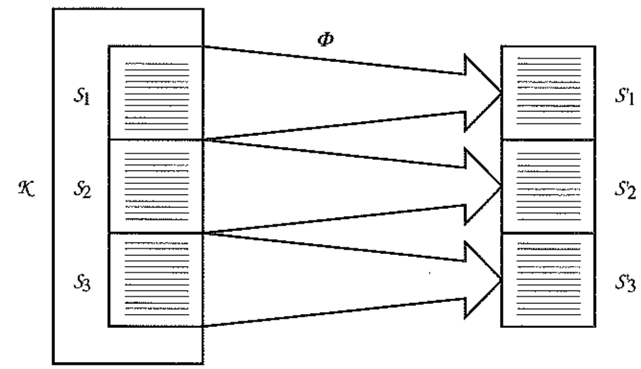

Figure 5.16