3

The Demons of Passive Observation

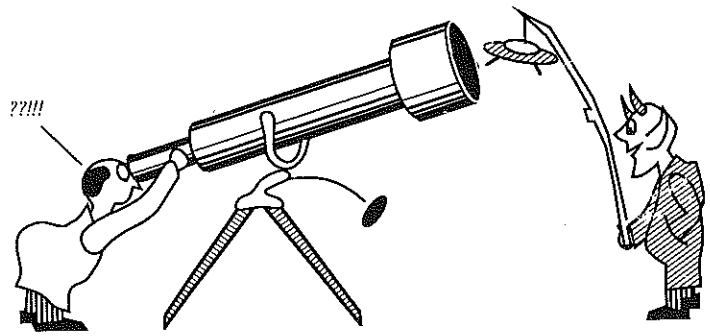

1. Introduction

In the preceding chapter, several notions of logical reliability were introduced. It was proposed that degrees of underdetermination be defined as the impossibility of reliability in various different senses. In this chapter, I will categorize some inductive problems by their degrees of underdetermination. Thus, we will be engaged in questions of the third level, which demand that inductive problems be shown unsolvable by all possible methods.

Recall that for all the scientist supposes, the actual world may be any world in $\mathcal{K}$. The hypothesis h under investigation is correct for some of the worlds in $\mathcal{K}$ and incorrect for others. The relation C of *correctness* may be truth, empirical adequacy, predictiveness, or any other relation depending only on the hypothesis, the world, and the method the scientist employs. Each world interacts with the scientist through some data protocol and emits a stream of data, of which the scientist can only scan the initial segment produced up to the current stage of inquiry. When the data protocol is historical, we have a simpler situation in which the scientist's conjectures and experimental acts feed back into the world's experimental protocol to generate more data, and so forth (*Fig. 3.1*).

Until chapter 14, I will assume that truth or correctness is an *empirical* relation, depending only on the hypothesis in question and the actual data stream, and that the scientist is a passive observer, so that the data stream does not depend on what the scientist does. Given these simplifying assumptions, the machinery of worlds and data protocols no longer does any work, since each world determines a unique data stream and truth or correctness depends only on the data stream produced. Hence, worlds can be identified with the unique data streams they produce, so that $\mathcal{K}$ is a set of possible data streams

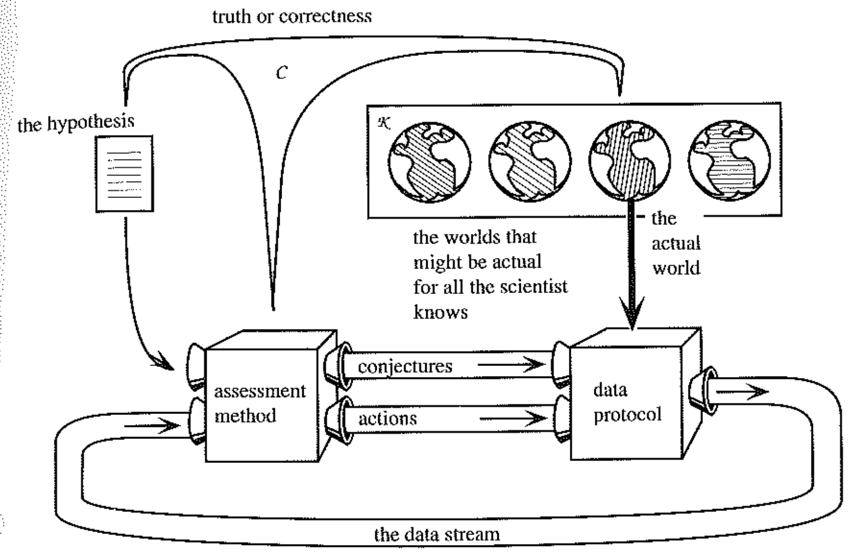

Figure 3.1

and the relation of correctness C holds directly between data streams and hypotheses. We can also ignore the fact that the scientist produces actions, since these actions are irrelevant to correctness and to the data received.

I will assume that the entries in the data stream are natural numbers. These may be thought of as code numbers for other sorts of discrete data. Let ω denote the set of all natural numbers. Let ω^* denote the set of all finite sequences of natural numbers. The set of all infinite data streams is denoted by $\mathcal{N} = \omega^\omega$.[1] Define (*Fig. 3.2*):

ε_n = *the datum occurring in position n of* ε.

$\varepsilon|n$ = *the finite, initial segment of* ε *through position* $n - 1$.

$n|\varepsilon$ = *the infinite tail of* ε *starting with position n.*

$lh(\mathbf{e})$ = *the number of positions occurring in finite sequence* $\mathbf{e}$.

$\mathbf{e}-$ = *the result of deleting the last entry in* $\mathbf{e}$.

$\text{last}(\mathbf{e})$ = *the last item occurring in* $\mathbf{e}$.

$\mathbf{0}$ = the empty sequence

Given the assumption that worlds are interchangeable with data streams, let $\mathcal{K}$ be some subset of $\mathcal{N}$ and let truth or correctness be a relation C between

[1] Each element of the infinite cross product of ω with itself is an infinite sequence of natural numbers and hence is a data stream.

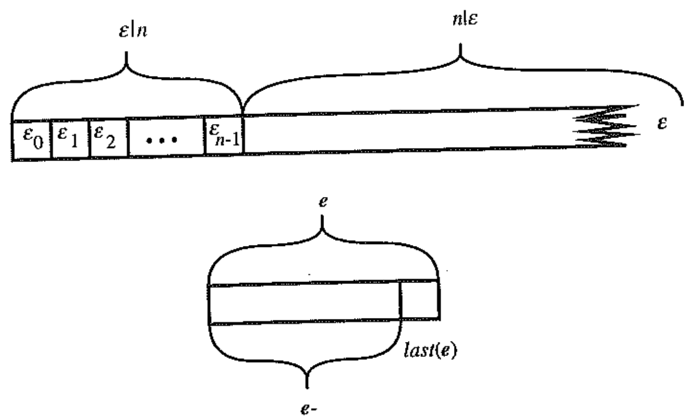

Figure 3.2

data streams and hypotheses. I will also assume that hypotheses are discrete objects coded by natural numbers, so that $H \subseteq \omega$. Thus $C \subseteq \mathcal{N} \times \omega$.

Finally, since I am ignoring the actions recommended by the scientist's assessment method α, I will assume that such a method is an arbitrary, total[2] map from hypotheses and finite data sequences to conjectures (*Fig. 3.3*). *Conjectures* (outputs of the assessment method) will be rational numbers between 0 and 1, together with a special *certainty symbol* '!'.

Now we may consider a traditional inductive problem that is idealized enough to clarify the logic of the approach and yet rich enough to illustrate many different senses of solvability and unsolvability.

Example 3.1 The infinite divisibility of matter

Since classical times, there have been debates concerning the nature of matter. In the seventeenth and eighteenth centuries, this debate took on a renewed urgency. The *corpuscularian* followers of René Descartes believed that matter is continuously extended and hence infinitely divisible. *Atomists* like Newton thought the opposite. Kant held the more measured opinion that the question goes beyond all possible experience:

> For how can we make out by experience ... whether matter is infinitely divisible or consists of simple parts? Such concepts cannot be given in any experience, however extensive, and consequently the falsehood either of the affirmative or the negative proposition cannot be discovered by this touchstone.[3]

The issue has not gone away. It seems that every increase in the power of particle accelerators may lead to a new, submicroscopic universe. Who was right, Descartes, Newton, or Kant?

[2] One could also countenance partial maps and require for convergence that the map be defined on each finite initial segment of the data stream. Nothing of interest depends on this choice.

[3] Kant (1950): 88.

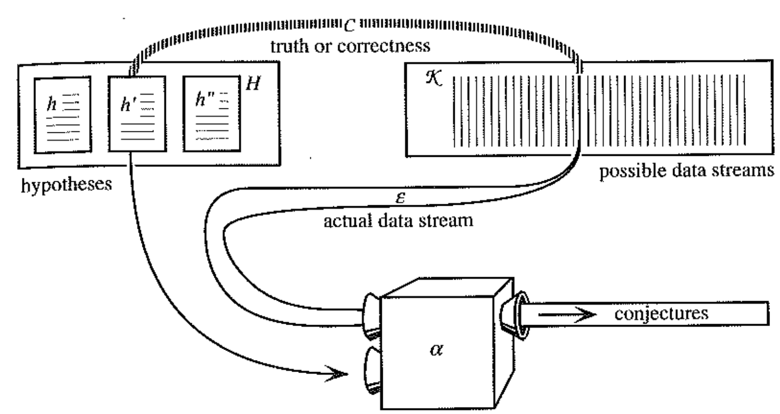

Figure 3.3

Our scientist is a theorist at the IRBP (Institute for Really Big Physics). His attitude toward the infinite divisibility of matter is frankly operational. If no accelerator of any power can split a particle, then he takes it as metaphysical nonsense to say that the particle is divisible, so indivisibility is just indivisibility by some possible accelerator that the IRBP, with its unlimited research budget, could construct. To investigate Kant's question, the theorist is commited to the following experimental procedure.

The lab maintains a list of the particles that have been obtained from previous splits but that have not yet been split themselves. This list is initialized with the original particle p under investigation. At stage n, the lab attempts to split the particle currently at the head of the list. If they succeed, then they remove the split particle from the head of the list and add the new fragments to the tail of the list, write 1 on the data tape, and proceed to stage $n + 1$ without building a larger accelerator. If, however, they fail to split the current particle into any fragments, they report 0. Then they place it at the end of the list; build a new, bigger accelerator; and proceed to stage $n + 1$.

Given no assumptions at all, the reports of the lab may have nothing to do with actual divisions of particles, so that all questions about the divisibility of particles are globally underdetermined: the alleged fragments of particles fed to these accelerators may be artifacts of faulty theory and irrelevant but misleading machinery. To make the issue more interesting, we will grant our theorist some strong theoretical assumptions.

First, we will grant him that the lab accurately reports the result of each attempted cut. It never reports cuts in particles when no such particles or cuts exist, and it always reports cuts that occur. Second, we will grant him that any physically possible cut in a particle given to the lab will eventually be made by his lab. This reflects our theorist's operationist leanings. We will not grant, however, that the very next attempt will succeed if a cut is possible. It may take some time before even the IRBP can secure funding for a sufficiently large accelerator.

Accordingly, if we let h_{inf} denote the hypothesis that particle p is infinitely divisible and let h_{fin} be the hypothesis that p is only finitely divisible, then our assumptions imply that

$$C(\varepsilon, h_{inf}) \Leftrightarrow \textit{infinitely many 1s occur in } \varepsilon.$$

$$C(\varepsilon, h_{fin}) \Leftrightarrow \textit{only finitely many 1s occur in } \varepsilon.$$

The scientist can entertain a variety of hypotheses about the fundamental structure of matter. Let h_{simple} be the hypothesis that the particle p is physically indivisible, and let $h_{divisible}$ be the negation of h_{simple}. Then the scientist's assumptions imply that

$$C(\varepsilon, h_{simple}) \Leftrightarrow \textit{only 0s occur in } \varepsilon.$$

$$C(\varepsilon, h_{divisible}) \Leftrightarrow 1 \textit{ occurs somewhere in } \varepsilon.$$

Let $h_{division\ at\ t}$ be the hypothesis that a successful division will occur at t. Then:

$$C(\varepsilon, h_{division\ at\ t}) \Leftrightarrow \varepsilon_t = 1.$$

$$C(\varepsilon, \neg h_{division\ at\ t}) \Leftrightarrow \varepsilon_t = 0.$$

Let $h_{LRF(1)=r}$ be the hypothesis that the limiting relative frequency of successful cuts in the experiment under consideration exists and is equal to r.

$$C(\varepsilon, h_{LRF(1)=r}) \Leftrightarrow LRF_\varepsilon(1) \textit{ exists and is equal to } r.$$

$$C(\varepsilon, \neg h_{LRF(1)=r}) \Leftrightarrow LRF_\varepsilon(1) \textit{ does not exist or exists but is } \neq r.$$

Let H be the set of all these hypotheses. We will assume that the huge accelerators of the IRBP are new and so unprecedented in power that the theorist has to admit that any sequence of 0s and 1s is possible for all he knows or assumes. So $\mathcal{K}$ is just the set 2^ω of all infinite sequences of 0s and 1s.

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

Now we have a set of hypotheses H, a set of possible data streams $\mathcal{K}$, and an empirical relation of correctness C, reflecting something of the character of Kant's claim. The question is whether Kant was right, and if so, then in what sense?

2. Decidability with a Deadline

We would certainly like science to be guaranteed to succeed by some fixed time, namely, the time we need the answer by. This corresponds to a trivial notion

of convergence, where convergence to an answer by stage n of inquiry means simply producing that answer at n. Let b range over possible conjectures. Then define:

$$\alpha \textit{ produces } b \textit{ at } n \textit{ on } h, \varepsilon \Leftrightarrow \alpha(h, \varepsilon|n) = b.$$

This notion of convergence says nothing about being right or wrong. *Success* is a matter of converging to the right conjecture on a given data stream. Traditionally, methodologists have recognized three different notions of contingent success. Verification requires convergence to 1 when and only when the hypothesis is correct. Refutation demands convergence to 0 when and only when the hypothesis is incorrect. Decision requires both. In the present case we have:

$$\alpha \textit{ verifies}_C\ h \textit{ at } n \textit{ on } \varepsilon \Leftrightarrow [\alpha \textit{ produces } 1 \textit{ at } n \textit{ on } h, \varepsilon \Leftrightarrow C(\varepsilon, h)].$$

$$\alpha \textit{ refutes}_C\ h \textit{ at } n \textit{ on } \varepsilon \Leftrightarrow [\alpha \textit{ produces } 0 \textit{ at } n \textit{ on } h, \varepsilon \Leftrightarrow \neg C(\varepsilon, h)].$$

$$\alpha \textit{ decides}_C\ h \textit{ at } n \textit{ on } \varepsilon \Leftrightarrow \alpha \textit{ both verifies}_C \textit{ and refutes}_C\ h \textit{ by } n \textit{ on } \varepsilon.$$

So α can verify h at n on ε by producing anything but 1 (e.g., 0.5) when h is incorrect.

Reliability specifies the range of possible worlds over which the method must succeed. The logical conception of reliability demands success over all possible worlds in $\mathcal{K}$:

$$\alpha \textit{ verifies}_C\ h \textit{ at } n \textit{ given } \mathcal{K} \Leftrightarrow \textit{for each } \varepsilon \in \mathcal{K}, \alpha \textit{ verifies}_C\ h \textit{ at } n \textit{ on } \varepsilon \textit{ (Fig. 3.4).}$$

$$\alpha \textit{ refutes}_C\ h \textit{ at } n \textit{ given } \mathcal{K} \Leftrightarrow \textit{for each } \varepsilon \in \mathcal{K}, \alpha \textit{ refutes}_C\ h \textit{ at } n \textit{ on } \varepsilon \textit{ (Fig. 3.5).}$$

$$\alpha \textit{ decides}_C\ h \textit{ at } n \textit{ given } \mathcal{K} \Leftrightarrow \textit{for all } \varepsilon \in \mathcal{K}, \alpha \textit{ decides}_C\ h \textit{ at } n \textit{ on } \varepsilon \textit{ (Fig. 3.6).}$$

Given definitions of verification, refutation, and decision *at* a given time, we can define verification, refutation, and decision *by* a given time.

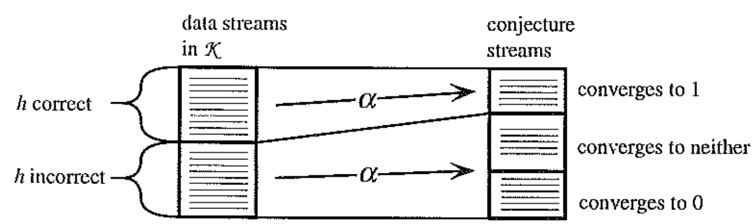

Figure 3.4

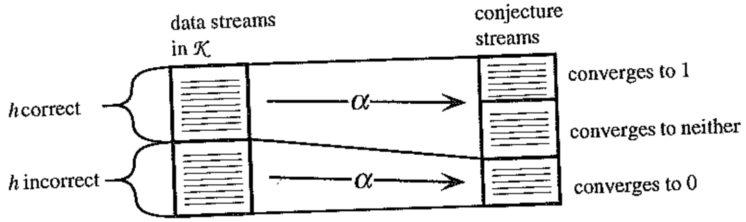

Figure 3.5

$$\alpha \begin{bmatrix} verifies_C \\ refutes_C \\ decides_C \end{bmatrix} h \text{ by } n \text{ given } \mathcal{K} \Leftrightarrow \text{there is } m \leq n \text{ such that } \alpha \begin{bmatrix} verifies_C \\ refutes_C \\ decides_C \end{bmatrix} h \text{ at } m \text{ given } \mathcal{K}.$$

There is another sense of guaranteed success. It concerns the *range* H of hypotheses that the method can reliably assess.[4] $\mathcal{K}$ reflects uncertainty about the world under study, whereas H reflects uncertainty about one's next scientific job assignment. Accordingly, we define:

$$\alpha \begin{bmatrix} verifies_C \\ refutes_C \\ decides_C \end{bmatrix} H \text{ by } n \text{ given } \mathcal{K} \Leftrightarrow \text{for all } h \in H, \; \alpha \begin{bmatrix} verifies_C \\ refutes_C \\ decides_C \end{bmatrix} h \text{ by } n \text{ given } \mathcal{K}.$$

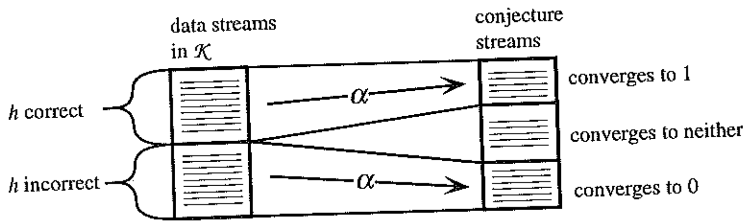

Figure 3.6

[4] Some reliabilist epistemologists seem to confuse these two senses, so that they take a method to be reliable if it *happens* to succeed over a wide range of hypotheses in the actual world. In my view, lucky success over lots of hypotheses does not add up to reliability. It adds up only to lots of lucky success. On my view, a method guaranteed to succeed on a single hypothesis is reliable for that hypothesis even if it fails for every other. I propose that the philosophers in question have confused alleged evidence for reliability with reliability itself. It may be hard to believe that a method could accidentally succeed over a wide range of hypotheses, but what is hard to believe can be true, and here we are defining what reliability *is*, not how we come to believe that a given system is reliable by watching its actual behavior on actual data.

Finally, given a definition of reliability for a particular method, we can immediately ascend to a notion of inductive problem *solvability*. It will be interesting to consider solvability by methods of various different kinds. Let $\mathcal{M}$ be a collection of assessment methods.

$$H \text{ is } \begin{bmatrix} verifiable_C \\ refutable_C \\ decidable_C \end{bmatrix} \text{ by } n \text{ given } \mathcal{K} \text{ by a method in } \mathcal{M}$$
$$\Leftrightarrow \text{there is an } \alpha \in \mathcal{M} \text{ such that } \alpha \begin{bmatrix} verifies_C \\ refutes_C \\ decides_C \end{bmatrix} H \text{ by } n \text{ given } \mathcal{K}.$$

When $H = \{h\}$, we will simply speak of the hypothesis h as being verifiable, refutable, or decidable by a given time. When reference to $\mathcal{M}$ is dropped, it should be understood that $\mathcal{M}$ is the set of all possible assessment methods.

Now we may think of a quadruple $(C, \mathcal{K}, \mathcal{M}, H)$ as an *inductive problem* that is either solvable or unsolvable in the sense described. Each notion of problem solvability to be considered may be analyzed into criteria of convergence, success, reliability, and range of application in this manner. A specification of these criteria will be referred to as a *paradigm*. The paradigm just defined is the *bounded sample decision* paradigm.

A basic fact about this paradigm is that verifiability, refutability, and decidability collapse for ideal methods. So henceforth, we need consider only decidability by n unless restrictions on the method class $\mathcal{M}$ are imposed that preclude the following, trivial construction (*Fig. 3.7*).

Proposition 3.2

(a) *The following statements are equivalent:*

(i) *H is decidable$_C$ by n given $\mathcal{K}$.*

(ii) *H is verifiable$_C$ by n given $\mathcal{K}$.*

(iii) *H is refutable$_C$ by n given $\mathcal{K}$.*

(b) *H is decidable$_C$ by n given $\mathcal{K} \Rightarrow H$ is decidable$_C$ by each $n' \geq n$ given $\mathcal{K}$.*

Proof: (b) is immediate. (a) (i) ⇒ (ii) & (iii) is immediate. (ii) ⇒ (i): Suppose α verifies$_C$ H by n given $\mathcal{K}$. Define α' so that α' produces 1 when α does, and produces 0 otherwise. α' refutes$_C$ H by n given $\mathcal{K}$. (iii) ⇒ (i) is similar. ■

The bounded sample decision paradigm is strict, but it is not impossible to meet. The hypothesis $h_{division\ at\ n}$, which states that the particle at hand is split at stage n, is clearly decidable by stage n simply by running the experiment and by looking at what happens at stage n. It is not hard to see that this hypothesis is not decidable by stage $n - 1$ when $\mathcal{K} = 2^\omega$. Suppose that some hypothesis

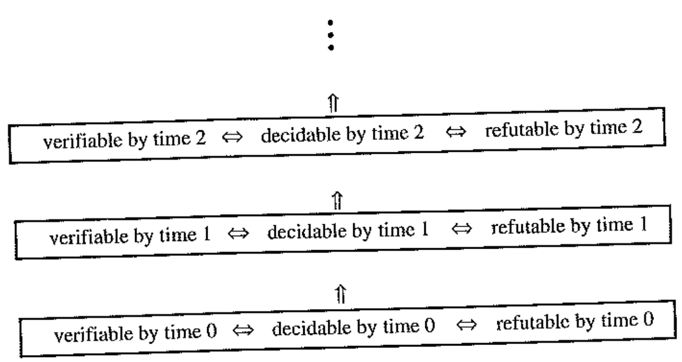

Figure 3.7

assessor α and some stage n are given. Imagine the scientist as being given a tour of a railroad switching yard, in which each track is an infinite data stream. The engineer of the train is the *inductive demon*, who is fiendishly devoted to fooling the scientist. The scientist makes his guesses, and the demon is free to turn switches in the switching yard wherever and whenever he pleases, so long as the infinite path ultimately taken by the train ends up in $\mathcal{K}$. That is, the only constraint on the demon is that he ultimately produce a data stream consistent with the scientist's background assumptions.

The demon proceeds according to the following, simple strategy (*Fig. 3.8*). He presents the scientist's method α with nothing but 0s up to stage $n-1$. Method α must conjecture 0 or 1 at stage $n-1$ in order to succeed by that stage. But if α says 0, the demon presents 1 at stage n and forever after; if α says 1, the demon presents 0 at stage n and forever after. So no matter what α does, it fails on the data stream produced by the demon. But this data stream is in $\mathcal{K}$ because any infinite sequence of 1s and 0s is. Thus α fails to decide$_C$ $h_{division\ at\ n}$ by stage $n-1$ given $\mathcal{K}$. Since α is arbitrary, no possible assessment method succeeds in the required sense.

This is precisely Hume's argument for inductive skepticism. No matter how many sunrises we have seen up to stage $n-1$, the sun may or may not rise at

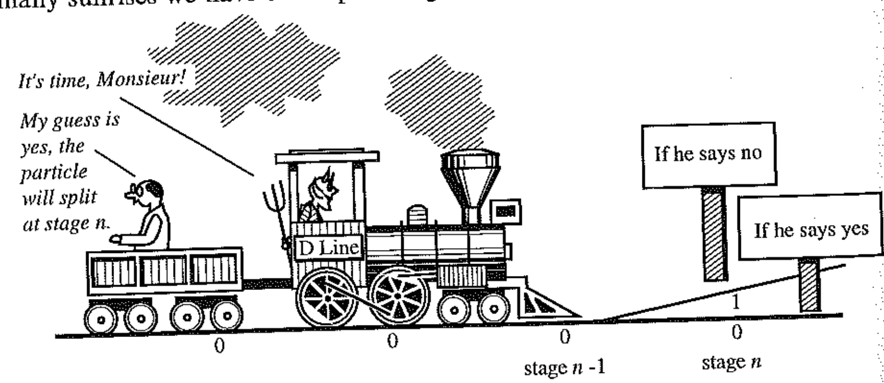

Figure 3.8

stage n. If $\mathcal{K}$ admits both possibilities, the argument is a proof against decidability by stage $n-1$. The same argument shows that $h_{division\ at\ n}$ is not verifiable or refutable by stage $n-1$, either. Since n is arbitrary, this argument shows that the relation given in proposition 3.2(b) is optimal, and hence the levels in Figure 3.7 do not collapse. This is the simplest possible demonic argument against all possible scientific methods. As we shall see, the demon's task gets more difficult as the operative notion of convergence is weakened and as the scientist's knowledge is strengthened.

Scientific realists faced with demonic arguments sometimes object that there *is* no demon, as though one must be justified (in the realist's elusive sense!) in believing that a demon exists before such arguments must be heeded. But the real issue is the existence of a logically reliable method for the scientist, and demons are just a mathematical artifice for proving that no such method exists. If the scientist is actually a demonologist interested in whether or not a demon employing some strategy actually exists, then demonic arguments against this scientist would have to invoke metademons that try to fool the scientist about the demons under study.

The argument just presented is driven by the fact that $\mathcal{K}$ contains data streams presenting a fork to the scientist at stage $n-1$, together with the demand that the scientist arrive at the truth by stage $n-1$ (*Fig. 3.9*). So it can be overturned either by removing data streams from $\mathcal{K}$ (i.e., by adding background assumptions) or by relieving the stringent demand of convergence by a fixed time.

It is only slightly more complicated to show that $h_{divisible}$, which says that p can be divided, is not verifiable, refutable, or decidable by an arbitrary, fixed time given $\mathcal{K}$. Let's consider the case of verification by stage n. Let α be an arbitrary assessment method. The demon's strategy is to present 0s up to stage $n-1$, and to wait for α to produce its conjecture at stage n. If the conjecture is 1, then the demon feeds 0s forever after. If the conjecture is anything but 1, the demon feeds 1s forever after (*Fig. 3.10*). In the former case, α produces 1 at n when $h_{divisible}$ is false, and in the latter case, α fails to produce 1 at n when $h_{divisible}$ is true. Either way, α fails to verify or to decide $h_{divisible}$ at stage n. But α and n are arbitrary, so no possible method can verify or decide $h_{divisible}$ by any fixed time. A similar argument works for the case of refutation, except that the demon produces 1 when the scientist produces anything but 0 and produces 0 otherwise. The same argument works for the hypotheses h_{inf}, h_{fin}, and $h_{LRF(1)=r}$, none of which is verifiable, refutable, or decidable by any fixed time.

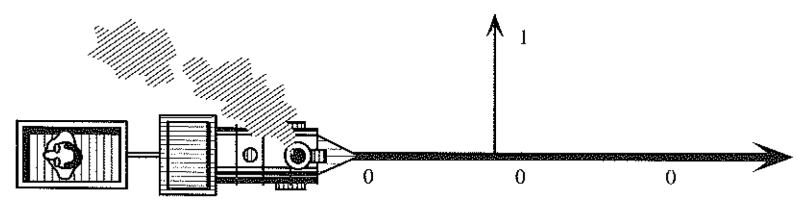

Figure 3.9

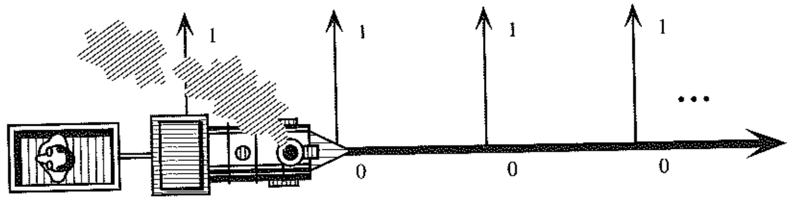

Figure 3.10

Keynes was correct to observe that those of us who are now alive are all dead by some fixed time. But Sextus was equally correct to maintain that by any fixed time we are all unreliable unless we already assume quite a bit. For example, if we assume that if all cuts up to stage n fail, all later cuts fail as well, then $h_{divisible}$ is decidable by stage n.

For those who assume no such thing and who would prefer that inquiry have some logical connection with the truth, it is of interest to investigate weaker senses of convergence that do not demand success by a fixed time, even though it is clear that we would prefer a method guaranteed to succeed in the near future, if such a method were to exist.

3. Decidability, Verifiability, and Refutability with Certainty

If a method is guaranteed to succeed by a given stage n, then given our assumptions $\mathcal{K}$, the method has the truth at stage n. Thus, the method confers a kind of certainty upon its user at n, given $\mathcal{K}$. This certainty is very desirable when it can be had. Since we can't have certainty about infinite divisibility at some fixed time, perhaps we can have it by *some* time that cannot be specified in advance. Since we no longer know a priori when the method will be sure, it must tell us when it is sure by some sign (*Fig. 3.11*).

Any convention will do. I require that the method write an exclamation mark '!' just prior to announcing its certain conjecture. Nothing prevents a method from producing this mark of certainty several times, so we will take its certain conjecture to be the one immediately following the *first* occurrence of '!'. The others will be ignored, as though the method has finished its job and is merely producing irrelevant noise thereafter. Let $b = 0$ or 1. Then define:

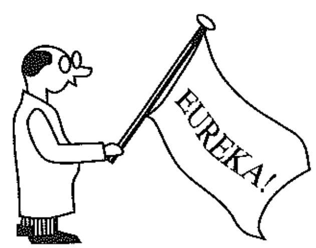

Figure 3.11

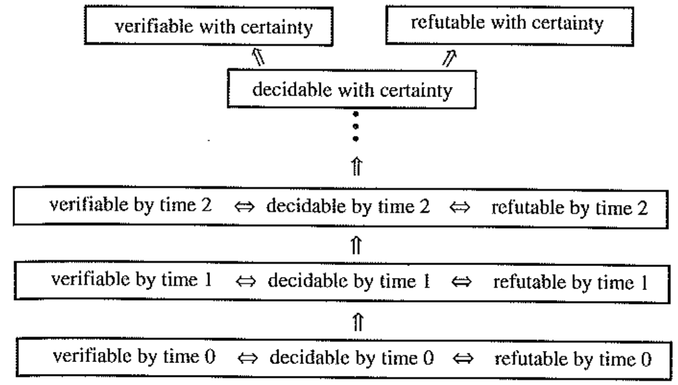

Figure 3.12

α produces b with certainty on h, $\varepsilon \Leftrightarrow$ there is an n such that

(a) $\alpha(h, \varepsilon|n) =$ '!',

(b) $\alpha(h, \varepsilon|n + 1) = b$, *and*

(c) *for each* $m < n$, $\alpha(h, \varepsilon|m) \neq$ '!'.

This concept of convergence induces three notions of success, just as before:

α verifies$_C$ h with certainty on ε
$\Leftrightarrow$ [*α produces* 1 *with certainty on* h, $\varepsilon \Leftrightarrow C(\varepsilon, h)$];

α refutes$_C$ h with certainty on ε
$\Leftrightarrow$ [*α produces* 0 *with certainty on* h, $\varepsilon \Leftrightarrow \neg C(\varepsilon, h)$];

α decides$_C$ h with certainty on ε
$\Leftrightarrow$ *α verifies$_C$ h with certainty on ε and α refutes$_C$ h with certainty on ε.*

These notions of success give rise to corresponding definitions of reliability, range of applicability, and problem solution just as before. Let $\bar{\mathscr{C}}$ denote the complement of $\mathscr{C}$. Then we have the following basic facts relating these paradigms (*Fig. 3.12*).[5]

Proposition 3.3

(a) *H is verifiable$_C$ with certainty given $\mathcal{K}$*
$\Leftrightarrow$ *H is refutable$_{\bar{C}}$ with certainty given $\mathcal{K}$.*

(b) *H is decidable$_C$ with certainty given $\mathcal{K}$*
$\Leftrightarrow$ *H is both verifiable$_C$ and refutable$_C$ with certainty given $\mathcal{K}$.*

(c) *For each n, if H is decidable$_C$ by time n given $\mathcal{K}$ then H is decidable$_C$ with certainty given $\mathcal{K}$.*

[5] In this case, the result (b) is trivial, but we will see that it fails in more complex paradigms.

Proof: (a) (⇒) Let α_1 verify$_C$ H with certainty given $\mathcal{K}$. Define α_2 to conjecture 0 when α_1 conjectures 1 and to conjecture 1 when α_1 conjectures 0. All other conjectures (including '!') are left unaltered. α_2 refutes H with certainty given $\mathcal{K}$ with respect to $\bar{C}$. (⇐) is similar.

(b) (⇒) Immediate. (⇐) Let α_1 verify$_C$ H with certainty given $\mathcal{K}$, and let α_2 refute$_C$ H with certainty given $\mathcal{K}$. Let method α conjecture 0.5 until either α_1 or α_2 conjectures '!'. Then α conjectures '!', and repeats forever whatever the machine that produced '!' says next.

(c) Suppose α_1 decides$_C$ H at n given $\mathcal{K}$. Define α_2 to mimic α_1 exactly, except at $n - 1$, when α_2 produces '!' no matter what. α_2 decides$_C$ H with certainty given $\mathcal{K}$. ■

We have seen that the hypothesis $h_{divisible}$ is not verifiable, refutable, or decidable by any fixed time given 2^ω. We can also show that $h_{divisible}$ is not decidable with certainty given 2^ω, even when we do not insist on success by a fixed time. Once again, a demonic argument suffices. Let assessment method α be given. The demon's strategy is to present the everywhere 0 sequence (indicating no cuts) until the method produces its first mark of certainty. If the next conjecture is 1, the demon feeds 0 forever after. If the next conjecture is 0, the demon feeds 1 forever after. And if the next conjecture is anything else, it doesn't matter what the demon does, so long as he produces a 0–1 data stream, so we will arbitrarily have him produce all 0s. If α never produces a mark of certainty followed by a 0 or a 1, then α clearly fails to decide $h_{divisible}$ with certainty given $\mathcal{K}$. If α produces '!' followed immediately by 0, then $h_{divisible}$ is true so α is wrong. If α produces '!' followed immediately by 1, then $h_{divisible}$ is false, so α is wrong again (*Fig. 3.13*). So in any case, α fails to succeed on the data stream produced by the demon, which is constructed so as to be in $\mathcal{K}$. Since α is arbitrary, no possible assessment method can decide $h_{divisible}$ with certainty given $\mathcal{K}$.

It follows that h_{simple}, the negation of $h_{divisible}$, is not decidable with certainty given $\mathcal{K}$, either. This is just the problem of universal generalization, and the preceding demonic argument was already familiar to philosophers before the time of Sextus Empiricus. On the other hand, $h_{divisible}$ is verifiable with certainty. The method need only produce conjecture 0 until a successful cut is observed, and then produce '!', 1, 1,... forever after. Dually, h_{simple} is refutable with certainty. The method just reverses the conjectures of the method for $h_{divisible}$. It follows that $h_{divisible}$ is not refutable with certainty and h_{simple} is not verifiable with certainty. This is an improvement over the situation regarding decidability,

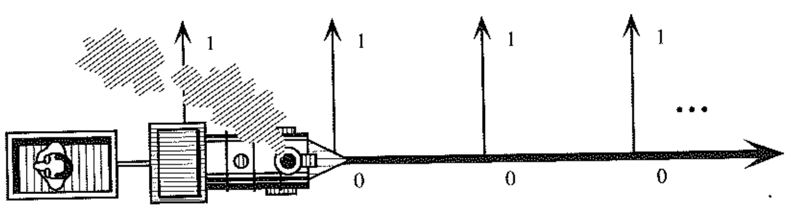

Figure 3.13

verifiability, and refutability by a fixed time, for the same problems are not solvable in those senses. To enhance the scope of reliable inquiry by weakening the skeptic's proposed standard of success (decision with certainty or decision now) is the basic idea behind Popper's falsificationist philosophy of science. Verificationism, the idea that theoremhood can be verified but not refuted with certainty, is the fundamental idea behind the philosophy of mathematical proof.

Let's return to Kant's question about the infinite divisibility of matter. It is readily seen to be neither verifiable nor refutable with certainty. For given an assessment method α, the demon can present α with nothing but successful cuts until α says preceded by '!', at which time the demon provides no more cuts, and α fails to verify infinite divisibility with certainty.

Similarly, the demon can provide only failed cuts until α says 0, after which only successful cuts are announced, so α fails to refute the infinite divisibility hypothesis with certainty. Since α is arbitrary, no possible method can verify or refute infinite divisibility with certainty. Hence, no possible method can decide infinite divisibility with certainty, either. So although verificationism and falsificationism expand the scope of reliable inquiry, they do not save infinite divisibility from Kant's charge of running beyond all possible (local) experience. The situation with limiting relative frequency is even worse. Perhaps we can bring these hypotheses under the purview of logically reliable inquiry by weakening the operative concept of convergence still further.

4. Verification, Refutation, and Decision in the Limit

Certainty demands that a method eventually give a sign that it has arrived at a correct answer. The impossibility of giving such a sign for universal hypotheses is the issue that drives Plato's *Meno* paradox.

> *Meno*: And how will you inquire, Socrates, into that which you do not know? What will you put forth as the subject of inquiry? And if you find what you want, how will you ever know that this is the thing which you did not know?[6]

Meno's underlying assumption is that inquiry is worthless unless it can produce a determinate sign that it has succeeded. But it is possible for a method to be guaranteed to arrive at the truth and to stick with it forever after without ever giving such a sign. Stabilization requires only that inquiry eventually settle down to some fixed assessment value, even though the user of the method may never be sure when stabilization has occurred, since a future reversal is always possible for all he knows (*Fig. 3.14*).

α stabilizes to b for h on ε
⇔ there is a stage n such that for each later stage $m \geq n$, $\alpha(h, \varepsilon|m) = b$.

Nonetheless, after the method has stabilized, no such upsets ever occur

[6] Plato (1949): 36.

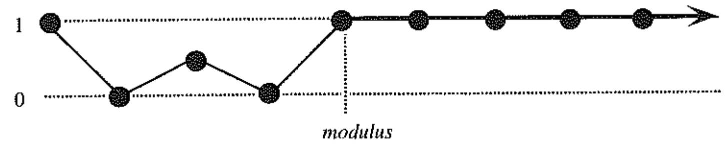

Figure 3.14

again. We will refer to the earliest time after which all conjectures are the same as the *modulus of convergence*.

$$modulus_\alpha(h, \varepsilon) = \text{the least } n \text{ such that for all } m \geq n, \alpha(h, \varepsilon|m) = \alpha(h, \varepsilon|n).$$

The convergence criterion of stabilization in the limit gives rise in the usual way to three different concepts of success on a data stream:

α *verifies*$_C$ *h in the limit on* $\varepsilon \Leftrightarrow [\alpha$ *stabilizes to* 1 *on* $h, \varepsilon \Leftrightarrow C(\varepsilon, h)]$.

α *refutes*$_C$ *h in the limit on* $\varepsilon \Leftrightarrow [\alpha$ *stabilizes to* 0 *on* $h, \varepsilon \Leftrightarrow \neg C(\varepsilon, h)]$.

α *decides*$_C$ *h in the limit on* ε
$\Leftrightarrow \alpha$ *both verifies*$_C$ *and refutes*$_C$ *in the limit on* h, ε.

The corresponding definitions of reliability, range of applicability, and problem solvability proceed just as before, and need not detain us here.

The idea that inquiry should stabilize to the truth has appealed to philosophers for a long time. Indeed, Plato made it a centerpiece of his theory of knowledge, as presented in the *Meno*.

> While [true opinions] abide with us they are beautiful and fruitful, but they run away out of the human soul, and do not remain long, and therefore they are not of much value until they are fastened by the tie of the cause [reason-why]; and this fastening of them, friend Meno, is recollection, as you and I have agreed to call it. But when they are bound, in the first place, they have the nature of knowledge; and, in the second place, they are abiding. And this is why knowledge is more honorable and excellent than true opinion, because fastened by a chain.[7]

One reading of the passage is that belief formed by a method guaranteed to stabilize to the truth (e.g., Platonic recollection) is eventually stable, and when it stabilizes, so that it does not vacillate in the face of true data, it is knowledge. Reliabilist methodology need not endorse this or any other account of knowledge, but insofar as knowledge is held to be at least stable, true belief formed by a reliable process, the analysis of stabilization to the truth is relevant to the theory of knowledge. The basic logical relations among the limiting assessment paradigms are as follows (*Fig. 3.15*):

[7] Plato (1949): 58–59.

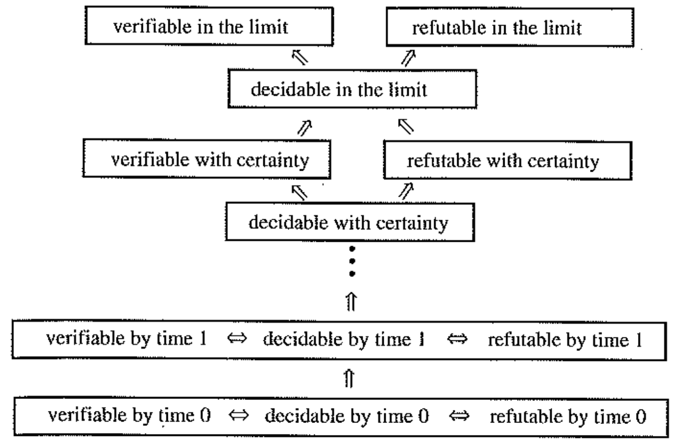

Figure 3.15

Proposition 3.4

(a) *H is verifiable*$_C$ *in the limit given* $\mathcal{K}$
$\Leftrightarrow$ *H is refutable*$_{\bar{C}}$ *in the limit given* $\mathcal{K}$.

(b) *H is decidable*$_C$ *in the limit given* $\mathcal{K}$
$\Leftrightarrow$ *H is both verifiable*$_C$ *and refutable*$_C$ *in the limit given* $\mathcal{K}$.

(c) *H is decidable*$_C$, *verifiable*$_C$, *or refutable*$_C$ *with certainty given* $\mathcal{K}$
$\Rightarrow$ *H is decidable*$_C$ *in the limit given* $\mathcal{K}$.

Proof: Let α_1 verify$_C$ H in the limit given $\mathcal{K}$. Define $\alpha_2(h, e) = 1 - \alpha_1(h, e)$. α_2 refutes H in the limit given $\mathcal{K}$ with respect to $\bar{C}$.

(b) ($\Rightarrow$) Immediate. ($\Leftarrow$) Let α_1 verify$_C$ H in the limit given $\mathcal{K}$, and let α_2 *verify*$_{\bar{C}}$ H in the limit given $\mathcal{K}$, by (a). Define method α as follows. α simulates α_1 and α_2 on each initial segment of the current data $\boldsymbol{e}$. If α_1 says something other than 1 more recently than α_2, then α conjectures 0. Otherwise α conjectures 1.

Let $\varepsilon \in \mathcal{K}$, $h \in H$. If $C(\varepsilon, h)$, then α_1 stabilizes to 1 and α_2 does not, so some time after the modulus of convergence of α_1 on ε, α_2 produces a non-1 conjecture, and α produces 1 thereafter. If $\neg C(\varepsilon, h)$, then α_2 stabilizes to 1 and α_1 does not, so by a similar argument α stabilizes to 0.

(c) Let α_1 verify$_C$ H with certainty given $\mathcal{K}$. To verify$_C$ H in the limit given $\mathcal{K}$, let α conjecture 0 until α_1 produces '!', followed by 1, after which α says 1 no matter what. The decision case follows immediately. A dual construction works for the case of refutation. ■

Despite the lenience of decidability in the limit, it is not a panacea for local underdetermination. For example, without extra background knowledge, the hypothesis that matter is infinitely divisible is not decidable in the limit. The proof is a limiting generalization of the short-run demonic arguments presented earlier.

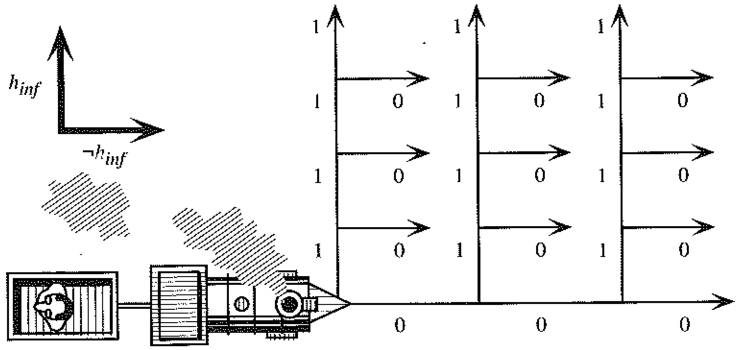

Figure 3.16

Recall the picture of $\mathcal{K}$ as an infinite switchyard through which the inductive demon takes the scientist for a ride. In our problem, $\mathcal{K} = 2^\omega$, so the switchyard is an infinite, binary-branching tree of 1s and 0s (*Fig. 3.16*). The demon starts out by presenting 111 ... until α says 1. If α says 1, the demon switches to all 0s until α says 0. If α then says 0, the demon starts presenting 1s until α says 1. This process continues forever. Either α stabilizes to 1 or to 0 on the data stream provided, or not. Suppose α stabilizes to 1. Then the demon feeds 0s forever after α stabilizes to 1, so h_{inf} is false, and α is wrong. Suppose α stabilizes to 0. Then the demon feeds 1s forever after the modulus of convergence, so h_{inf} is true, and α again fails. But if α does not stabilize to 0 or to 1, α fails again, since α must stabilize to one or the other so long as the data stream contains only 1s and 0s, which it will according to the demon's strategy. So in any event, α fails. Since α is arbitrary, we know that no possible ideal assessment procedure can decide h_{inf} given $\mathcal{K}$ in the limit. This result seems to vindicate Kant's opinion that the question about infinite divisibility goes beyond all possible experience.

On the other hand, h_{inf} is refutable in the limit given $\mathcal{K}$ by a trivial method. Let α_{repeat} simply repeat the last datum it has seen when assessing h_{inf}. If its current datum is a 1, α_{repeat} conjectures 1. If its current datum is 0, α_{repeat} conjectures 0. Let $\varepsilon \in 2^\omega$. Suppose h_{inf} is correct for ε. Then infinitely many 1s occur, so α_{repeat} correctly fails to stabilize to 0. Suppose h_{inf} is incorrect for ε. Then only finitely many 1s occur in ε. After the last 1 is seen, α_{repeat} converges to 0. Thus, the trivial method α_{repeat} refutes h_{inf} in the limit. So in another sense, Kant was wrong. By proposition 4.1, it follows from the fact that h_{inf} is refutable but not decidable in the limit that h_{inf} is not verifiable in the limit. Thus, its negation h_{fin} is verifiable but neither refutable nor decidable in the limit.

There is a tendency for philosophers to suppose that the demon is omnipotent so that the deck is stacked against the scientist. But this is not the case for the *inductive* demons under consideration. Since their outputs are determined locally, just as the scientist's are, the scientist may also be in a position to fool every possible demon. That is just what happens when the scientist uses a logically reliable method. For example, consider the case of verifying h_{fin} in the limit. A would-be demon would have to produce a data

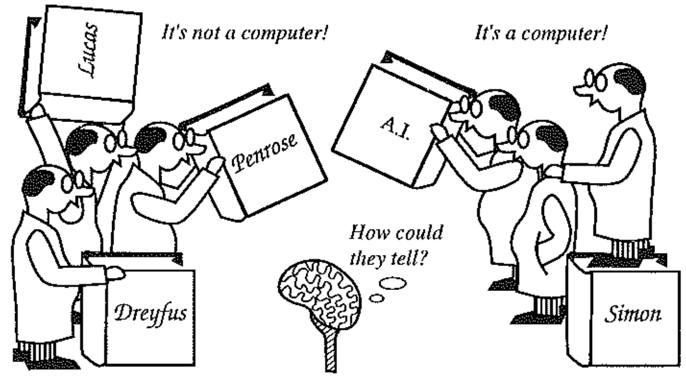

Figure 3.17

stream in $C_{h_{fin}}$ just in case the scientist's conjecture stream does not stabilize to 1. Since the demon only sees the scientist's conjectures as they are produced, he can't tell if the scientist's vacillations will continue forever, or stop. The scientist's successful method of conjecturing the current datum is guaranteed to outwit the demon, since the method converges to 1 if and only if the data stream constructed by the demon does not.

Example 3.5 Cognitive psychology

For many years, cognitive psychologists have been concerned to discover the general computational architecture of the mind. One might object to the implicit assumption that there is such a program to be found. If human behavioral dispositions are actually uncomputable, there is no such program. So it would seem that the first question is whether it is possible, even in the limit, to decide whether or not an arbitrary input-output behavior is computable (*Fig. 3.17*).

J. R. Lucas attempted to give metaphysical arguments based on Gödel's theorem to the effect that human behavior cannot be computable, and recently the physicist R. Penrose has followed suit. A. Newell and H. Simon have responded that the issue is an empirical one: when increasing fragments of human behavior are duplicated by machines, the evidence for the computability of cognition increases.[8] The philosopher H. Dreyfus has answered that individual successes on "microworld" problems will never add up to real intelligence.

> [A]n overall pattern has emerged: success with simple mechanical forms of information processing, great expectations, and then failure when confronted with more complicated forms of behavior. Simon's predictions fall into place as just another example of the phenomenon which Bar-Hillel has called the "fallacy of the successful first step." Simon himself, however, has drawn no such sobering conclusions.[9]

[8] Lucas (1961), Penrose (1989), Newell and Simon (1976).
[9] Dreyfus (1979): 129.

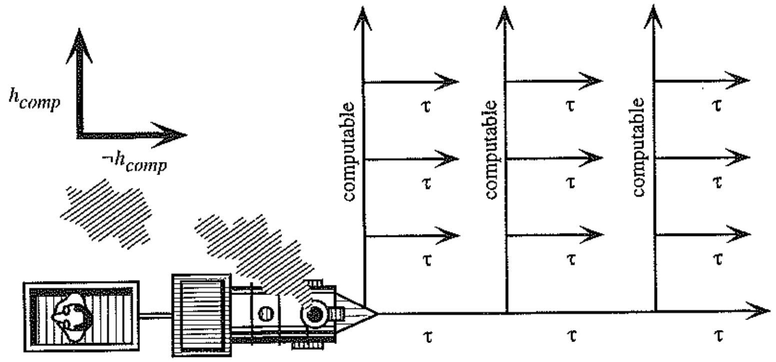

Figure 3.18

The situation is reminiscent of the one concerning the infinite divisibility of matter. Does the question lie beyond all possible experience? In this case, the lab feeds an effective ordering of all possible stimuli to the subject and records the response. To avoid equivocal descriptions of responses, the subject is tied into a Skinner box with only one finger free to press a button (response 0) or to refuse to press it (response 1) within a second. The hypothesis h_{comp} is correct if and only if the response sequence is a computable sequence of 1s and 0s. For all the scientist assumes a priori, any response sequence may arise, so again we let $\mathcal{K} = 2^{\omega}$.

Let α be an assessment method. The demon has at his disposal an infinite list of all the computable data streams in 2^{ω}.[10] He also has in hand a fixed, noncomputable data stream $\tau \in 2^{\omega}$.[11] He feeds α data generated by the first compatable data stream until α eventually produces a non-0 conjecture. He starts presenting τ from the beginning until α says 0. Then he finds the next computable data stream that agrees with the data produced so far and presents it until α makes a non-0 output, and so forth (*Fig. 3.18*). Such a data stream can always be found, since each finite chunk $\boldsymbol{e}$ of data can be recorded in a lookup table by some computer program that computes a data stream extending $\boldsymbol{e}$.

If α stabilizes neither to 1 nor to 0, then α fails to decide h_{comp} in the limit on the data presented. If α stabilizes to 1, then the uncomputable data stream τ is presented in its entirety (after some finite data sequence has been presented). Thus, the overall data stream presented is uncomputable, so h_{comp} is false and α fails to decide h_{comp} in the limit on the data provided. If α stabilizes to 0, then the data stream presented is computable, so h_{comp} is true and α again fails. So in any case, α fails on the data stream presented. Since α is arbitrary, h_{comp} is not decidable in the limit given 2^{ω}.

On the other hand, h_{comp} is verifiable in the limit by an ideal method given 2^{ω}. The ideal method α_{cogsci} maintains a list of all the computable sequences of

[10] There are just countably many since there are at most countably many computer programs to compute them.

[11] Examples of uncomputable functions in 2^{ω} will be given in chapter 6.

1s and 0s. It initializes a pointer to the beginning of this list. Each time the data disagrees with the data stream currently pointed to, the pointer is bumped to the next computable data stream consistent with the data, and the method returns 0. If the data agrees with the sequence the pointer currently points to, the method outputs 1.

If h_{comp} is correct for data stream ε, then the pointer never bumps past some machine in α_{cogsci}'s enumeration, so the method correctly stabilizes to 1. If h_{comp} is incorrect for data stream ε, then ε matches no computable sequence in the list, so the pointer never stops bumping, and infinitely many 0s are output, which is again correct. This *bumping pointer* method is a very general technique for obtaining positive results concerning limiting verifiability, as we shall see in the next chapter.

So computable cognition turns out to be exactly dual to infinite divisibility, so far as ideal inquiry is concerned. The former is verifiable but not refutable in the limit, and the latter is refutable but not verifiable in the limit. Both questions are difficult, but some questions are even harder, as we shall now see.

Example 3.6 Limiting relative frequency again

Let o be a possible datum or outcome. Recall that $F_{\varepsilon}(o, n)$ = the number of occurrences of o in ε up to and including time n, $RF_{\varepsilon}(o, n) = F_n(o, n)/n$, and that the limiting relative frequency of o in ε is defined as follows:

$$LRF_{\varepsilon}(o) = r \Leftrightarrow \textit{for each } s > 0 \textit{ there is an } n \textit{ such that for each } m \geq n, \; |RF_{\varepsilon}(o, m) - r| < s.$$

Statisticians are often concerned with finding out whether some probability lies within a given range. Accordingly, for each set S of reals between 0 and 1, define:

$$LRF_S(o) = \{\varepsilon: LRF_{\varepsilon}(0) \in S\}.$$

I will refer to $LRF_S(o)$ as a *frequency hypothesis for o*. Special cases include *point hypotheses* of form $LRF_{\{r\}}(o)$ and *closed interval hypotheses* of form $LRF_{[r,r']}(o)$. There are also *open interval hypotheses* (e.g., $LRF_{(r,r')}(o)$) and more complicated hypotheses in which S is not an interval.

Now we may ask: What sorts of frequency hypotheses may be decided, verified, or refuted in the limit? When no background assumptions are given, the result is rather dismal.

Proposition 3.7

If $\mathcal{K} = 2^{\omega}$ then no nonempty frequency hypothesis is verifiable in the limit.

Proof: Let $r \in S$. Without loss of generality, we consider the case in which $r > 0$, the case in which $r < 1$ being similar. Let α be an arbitrary assessment method.

Let $q_0, q_1, \ldots, q_n, \ldots$ be an infinite, monotone, increasing sequence of rationals converging to r such that for each i, $0 < q_i < r$. The demon proceeds in stages as follows. At stage 0, the demon's plan is to repeatedly drive the relative frequency of o below q_0 and above q_1. At stage $n + 1$, the demon's plan is to drive the relative frequency of o between q_n and r and then to repeatedly drive the relative frequency of o below q_{n+1} and above q_{n+2} without ever going below q_n or above r. The demon moves from stage n to stage $n + 1$ when the following conditions are met:

(a) α has conjectured at least $n + 1$ non-1s on the data presented so far.
(b) Enough data has been seen to dampen the effect on relative frequency of a single datum so that it is possible for the demon to hold the relative frequency of o above q_n and below r during stage $n + 1$.

Suppose α stabilizes to 1. Then the demon ends up stuck for eternity at some stage $n + 1$, and hence presents a data stream in which the relative frequency of o oscillates forever below q_{n+1} and above q_n, so no limiting relative frequency exists and hence $LRF_S(o)$ is incorrect. So suppose α does not stabilize to 1. Then the demon runs through each stage in the limit, so the relative frequency of o is constrained in ever tighter intervals around r and hence $LRF_S(o)$ is correct (*Fig. 3.19*). So α is wrong in either case. Since α is arbitrary, the result follows. ■

Proposition 3.8

If $\mathcal{K} = 2^\omega$ then no nonempty frequency hypothesis is refutable in the limit.

Proof: Let $r \in S$. Without loss of generality, suppose that $r > 0$, the case in which $r < 1$ being similar. Choose ε so that $LRF_\varepsilon(o) = r$. Let α be an arbitrary assessment method. Let q, q' be rationals such that $0 < q < q' < r$. The demon starts by feeding ε to α until α says something other than 0. Then the demon proceeds to drive the relative frequency of o below q and above q', until α once again says 0 and the relative frequency has been driven below q and above q' at least once. When α says 0, the demon continues feeding ε, from where he left off last. If α stabilizes to 0, then the limiting relative frequency of o in the data stream is r since the data stream is just ε with some initial segment tacked on, so $LRF_S(o)$ is correct. If α produces a conjecture other than 0 infinitely often,

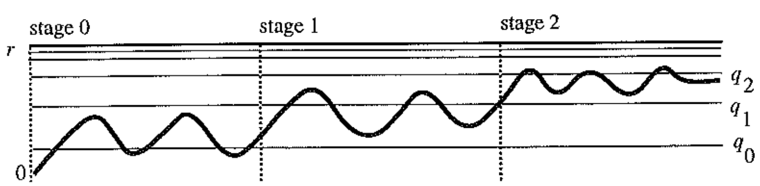

Figure 3.19

then the observed frequency of o vacillates forever outside of the fixed interval from q to q', so no limiting relative frequency exists, and hence $LRF_S(o)$ is false. So α is wrong in either case. Since α is arbitrary, the result follows. ■

In the standard theory of statistical significance tests, it is usually assumed that the frequency hypothesis in question is to be tested only against other frequency hypotheses, and not against the possibility that the limiting relative frequency does not exist. Accordingly, let $LRF(o)$ denote the set of all data streams in which the limiting relative frequency of outcome o is defined. This time the result is more optimistic.

Proposition 3.9

(a) *$LRF_{[r,r']}(o)$ is refutable in the limit given $LRF(o)$.*

(b) *$LRF_{(r,r')}(o)$ is verifiable in the limit given $LRF(o)$.*

Proof: (a) Method α works as follows. α comes equipped with an infinitely repetitive enumeration $q_0, q_1, \ldots, q_n, \ldots$ of the rationals in $[0, 1]$ (i.e., each such rational occurs infinitely often in the enumeration). α starts out with a pointer at q_0. On empty data, α arbitrarily outputs 1 and leaves the pointer at q_0 (*Fig. 3.20*).

Given finite data sequence e, α calculates the relative frequency w of o in e (i.e., $w = RF_e(o, lh(e))$). Let q_i be the rational number pointed to after running α on $e-$. Then α checks whether $w \in [r - q_i, r' + q_i]$. If so, then the pointer is moved one step to the right and α conjectures 1. Otherwise, the pointer is left where it is and α conjectures 0.

Let $\varepsilon \in LRF_{[r,r']}(o)$. Then for some s such that $r \leq s \leq r'$, $LRF_\varepsilon(o) = s$. So for each q_i, there is an n such that $RF_\varepsilon(o, n) \in [r - q_i, r' + q_i]$. So the pointer is bumped infinitely often and α correctly fails to converge to 0. Suppose $\varepsilon \in LRF(o) - LRF_{[r,r']}(o)$. Then since $\varepsilon \in LRF(o)$, there is an $s \notin [r, r']$ such that $LRF_\varepsilon(0) = s$. Without loss of generality consider the case in which $s > r'$. Pick some q_i such that $0 < q_i$ and $q_i + r' < s$.

For some n, we have that for each $m \geq n$, $RF_\varepsilon(o, m) \notin [r - q_i, r' + q_i]$. There is a j past the current pointer position at stage n such that $q_i = q_j$, since the enumeration is infinitely repetitive. Since for each $m \geq n$, $RF_\varepsilon(o, m) \notin$

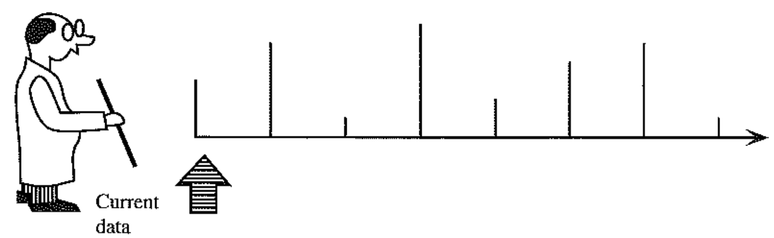

Figure 3.20

$[r - q_i, r' + q_i]$, the pointer is never bumped past j, so α converges to 0, as required.

(b) Similar argument, switching 1 for 0 and $[0, 1] - (r, r')$ for $[r, r']$. ■

Now the question is whether the strategy described could possibly be improved to decide frequency hypotheses in the limit, given that a limiting relative frequency exists. The answer is negative.

Proposition 3.10 (C. Juhl)

(a) *If $0 < r$ or $r' < 1$ then $LRF_{[r,r']}(o)$ is not verifiable in the limit given $LRF(o)$.*

(b) *$LRF_{(r,r')}(o)$ is not refutable in the limit given $LRF(o)$.*

Proof: In light of the preceding two results, it suffices to show that neither hypothesis is decidable in the limit given $LRF(o)$. (a) Without loss of generality, suppose that $r' < 1$. Let $q = 1 - r'$. Let $q_1, q_2, \ldots, q_n, \ldots$ be an infinite descending sequence of nonzero rationals that starts with q and that converges to 0 without ever arriving at 0. The demon starts by feeding data from some data stream ε such that $LRF_\varepsilon(o) = r'$ and such that the relative frequency of o never goes below r'. When α says 1, the demon presents data that drive the relative frequency of o ever closer to $r' + q_1$ without leaving the interval $[r', r' + q_1]$. When α says something other than 1, the demon resumes driving the observed frequency back to r' from above, without leaving the interval $[r', r' + q_1]$. This may not be possible immediately after α says 1, but it will be eventually, once sufficient data have been presented to dampen out the effect on observed relative frequency of a single observation. Therefore, the demon may have to wait some finite time before implementing this plan. The next time α says 1, the cycle repeats, but with q_2 replacing q_1, and so on, forever. If α stabilizes to 1, then the limiting relative frequency of o in the data stream presented lies outside of $[r, r']$, so $LRF_{[r,r']}(o)$ is false and α fails. If α does not stabilize to 1, then either α conjectures 1 infinitely often or not. If α does not conjecture 1 infinitely often, then the demon gets stuck at a given stage and drives the relative frequency to r' so $LRF_{[r,r']}(o)$ is true but α fails to converge to 1. If α does conjecture 1 infinitely often without stabilizing to 1, then the demon goes through infinitely many distinct stages so the limiting relative frequency of o is again r', and $LRF_{[r,r']}(o)$ is once again true (*Fig. 3.21*), so α fails in every case. Since α is arbitrary, no possible assessment method can verify $LRF_{[r,r']}(o)$ in the limit. (b) is similar. ■

The preceding results permit us to relate limiting verification, refutation, and decision to standards of reliability more frequently encountered in classical statistics. From a sufficiently abstract point of view, a *statistical test* is an assessment method that conjectures 0 or 1 on each finite data sequence, where

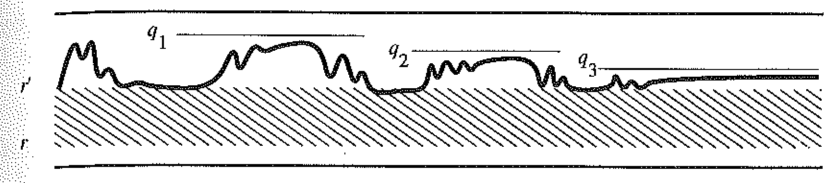

Figure 3.21

1 is read as *acceptance* and 0 is read as *rejection.* In statistical jargon, a *type 1 error* occurs when a true hypothesis is rejected and a *type 2 error* occurs when a false hypothesis is accepted. A standard goal of statistical testing is then to logically guarantee no more than a given limiting relative frequency r_1 of type 1 error, while minimizing the limiting relative frequency r_2 of type 2 error as much as possible given r_1. To minimize the limiting relative frequency of type 2 error is to maximize the limiting relative frequency $1 - r_2$ of rejecting the hypothesis when it is false. Neyman put the matter this way:

> It may often be proved that if we behave according to such a rule, then in the long run we shall reject h when it is true not more, say, than once in a hundred times, and in addition we may have evidence that we shall reject h sufficiently often when it is false.[12]

If such an r_1 exists, then it is said to be the *significance* of the test (with respect to h and $\mathcal{K}$). If such an r_2 exists, then $1 - r_2$ is called the *power* of the test (with respect to h and $\mathcal{K}$). A test is said to be *biased* when its power does not exceed its significance. In such a case, the limiting relative frequency of rejecting the hypothesis when it is false does not exceed the limiting relative frequency of rejecting it when it is true, so that flipping a fair coin to decide rejection would do as well as using the test so far as limiting relative frequencies of error are concerned. For this reason, biased tests are sometimes said to be *less than useless.* It is therefore of interest to examine the conditions under which an *unbiased* test exists.

It turns out that the relationship between unbiased testability and stabilization to the truth in the limit is sensitive to whether the significance level is strictly greater than or merely no less than the limiting relative frequency of type 1 error over all of $\mathcal{K}$. We may call the test *open* in the former case and *closed* in the latter. These ideas yield a learning-theoretic paradigm with a statistical flavor. Define:

$$\alpha(h, \varepsilon) = \tau \Leftrightarrow \forall n,\ \tau_n = \alpha(h, \varepsilon|n).$$

[12] Neyman and Pearson (1933): 142.

In other words, $\alpha(h, \varepsilon)$ denotes the infinite conjecture sequence produced by α as more and more of ε is read. Let $\mathscr{C}_h$ denote the set of all data streams for which h is correct. Then define:[13]

> *α is an open unbiased test$_C$ for h given $\mathcal{K}$ at significance r* $\Leftrightarrow \forall \varepsilon \in \mathcal{K} - C_h$, *$LRF_{\alpha(h,\varepsilon)}(0)$ exists and is $\geq r$ and* $\forall \varepsilon \in \mathcal{K} \cap C_h$, *$LRF_{\alpha(h,\varepsilon)}(0)$ exists and is $< r$.*
>
> *α is a closed unbiased test$_C$ for h given $\mathcal{K}$ at significance r* $\Leftrightarrow \forall \varepsilon \in \mathcal{K} - C_h$, *$LRF_{\alpha(h,\varepsilon)}(0)$ exists and is $> r$ and* $\forall \varepsilon \in \mathcal{K} \cap C_h$, *$LRF_{\alpha(h,\varepsilon)}(0)$ exists and is $\leq r$.*
>
> *h is open [closed] unbiased testable$_C$ given $\mathcal{K}$*
> $\Leftrightarrow$ *there is an α and an r such that α is an open [closed, clopen] unbiased test$_C$ for h given $\mathcal{K}$ at significance r.*

We may now consider the relationship between unbiased testability and limiting verification and refutation. The result is one of exact equivalence, depending on whether the limiting frequency of type 1 errror is $\leq r$ or $< r$.

Proposition 3.11 (with C. Juhl)

(a) *h is open unbiased testable$_C$ given $\mathcal{K}$*
$\Leftrightarrow$ *h is verifiable$_C$ in the limit given $\mathcal{K}$.*

(b) *h is closed unbiased testable$_C$ given $\mathcal{K}$*
$\Leftrightarrow$ *h is refutable$_C$ in the limit given $\mathcal{K}$.*

Proof: (a) ($\Rightarrow$) Let α be an open unbiased test$_C$ of h given $\mathcal{K}$ with significance level r. We construct a limiting verifier$_C$ of h given $\mathcal{K}$ that uses the successive conjectures of α. The proof of proposition 3.9(b) yields that $LRF_{[0,r)}(0)$ is verifiable in the limit given $LRF(0)$, say by method β. Let $\gamma = \beta \circ \alpha$. Let $\varepsilon \in C_h \cap \mathcal{K}$. Then $LRF_{\alpha(h,\varepsilon)}(0)$ exists and is $<r$. Let $\alpha(h, \varepsilon)$ denote the infinite sequence of conjectures $\alpha(h, \varepsilon|0)$, $\alpha(h, \varepsilon|1), \ldots$. Hence, $\alpha(h, \varepsilon) \in LRF_{[0,r)}(0)$, so β stabilizes to 1 on $\alpha(h, \varepsilon)$. So γ stabilizes to 1 on ε. Let $\varepsilon \in \mathcal{K} - C_h$. Then $LRF_{\alpha(h,\varepsilon)}(0)$ exists and is $\geq r$. Hence, $\alpha(h, \varepsilon) \notin LRF_{[0,r)}(0)$, so β does not stabilize to 1 on $\alpha(h, \varepsilon)$. So γ does not stabilize to 1 on ε.

($\Leftarrow$) Suppose that α verifies$_C$ h in the limit given $\mathcal{K}$. We construct an open, unbiased test with significance level 1 that uses the conjectures of α. β maintains a list $\boldsymbol{c}$ of natural numbers initialized to (0). β simulates α on the data as it feeds in and adds n to the end of the list when α has produced exactly n conjectures less than 1. β produces its current conjecture in accordance with

[13] The senses of statistical testability defined here are more general than usual in the sense that real tests produce conjectures depending only on the current sample, without reference to the samples taken before. They are more specific in the sense that, statistical tests are supposed to work when the truth of the hypothesis changes spontaneously from test to test, whereas in the paradigm just defined it is assumed that the truth of h is fixed for all time.

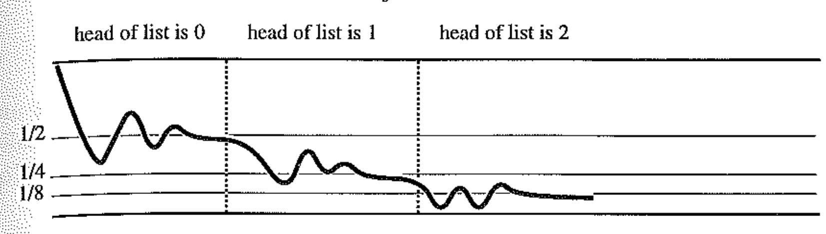

Figure 3.22

the number at the head of the list $\boldsymbol{c}$. If the number at the head of the list is currently n, then β produces conjectures in $\{0, 1\}$ in such a way as to drive the relative frequency of 1s to $1/2^{n+1}$ without ever allowing the observed frequency to exceed $1/2^n$. This is trivial when $n = 0$. If $n + 1$ occurs next in the list, then when (a) the current frequency of 1 conjectures is between $1/2^n$ and $1/2^{n+1}$ and (b) enough conjectures have been made so that it is possible not to exceed $1/2^{n+1}$ at the next stage, then the first item in $\boldsymbol{c}$ is deleted (*Fig. 3.22*).

Let $\varepsilon \in \mathcal{K} \cap C_h$. Then α stabilizes to 1. So only finitely many numbers are ever added to $\boldsymbol{c}$, so $LRF_{\beta(h,\varepsilon)}(1) > 0$ and hence $LRF_{\beta(h,\varepsilon)}(0) < 1$. Let $\varepsilon \in \mathcal{K} - C_h$. Then every number is eventually added to $\boldsymbol{c}$, so for each n, the relative frequency of 1s conjectured by β on ε eventually remains below $1/2^n$ and hence $LRF_{\beta(h,\varepsilon)}(1) = 0$, so $LRF_{\beta(h,\varepsilon)}(0) = 1$. Hence, β is an open unbiased test$_C$ of h given $\mathcal{K}$ with significance level 1. (b) follows by a similar argument. ■

Frequentists usually add to background knowledge the assumption that the data stream is random.[14] A *place selection* is a total function that picks a position in the data stream beyond position n when provided with a finite initial segment of the data stream of length n. Then if we run the place selection function over the entire data stream, it will eventually pick out some infinite subsequence of the original data stream. A place selection π is said to be a *betting system* for ε just in case the limiting relative frequency of some datum o in the subsequence of ε selected by π is different from the limiting relative frequency of o in ε. Let $\mathcal{PS}$ be a fixed, countable collection of place selections. Let ε be a data stream in which the limiting relative frequency of each datum occurring in ε exists. Then ε is *$\mathcal{PS}$-random* just in case there is no betting system for ε in $\mathcal{PS}$. Since randomness is a property of data streams, it is yet another empirical assumption that is subject to empirical scrutiny (cf. exercise 3.3). If it is assumed as background knowledge, however, then the ($\Rightarrow$) side of the preceding result may not hold, since there is no guarantee that the data stream produced by the demon in the proof is random. The effect of randomness assumptions upon logical reliability is an important issue for further study.

[14] Von Mises (1981).

5. Decision with *n* Mind Changes

Whenever α stabilizes to b for h on ε, α changes its conjecture about h only finitely many times. The number of times α changes its conjecture will be called the number of *mind changes* of α for h on ε, which I denote $mc_\alpha(h, \varepsilon)$. More precisely,

$$mc_\alpha(h, \varepsilon) = |\{n \in \omega: \alpha(h, \varepsilon|n) \neq \alpha(h, \varepsilon|n+1)\}|.$$

Decision in the limit countenances an arbitrary number of mind changes by α. The prospect of surprises in the future is what keeps the scientist's stomach churning. It would be nice if some a priori bound could be placed on the number of mind changes the scientist will encounter. Then if the method happens to use up all of its mind changes, the scientist can quit with certainty.

When mind changes are counted, it turns out to matter what sort of conjecture the method α starts with before seeing any data. In the case of a universal generalization such as "all ravens are black," the scientist can succeed with one mind change starting with 1 by assuming the hypothesis is true until it is refuted. In the case of an existential hypothesis, the scientist can succeed in one mind change starting with 0 until the hypothesis is verified. We may then speak of α deciding h in n mind changes *starting with b*. Recall that $\mathbf{0}$ denotes the empty data sequence.

(M1) *α decides$_C$ h with n mind changes starting with b on $\varepsilon \Leftrightarrow$*

(a) *α decides$_C$ h in the limit on ε and*

(b) *$mc_\alpha(h, \varepsilon) \leq n$ and*

(c) $\alpha(h, \mathbf{0}) = b$.

Defining the associated notions of reliability, applicability, and problem solvability is straightforward and is left to the reader.

The following proposition summarizes the elementary properties of the bounded mind-change paradigm (*Fig. 3.23*).

Proposition 3.12

For $n \geq 0$ and for all r such that $0 < r < 1$,

(a) *H is decidable$_C$ with n mind changes starting with 1 given $\mathcal{K}$*
$\Leftrightarrow$ H is decidable$_C$ with n mind changes starting with 0 given $\mathcal{K}$.

(b) *H is decidable$_C$ with n mind changes starting with r given $\mathcal{K}$*
$\Leftrightarrow$ H is decidable$_C$ with n mind changes starting with 1 given $\mathcal{K}$ and H is decidable$_C$ with n mind changes starting with 0 given $\mathcal{K}$.

(c) *H is decidable$_C$ with n mind changes given $\mathcal{K}$*
$\Leftrightarrow$ H is decidable$_C$ with n mind changes starting with 1 given $\mathcal{K}$ or H is decidable$_C$ with n mind changes starting with 0 given $\mathcal{K}$.

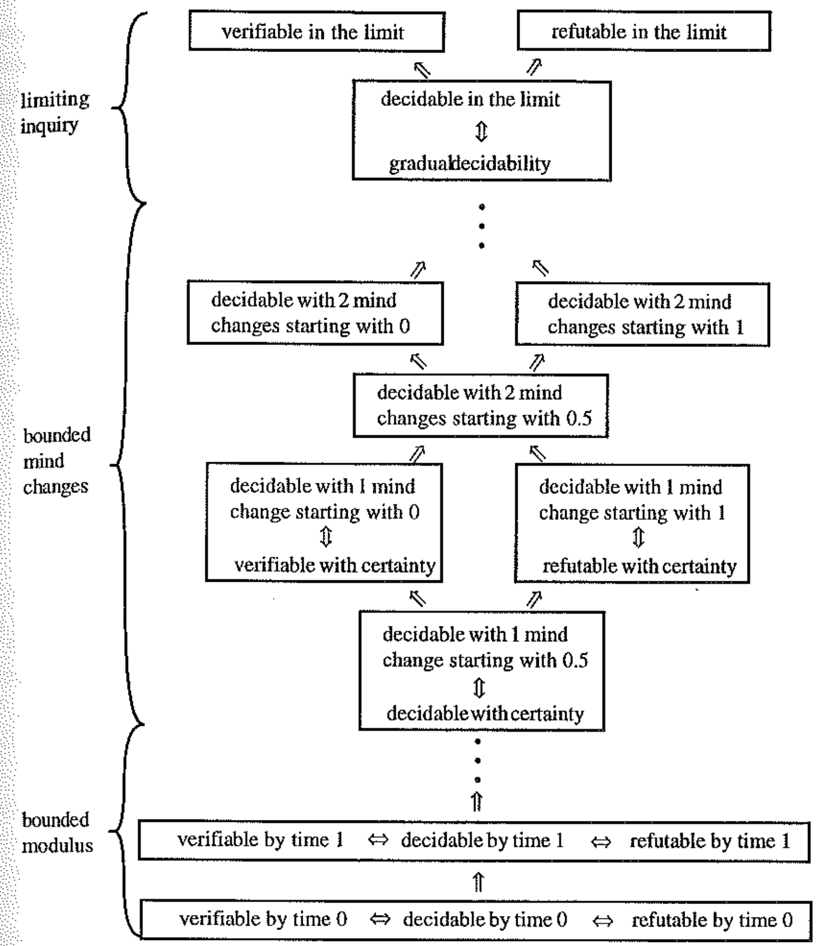

Figure 3.23

(d.1) *H is refutable$_C$ with certainty given $\mathcal{K}$*
$\Leftrightarrow$ H is decidable$_C$ with one mind change starting with 1 given $\mathcal{K}$.

(d.2) *H is verifiable$_C$ with certainty given $\mathcal{K}$*
$\Leftrightarrow$ H is decidable$_C$ with one mind change starting with 0 given $\mathcal{K}$.

(d.3) *H is decidable$_C$ with certainty given $\mathcal{K}$*
$\Leftrightarrow$ H is decidable$_C$ with one mind change starting with r given $\mathcal{K}$.

Proof: Exercise 3.2. ■

For each natural number n and $b \in \{0, 1\}$, there are hypotheses that can be decided with n mind changes starting with b, but that cannot be decided with fewer than n mind changes starting with b. For example, consider the hypothesis "there is exactly one black raven." A scientist can conjecture that this hypothesis is false until a black raven is seen. Then the scientist conjectures that it is true until another black raven is seen. In the worst case, this scientist changes his mind twice, starting with 0. The hypothesis "either there is exactly one black raven or there are exactly three black ravens" requires another mind change. Continuing the sequence in this manner shows that the implications among the mind-change paradigms are all proper.

6. Gradual Verification, Refutation, and Decision

We have seen that limiting relative frequencies fall beyond the scope of reliable verification or refutation in the limit. This suggests that another weakening of the notion of convergence should be considered. We will say that α approaches b just in case α's conjectures get closer and closer to b, perhaps without ever reaching it.

$$\alpha \text{ approaches } b \text{ on } h, \varepsilon \Leftrightarrow \text{for each rational } s \in (0, 1], \text{ there is a stage } n \text{ such that for each later stage } m \geq n, |b - \alpha(h, \varepsilon|m)| \leq s.$$

I refer to rationals in (0, 1] as *degrees of approximation*. If α approaches b on h, ε, then we may think of α as stabilizing with respect to each degree of approximation s by eventually remaining within s of b. Accordingly, the modulus of convergence for a given degree of approximation is the least time after which the conjectures of α remain always within that degree of b (*Fig. 3.24*).

$$modulus_\alpha(h, s, b, \varepsilon) = \text{the least } n \text{ such that for all } m \geq n, |b - \alpha(h, \varepsilon|m)| \leq s.$$

Success may now be defined as follows:

$$\alpha \text{ verifies}_C\ h \text{ gradually on } \varepsilon \Leftrightarrow [\alpha \text{ approaches } 1 \text{ on } h, \varepsilon \Leftrightarrow C(\varepsilon, h)].$$

$$\alpha \text{ refutes}_C\ h \text{ gradually on } \varepsilon \Leftrightarrow [\alpha \text{ approaches } 0 \text{ on } h, \varepsilon \Leftrightarrow \neg C(\varepsilon, h)].$$

$$\alpha \text{ decides}_C\ h \text{ gradually on } \varepsilon \Leftrightarrow [\alpha \text{ verifies}_C \text{ and refutes}_C\ h \text{ gradually on } h, \varepsilon].$$

The corresponding definitions of reliability, range of applicability, and problem solvability are again left to the reader.

When we turn to the logical relations among these paradigms, we discover

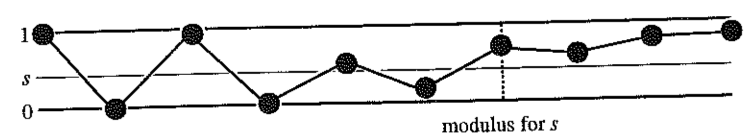

Figure 3.24

that gradual decidability is equivalent to decidability in the limit. Thus, gradual verifiability and refutability do not jointly imply gradual decision in the limit. In other words, it is not in general possible to construct a gradual decider out of a gradual verifier and a gradual refuter, contrary to the situation in the limiting case! It will be seen that these implications cannot be reversed or strengthened.

Proposition 3.13

(a) *H is verifiable$_C$ gradually given $\mathcal{K}$*
$\Leftrightarrow$ *H is refutable$_{\bar{C}}$ gradually given $\mathcal{K}$.*

(b) *H is decidable$_C$ gradually given $\mathcal{K}$*
$\Leftrightarrow$ *H is decidable$_C$ in the limit given $\mathcal{K}$.*

(c) *H is decidable$_C$, verifiable$_C$, or refutable$_C$ in the limit given $\mathcal{K}$*
$\Rightarrow$ *H is verifiable$_C$ and refutable$_C$ gradually given $\mathcal{K}$.*

Proof: (a) As usual. (b) ($\Leftarrow$) trivial, since stabilization to b implies approach to b. ($\Rightarrow$) Let α gradually decide$_C$ H given $\mathcal{K}$. Define:

$$\beta(h, e) = \begin{cases} 1 & \text{if } \alpha(h, e) > 0.5 \\ 0 & \text{otherwise.} \end{cases}$$

β stabilizes to 1 if and only if α approaches 1, and similarly for 0. Thus β decides$_C$ H in the limit given $\mathcal{K}$. (c) It is immediate that verification in the limit implies gradual verification, and similarly for refutation. Suppose α_1 verifies$_C$ H in the limit given $\mathcal{K}$. Let α_2 proceed as follows. α_2 simulates α_1 on each initial segment of the current data e, and counts how many times α_1 makes a conjecture other than 1. Call the count k. Then α_2 conjectures $1/2^k$ (*Fig. 3.25*).

Let $\varepsilon \in \mathcal{K}$, $h \in H$, and suppose $C(\varepsilon, h)$. Then α_1 stabilizes to 1, so α_2 stabilizes to some value $1/2^k$, and hence does not approach 0. If $\neg C(\varepsilon, h)$, then α_1 does

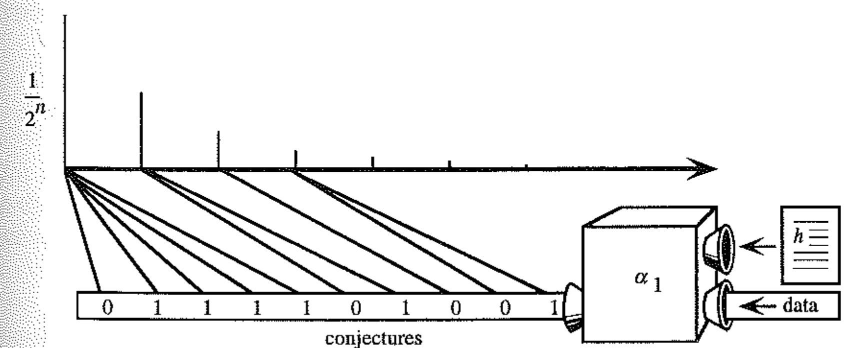

Figure 3.25

not stabilize to 1, and hence emits infinitely many non-1 conjectures, so α_2 approaches 0. Thus, α_2 approaches 0 if and only if $\neg C(\varepsilon, h)$, and hence gradually refutes$_C$ H given $\mathcal{K}$. The implication from refutability in the limit to gradual verifiability may be established in a similar manner. ■

We have already seen several examples of hypotheses that are verifiable but not refutable or decidable in the limit. By proposition 3.13(c), such hypotheses are both gradually verifiable and gradually refutable. But by proposition 3.13(b), such hypotheses are not gradually decidable. Thus, 3.13(c) cannot be strengthened. *Figure 3.26* summarizes all the paradigms introduced

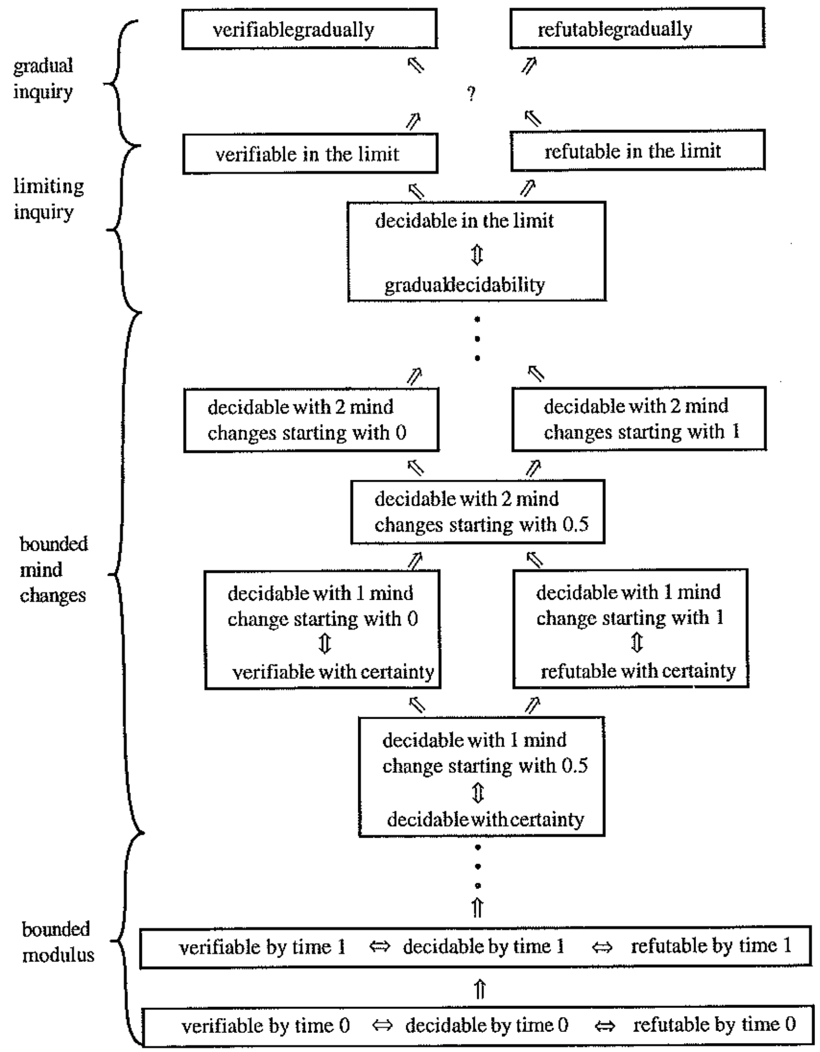

Figure 3.26

so far. The question mark indicates where we might have expected gradual decidability to be, in analogy with the lower levels of the diagram.

Unlike gradual decidability, gradual verification and refutation do increase the scope of reliable inquiry. In particular, limiting relative frequencies are gradually verifiable given 2^ω, as we shall now see. Recall that $LRF_{\{r\}}(o)$ says that the actual limiting relative frequency of o will be r.

Proposition 3.14

(a) *$LRF_{\{r\}}(o)$ is verifiable gradually given 2^ω.*

(b) *$LRF_{\{r\}}(o)$ is not refutable gradually given 2^ω.*

Proof: (a) Let $\alpha(e) = 1 - |RF_e(o, lh(e)) - r|$. Method α evidently verifies $LRF_{\{r\}}(o)$ gradually given 2^ω.

(b) Either $r \neq 1$ or $r \neq 0$. Without loss of generality, suppose $r < 1$. Let $q = 1 - r'$. Let $q_1, q_2, \ldots, q_n, \ldots$ be an infinite, descending sequence of nonzero rationals that starts with q and that converges to 0. The demon proceeds as follows. When α produces a conjecture between $1/2^n$ and $1/2^{n+1}$, the demon adopts the plan of bouncing the observed relative frequency on either side of $[r + q_{n+2}, r + q_{n+1}]$ without leaving the interval $[r, r + q_n]$. The plan is not implemented until enough stages have elapsed to dampen the effect on relative frequency of a single datum so the plan can be implemented. The plan is then implemented at least long enough for one bounce to be accomplished and the demon remembers what α does while the plan is in place. Then the demon puts the plan corresponding to α's next conjecture into place, and so forth. If α approaches 0, then the demon drives the observed frequency into smaller and smaller intervals around r, so that $LRF_{\{r\}}(o)$ is true and α is wrong. If α fails to approach 0, then α produces a conjecture greater than $1/2^n$ infinitely often, so the demon makes the data bounce forever to either side of $[r + q_n, r + q_{n+1}]$. Thus no limiting relative frequency exists. ■

So gradual inquiry extends the scope of reliable inquiry to statistical point hypotheses, at least in the one-sided sense of verifiability.

7. Optimal Background Assumptions

One important task of inductive methodology is to optimize inductive methods so that they are reliable given the weakest possible assumptions. This raises the question whether there are such assumptions. Define:

$\mathcal{K}$ is the optimum assumption for verifying$_C$ H in the limit $\Leftrightarrow$

(a) *H is verifiable$_C$ in the limit given $\mathcal{K}$ and*

(b) *for all $\mathcal{J}$ such that H is verifiable$_C$ in the limit given $\mathcal{J}$, $\mathcal{J} \subseteq \mathcal{K}$,* and similarly for each of the other standards of success. Of course, whenever H is verifiable$_C$ in the limit given $\mathcal{N}$, $\mathcal{N}$ represents the optimum assumption for the problem, since $\mathcal{N}$ contains all the data streams we are considering in our setting. Are there any other examples of optimum assumptions? In fact there are none. And the situation is even worse than that. Define:

$\mathcal{K}$ is optimal for verifying$_C$ H in the limit $\Leftrightarrow$

(a) *H is verifiable$_C$ in the limit given $\mathcal{K}$ and*

(b) *for all $\mathcal{J}$ such that $\mathcal{K} \subset \mathcal{J}$, H is not verifiable$_C$ in the limit given $\mathcal{J}$.*

The optimum assumption (if it exists) is optimal, but the converse may fail if there are several, distinct, optimal sets of data streams. It will be useful to have a simple way to refer to the set of all data streams on which a given method α succeeds. Accordingly, define:

$$limver_{C,H}(\alpha) = \{\varepsilon: \alpha \text{ verifies}_C\ H \text{ in the limit on } \varepsilon\}.$$

If $limver_{C,H}(\alpha)$ is the optimum assumption for verifying$_C$ H in the limit, then we say that α is an *optimum* limiting verifier$_C$ for H. If $limver_{C,H}(\alpha)$ is optimal for verifying$_C$ H in the limit, then we say that α is an *optimal* limiting verifier$_C$ for H. Now it turns out that:

Proposition 3.15

If H is not verifiable$_C$ in the limit given $\mathcal{N}$ then

(a) *no $\mathcal{K} \subseteq \mathcal{N}$ is optimal for verifying$_C$ H in the limit, and hence*

(b) *no α is an optimal limiting verifier$_C$ for H.*

Proof: Suppose H is not verifiable$_C$ in the limit given $\mathcal{N}$. Let α be a method and let $\mathcal{K} = limver_{C,H}(\alpha)$. Then $\mathcal{K} \subset \mathcal{N}$. Choose $\varepsilon \in \mathcal{N} - \mathcal{K}$. Define:

$$\alpha'(h, e) = \begin{cases} 1 & \text{if } e \subset \varepsilon \text{ and } \varepsilon \in C_h \\ 0 & \text{if } e \subset \varepsilon \text{ and } \varepsilon \notin C_h \\ \alpha'(h, e) & \text{otherwise.} \end{cases}$$

α' verifies$_C$ H in the limit given $\mathcal{K} \cup \{\varepsilon\}$. ■

One can view the situation described in proposition 3.15 either as a disaster (no optimal methods are available) or an an embarrassment of riches, as when a genie offers any (finite) amount of money you please. Any amount asked for could have been larger, but it is hard to call the situation bad, except insofar as a kibbitzer can always chide you for not asking for more. This situation is typical of worst-case methodology and will be seen to arise in many different

ways when we consider computable methods. The patching argument of proposition 3.15 also applies to refutation and decision in the limit, and to gradual refutation and verification. It does not apply in the case of verification with certainty, however, as we shall see in the next chapter with the help of some topological concepts.

Exercises

*3.1. Recall Kant's claim that infinite divisibility and composition by simples are contradictories. This isn't so clear. Leibniz's model of the plenum packed an infinity of spheres in a finite space by filling interstices between larger spheres with smaller ones, and so on, until every point in the volume is included in some sphere (*Fig. 3.27*). Assuming that these spheres are simple particles, we have an infinitely divisible finite body that is composed of simples, so the contradiction disappears.

Imagine the IRBP attempting splits at higher and higher energies, and then attempting to split the results at higher and higher energies, etc. Given the theorist's assumption that each possible split is eventually found, a mass is composed of simples just in case this procedure leads in the limit to a history of splits such that each product of fission is either simple itself or leads eventually to a simple particle. In the case of a finite composite of simples, each path in the history of splits terminates in an indivisible particle, but in an infinite composite of simples like Leibniz' plenum, there will be infinite paths of splits. Nonetheless, each fission product is either simple or gives rise to a simple particle later (*Fig. 3.28*).

The denial of composition by simples is not just infinite divisibility, but rather the possession of a fragment or region that contains no simples. This amounts to the existence of a fragment, every subfragment of which is divisible. Such a mass is *somewhere densely divisible*. A somewhere densely divisible mass may either *contain a simple* (i.e., some indivisible fragment is eventually reached) or *be everywhere densely divisible* (i.e., each part is divisible). Nothing prevents a mixture of simples together with densely divisible parts. In fact, J. J. Thomson's *raisin pudding* theory of the atom had just this character, for electrons were envisioned as simple particles floating in an undifferentiated "smear" of positive charge (*Fig. 3.29*).

So perhaps Kant was merely being sloppy when he opposed composition by simples with infinite divisibility, rather than with somewhere dense divisibility. To study such hypotheses, the scientist must modify the experimental design at the IRBP. Instead of merely writing down 1 or 0 to indicate *successful cut* and *failed cut*, respectively, he instructs the lab to assign a new name to each new particle discovered, and to write down at each stage the list $(p_k, (p_{m_1}, \ldots, p_{m_n}))$, where p_k is the particle placed in the accelerator and $p_{m_1}, \ldots, p_{m_n}$ are the n fission products resulting from the split of p_k. If the split fails, the message $(p_k, ())$ is returned. Assume as before that the lab's data is true and that the lab's technique eventually discovers every physically possible split. For

Figure 3.27

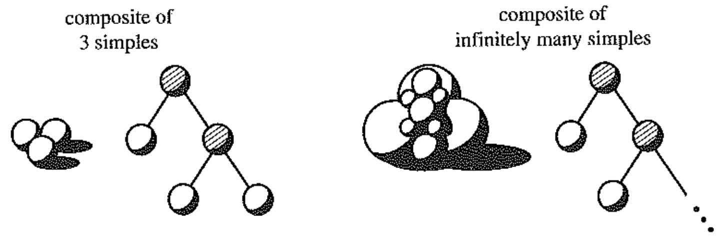

Figure 3.28

each of the following hypotheses, find the strongest sense in which its truth can be reliably determined. This requires showing that the problem is solvable in a given sense and that the demon has a winning strategy if the problem is to be solved in any stronger sense.

a. m is everywhere densely divisible.
b. m is a raisin pudding.
c. m is either a composite of 1 simple or a composite of 3 simples.
d. m is composed of simples.
e. m is a finite composite of simples.

3.2. Prove proposition 3.12.

3.3. In what sense can we reliably investigate $\mathcal{PS}$-randomness given that $\mathcal{PS}$ is countable and the limiting relative frequency of each datum occurring in the data stream exists?

*3.4. Suppose you have a system involving a car on a track subject to unknown forces (*Fig. 3.30*).

Consider the following hypotheses:

h_1: *the position of the car at t is x.*
h_2: *the velocity of the car at t is v.*
h_3: *the acceleration of the car at t is a.*

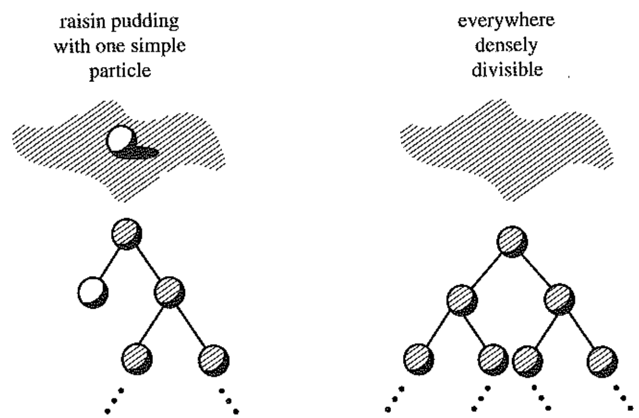

Figure 3.29

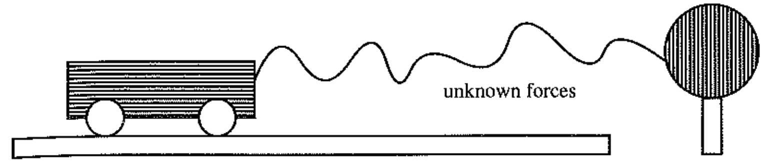

Figure 3.30

h'_1: *the trajectory of the car through time is $f(t)$.*
h'_2: *the velocity of the car through time is $v(t)$.*
h'_3: *the acceleration of the car through time is $a(t)$.*

Consider the following bodies of background knowledge:

$\mathcal{K}_1$: *f is any function from reals to reals.*
$\mathcal{K}_2$: *f is continuous.*
$\mathcal{K}_3$: *f is differentiable.*
$\mathcal{K}_4$: *f is twice differentiable.*

Now consider the following data protocols:

P_1: We can measure time and position exactly and continuously so that by time t we have observed $f|t = \{(x, y): f(x) = y \text{ and } x \leq t\}$.

P_2: We can have the lab repeat the motion exactly at will and set a super camera to record the exact position of the car at a given time. We can only run finitely many trials in a given interval of time, however.

P_3: Like P_2 except that the camera technology is limited so that a given camera can only determine position to within some interval Δx. The lab can improve the camera at will, however, to decrease the interval to any given size greater than 0.

P_4: We must now use a fixed camera that has no guaranteed accuracy. However, we know that if the car is in position x at time t, then an infinite sequence of repeated trials of the experiment will have limiting relative frequencies of measurements at t satisfying a normal distribution with mean x and variance v.

Provide a logical reliabilist analysis of the various inductive problems that result from different choices of hypothesis, background knowledge, and protocol.