1

Introduction: Radical probabilism

Adopting a central feature of Stoic epistemology, Descartes treated belief as action that might be undertaken wisely or rashly, and enunciated a method for avoiding false belief, a discipline of the will "to include nothing more in my judgments than what presented itself to my mind with such clarity and distinctness that I would have no occasion to put it in doubt."[1] He called such acts of the will "affirmations," i.e., acts of accepting sentences or propositions as true. (Essay 2 argues against "cognitive" uses of decision theory to choose among such replacements of considered probabilities by specious certainties.)

What do "belief," "acceptance," and "affirmation" mean in this context? I don't know. I'm inclined to doubt that anyone else does, either, and to explain the general unconcern about this lack of understanding by familiarity of the acceptance metaphor masquerading as intelligibility, perhaps as follows: "Since it's clear enough what's meant by accepting other things – gifts, advice, apologies – and it's clear enough what's meant by sentences' being true, isn't it clear what's meant by accepting sentences as true? Doesn't Quine make 'holding' sentences true the very pivot of his epistemology? And isn't affirmation just a matter of saying 'Yes'?"

The ("Bayesian") framework explored in these essays replaces the two Cartesian options, affirmation and denial, by a continuum of judgmental probabilities in the interval from 0 to 1, endpoints included, or – what comes to the same thing – a continuum of judgmental odds in the interval from 0 to ∞, endpoints included. Zero and 1 are probabilities no less than 1/2 and 99/100 are. Probability 1 corresponds to infinite odds, 1 : 0. That's a reason for thinking in terms of odds: to remember how momentous it may be to assign probability 1 to a hypothesis. It means you'd stake your all on its

1. *Discourse on the method . . .* , part 2.

truth, if it's the sort of hypothesis you can stake things on. To assign 100% probability to success of an undertaking is to think it advantageous to stake your life upon it in exchange for any petty benefit. We forget that when we imagine that we'd assign probability 1 to whatever we'd simply state as true.[2]

What is involved in attributing particular judgmental probabilities to sentences? Essays 13 and 14 answer in terms of a theory of preference seen as a relation between sentences or propositions: preference for truth of one sentence ("Cameroon wins") to truth of another ("Britain wins"). This theory is subjectivistic in addressing only the effects of such probability judgments, without saying how those judgments ought to be arrived at. The theory doesn't prejudge attempts like Carnap's to supply norms for forming such judgments; and indeed Carnap accepted this subjectivistic theory as an account of how judgmental probabilities are to be applied, once formed.[3]

Broadly speaking, a Bayesian is a probabilist, a person who sees making up the mind as a matter of either adopting an assignment of judgmental probabilities or adopting certain features of such an assignment, e.g., the feature of assigning higher conditional probability to 5-year survival on a diagnosis of ductal cell carcinoma than on a diagnosis of islet cell carcinoma. Some insist on restricting the term "Bayesian" narrowly to those who see conditioning (or "conditionalization") as the only rational way to change the mind; I don't. (Essays 3, 4, 5.) Rationalistic Bayesianism – hereafter, "rationalism" – is a subspecies of the narrow Bayesianism just noted, according to which there exists a (*logical, a priori*) probability distribution that would define the state of mind of a perfect intelligence, innocent of all experience. Notable subscribers: Bayes,

2. As probabilities p range over the unit interval [0,1], the corresponding odds $o = p/(1-p)$ range over the extended nonnegative reals $[0,\infty]$, enhancing resolution high in the scale. Thus, probabilities 99%, 99.9%, 99.99% correspond to odds 99, 999, 9999. But at the low end, where odds $p/(1-p)$ are practically equal to probabilities p, there is no such increase in resolution. Logarithms of odds, increasingly positive above $p = .5$ and symmetrically negative below, yield the same resolution at both ends. (Odds of 99, 999, 9999 become log odds of approximately 2, 3, 4, and at the low end, probabilities of .01, .001 .0001 become log odds of approximately -2, -3, -4.)
3. In his introduction to F. P. Ramsey's *Philosophical Papers* (Cambridge University Press, 1990, p. xviii) D. H. Mellor, who ought to know better, says that "modern Bayesian decision theory (e.g., Jeffrey's *The Logic of Decision* 1965), tells us to 'act in the way we think most likely to realize the objects of our desires' whether or not those thoughts and desires are either reasonable or right."

Laplace, W. E. Johnson, J. M. Keynes, Carnap, John Harsanyi.

Rationalism and empiricism are two sides of the same Bayesian coin. One side is a purely rational, "logical" element, a prior probability assignment **M** characterizing the state of mind of a newborn Laplacean intelligence. Carnap spent his last 25 years trying to specify **M**. The other side is a purely empirical element, a comprehensive report D of all experience to date. Together, these determine the experienced Laplacean intelligence's judgmental probabilities, obtained by conditioning the "ignorance prior" **M** by the *Protokollsatz* D. Thus $\mathbf{M}(H|D)$ is the correct probabilistic judgment about H for anyone whose experiential data base is D.

Radical probabilism makes no attempt to analyze judgment into a purely rational component and a purely empirical component, without residue. It rejects the empiricist myth of the sensuously given data proposition D as well as the rationalist myth of the ignorance prior **M**; it rejects the picture of judgment as a coin with empirical obverse and rational reverse. Let's see why.

On the empirical side, reports of conscious experience are too thin an abstract of our sensory inputs to serve adequately as the first term of the equation

(1) $$\textit{Experience} + \textit{reason} = \textit{judgment}.$$

Example: Blindsight.[4] In humans, monkeys, etc., some 90% of optic nerve fibers project to the striate cortex at the very back of the brain via the dorsal lateral geniculate nucleus in the midbrain. But

> while the geniculo-striate pathway constitutes the major portion of the optic nerve . . . there are at least 6 other branches that end up in the midbrain and subcortical regions . . . , and one of these contains about 100 000 fibres, by no means a trivial pathway – it is larger than the whole of the auditory nerve. . . . Therefore, if striate cortex is removed or its direct input blockaded, one should expect that some visual capacity should remain because all of those non-geniculo-striate pathways are left intact. The paradox is that in man this is usually not so: destruction of occipital cortex . . . characteristically causes blindness in that part of the visual field corresponding to the precise projection map of the retina on to the striate cortex. . . . Admittedly, some primitive visual reflexes can be de-

4. *Blindsight, a Case Study and Implications,* by L. Weiskrantz, Clarendon Press, Oxford, 1986, pp. 3–6, 24, and 168–169. See also Patricia Smith Churchland, *Neurophilosophy*, MIT Press, Cambridge, Mass., 1986, pp. 224–228.

tected . . . but typically the patient himself does not appear to discriminate visual events or to have any awareness of them.

The non-geniculo-striate 10% of optic nerve fibres seem to provide visual capacities of which such patients are unaware – capacities which they dismiss as mere guesswork, and which the rest of us need never distinguish as a special category. Thus, although a patient ("D. B.") whose right occipital lobe had been surgically removed "could not see one's outstretched hand, he seemed to be able to reach for it accurately. We put movable markers on the wall to the left of his fixation, and again he seemed to be able to point to them, although he said he did not actually see them. Similarly, when a stick was held up in his blind field either in a horizontal or vertical position, and he was asked to guess which of these two orientations it assumed, he seemed to have no difficulty at all, even though again he said he could not actually see the stick." This sort of thing looks like bad news for the New Way of Ideas, empiricism based on sense data: D. B. has factual sensitivity, a basis for probabilistic judgment, with no corresponding phenomenal sensitivity.

If sense data won't do for the first term of formula (1), perhaps the triggering of sensory receptors will. Quine seems to think so: "By the stimulation undergone by a subject on a given occasion I just mean the temporally ordered set of all those of his exteroceptors that are triggered on that occasion."[5] The experiential data base D might then correspond to an ordered set of such ordered sets, whence the Carnapian judgmental probability $\mathbf{M}(H|D)$ might be calculable. But no; not even "in principle." The trouble is that the temporal record of exteroception makes no perceptual sense without a correlated record of interoception; thus, interpretation of a record of activity in the optic nerve requires a correlated record of relative orientations of eye, head, trunk, etc. When Quine's bit of neurophysiology is put in context, his exteroceptive data base looks no more adequate for its purpose than did the sense data base it replaced.

Example: Proprioception and visual perception.[6] Exteroceptive nervous activity is interpreted in the light of concurrent interoceptive

5. W. V. Quine, *The Pursuit of Truth*, Harvard University Press, Cambridge, Mass., 1990, p. 2.
6. J. Allen Hobson, *The Dreaming Brain*, Basic Books, New York, 1988, pp. 110–112.

activity. Thus, optic nerve input is interpreted in the light of concurrent activity in oculomotor brain-stem neurons sending axons directly to the eye muscles, whose activity is coordinated by interactions of nuclei commanding vertical, oblique, and lateral determinants of gaze. The vestibular neurons in the brain stem relay information about head position from inner ear to oculomotor neurons. . . . Head and eye position are related, in turn, to spinal control of posture by the reticular formation. "Without the constant and precise operation of these three systems, we could neither walk and see, nor sit still and read. . . . Together with the cerebellum, the integrated activity of these brain-stem systems is responsible for giving sighted animals complex control of their acts." Quite apart from the question of awareness, it seems that the neurological analog of sense data must go beyond irritations of sensory surfaces. In the Cartesian mode it must treat the observer's body as a part of the "external" world providing the mind with inputs to be coordinated with exteroceptive inputs by innate neurological circuitry that is fine-tuned mostly in utero and in the earliest years of extrauterine life.

From Carnap to Quine, it is ordinary thing-languages to which physicalists have looked for observation sentences, whose imputed truth values (or probability values) are to be propagated through the confirmational net by conditioning (or generalized conditioning). Quine gestures toward temporally ordered sets of triggered exteroceptors as an empirical substrate for the real epistemological action, Cartesian affirmations of ordinary observation sentences. But the proffered substrate, once mentioned, plays no further rôle in Quine's epistemology. It is anyway incapable of providing an empirical footing for his holdings true until enriched by a coördinated efferent substrate. The full-blown afferent-efferent substrate would provide a footing ("neurological solipsism") upon which holdings true and holdings probable to various degrees could supervene, but it would play no rôle, either. Bag it.

So much for the empirical side of the epistemological coin. On the other side, radical probabilism abandons Carnap's search for the fountain of rationality in a perfect ignorance prior, at the same time abandoning the idea that conditioning, or generalized conditioning, is the canonical way to change your mind. Instead, radical probabilism offers a dynamic or synchronic point of view, from which the distinction between making up your mind and changing it becomes tenuous. The Carnapian motion picture is a sequence of instantaneous frames, your successive complete probability assignments to all sentences of your language, beginning with **M** and changing every time a new conjunct is added to your data base: **M**(—), **M**(—$|D_1$), **M**(—$|D_1$&D_2), and so on up to your present assignment, **M**(—$|D_1$&D_2& . . . &D_t). The radical probabilist picture is less detailed in each frame, and smoother or more structural across frames in the time dimension, more like a Minkowski diagram.

Thus, making up your mind probabilistically involves making up your mind about how you will change your mind. It's not that you must map that out in fine detail, any more than you must map out your instantaneous probabilities for all sentences of your language, frame by frame. But since it no longer goes without saying that you will change your mind by conditioning or generalized conditioning (probability kinematics) any more than it goes without saying that your changes of mind will be quite spontaneous or unconsidered, these are questions about which you may make up your mind about changing your mind in specific cases. You may decide to change your mind by generalized conditioning on some set of data propositions. According to the laws of probability logic ("the probability calculus") such a decision comes to the same thing as deciding to keep your conditional probabilities on the data propositions constant when your unconditional probabilities for them change. In case your probability for one of the data propositions changes to 1, this reduces to ordinary conditioning on the data proposition you've become sure of. (Essays 3, 6, 7.)

It needs to be emphasized that becoming sure of a sentence's truth doesn't guarantee that your new conditional probabilities based on it will be the same as they were before you became sure of it. That's why Carnap required that you condition only on sentences that you regard not only as true but as recording the whole of the relevant truth that you know about. For this to imply constancy of conditional probabilities there must be available to you an infinitely nuanced assortment of data propositions to condition upon. It strikes me as a fantasy, an epistemologist's pipe-dream, to imagine that such nuanced propositions are generally accessible to us. There need be no sentence you can formulate, that fits the description "the whole of the relevant truth that you know about."[7] But the diachronic perspective of radical probabilism reveals a different dimension of nuance that you can actually use in such cases to identify a set of data propositions relative to which you expect your conditional probabilities to be unchanged by an impending observation that you think will have the effect of changing your probabilities for some of the data propositions. That will be a case where updating by probability kinematics is appropriate.

Constancy of conditional probabilities opens other options for registering and communicating the effect of experience, e.g., the option of registering the ratios ("Bayes factors") $f(A,B)$ of odds between A and B afterward and beforehand:

$$f(A,B) = \frac{Q(A)/Q(B)}{P(A)/P(B)} \;. \tag{2}$$

What's conveyed by the Bayes factor is just the effect of experience, final odds with prior odds factored out. Others who accept your response to your experience, whether or not they share your prior opinion, can multiply their own prior odds between A and B by your Bayes factor to get their posterior odds, taking account of your experience.[8]

Example: Expert opinion. Jane Doe is a histopathologist who hopes to settle on one of the following diagnoses on the basis of microscopic examination of a section of tissue surgically removed from a pancreatic tumor. (To simplify matters, suppose she is sure that exactly one of the three diagnoses is correct.)

7. Given an old probability distribution, P, and a new one, Q, it is an open question whether, among the sentences D in your language for which your conditional probabilities are the same relative to Q as they are relative to P, there are any for which your new probability $Q(D)$ is 1. If so, and only then, Q can be viewed as coming from P by conditioning.
8. Such uses of Bayes factors were promoted by I. J. Good in chapter 6 of his book *Probability and the Weighing of Evidence*, Charles Griffin, London, 1950. See also *Alan Turing: The Enigma*, by Andrew Hodges, Simon and Schuster, New York, 1983, pp. 196–7. Good promotes *logarithms* of odds ("plausibilities") and of Bayes factors ("weights of evidence") as intelligence amplifiers which played a rôle in cracking the German "enigma" code during the second world war.

A = Islet cell carcinoma

B = Ductal cell carcinoma

C = Benign tumor

In the event, the experience does not drive her probability for any diagnosis to 1, but does change her probabilities for the three candidates from the following values (P) prior to the experience, to new values (Q):

	A	B	C
P	1/2	1/4	1/4
Q	1/3	1/6	1/2

Henry Roe, a clinician, accepts the pathologist's findings, i.e., he adopts, as his own, her Bayes factors between each diagnosis and some fixed hypothesis, say, C[9]:

$$f(A,C) = 1/3, \qquad f(B,C) = 1/3, \qquad f(C,C) = 1.$$

It is to be expected that, *given a definite diagnosis,* his conditional probabilities for the prognoses "live" (for 5 years) and "die" (within 5 years) are stable, unaffected by the pathologist's report. For definiteness, suppose those stable probabilities are as follows, where lowercase "p" and "q" are used for the clinician's prior and posterior probabilities, to distinguish them from the pathologist's.

$$q(\text{live} \mid D) = p(\text{live} \mid D) = .4,\ .6,\ .9 \text{ when } D = A,\ B,\ C$$

$$q(\text{die} \mid D) = p(\text{die} \mid D) = .6,\ .4,\ .1 \text{ when } D = A,\ B,\ C$$

Given his prior probabilities $p(D)$ for the diagnoses and his adopted Bayes factors, these conditional probabilities determine his new odds on 5-year survival.[10] It works out as follows.

9. Proposed by Schwartz, W. B., Wolfe, H. J., and Pauker, S. G., "Pathology and probabilities: a new approach to interpreting and reporting biopsies," *The New England Journal of Medicine* **305**(1981)917–923.

10. See Richard Jeffrey and Michael Hendrickson, "Probabilizing pathology," *Proceedings of the Aristotelian Society,* v. 89, part 3, 1988/89: p. 217, odds kinematics.

(3) $$\frac{q(\text{live})}{q(\text{die})} = \frac{p(\text{live} \mid A)p(A)f(A,C) + p(\text{live} \mid B)p(B)f(B,C) + p(\text{live} \mid C)p(C)}{p(\text{die} \mid A)p(A)f(A,C) + p(\text{die} \mid B)p(B)f(B,C) + p(\text{die} \mid C)p(C)}$$

If the clinician's prior distribution of probabilities over the three diagnoses was flat, $p(D) = 1/3$ for each diagnosis, then the imagined numbers given above raise his new odds on 5-year survival from $p(\text{live}) : p(\text{die}) = 19 : 11$ to $q(\text{live}) : q(\text{die}) = 37 : 13$, so that his probability for 5-year survival rises from 63% to 74%.

Prima facie, the task of eliciting Bayes factors looks more difficult than eliciting odds, for Bayes factors are ratios of odds.[11] For the same reason it may seem that the pathologist's Bayes factor, (posterior odds) : (prior odds), cannot be elicited if (as it may well be) she has no definite prior odds. But if her Bayes factors would be stable over a large range of prior odds, so as to be acceptable by colleagues with various prior odds, her Bayes factors are as easily elicited as her posterior odds if only she is willing and able to *adopt* definite odds prior to her observation. Thus, if she adopts even prior odds, $P(A)/P(C) = P(B)/P(C) = 1$, her Bayes factors will simply be equal to her posterior odds. But if it is only uneven priors straightforwardly based on statistical abstracts of past cases that are cogent for her, the extra arithmetic presents no problem.

The example illustrates two contrasts between the radical probabilism advocated here and the phenomenalism I have been deprecating. The less important contrast concerns the distinction between probability and certainty as basic attitudes toward *Protokollsätze*. The more important contrast concerns the status of those attitudes toward *Protokollsätze* (or toward what they report) as foundations for all of our knowledge. Here, C. I. Lewis wears the Cartesian black hat better than Carnap: "Subtract, in what we say that we see, or hear, or otherwise learn from direct experience, *all that could conceivably be mistaken;* the remainder is the given content of the experience inducing this belief. If there were no such hard kernel in experience – e.g., what we see when we think we see a deer but

11. In general, elicitation is a process of *drawing forth*. Here, authenticity does not require the elicited Bayes factors to have been present in the pathologist's mind before the process began; the process may well be one in which she is induced to form a judgment, *making up* her mind probabilistically.

there is no deer – then the word 'experience' would have nothing to refer to."[12]

This is the sort of empiricism dismissed above, in which the term "experience" is understood not in its ordinary sense, as the sort of thing that makes you an experienced doctor, sailor, lover, traveller, carpenter, teacher, or whatever, but in a new sense, the sensuously given, in which experience is bare phenomenology or bare irritation of sensitive surfaces. It presupposes a unitary faculty of reason, the same for all subject matter, which, added to the sensuously given, equals good judgment. The formula itself goes back much further than Descartes, e.g., to Galen: "When I take as my standard the opinion held by the most skillful and wisest physicians and the best philosophers of the past, I say: The art of healing was originally invented and discovered by the logos [reason] in conjunction with experience. And to-day also it can only be practiced excellently and done well by one who employs both of these methods."[13]

But in this formula reason is theory, and experience is gained by purges, surgery, etc., the sort of thing Hippocrates had called dire in his famous "experiment perilous" aphorism. For experience in Galen's formula, C. I. Lewis substitutes the given. Galen's formula is

Experience + reason = medical expertise.

There's a similar formula for other kinds of knowledge and technique, with "reason" and "experience" referring to other things than they do in the case of medicine.[14] But Lewis's formula is general:

The given + reason = good judgment.

Here "reason" needs to be understood as something like a successful outcome of the project to which Carnap devoted his last 25 years, of designing a satisfactory general inductive logic. For that I have no hope, for reasons given above under the headings of "blindsight" and "perception and proprioception."

12. *An Analysis of Knowledge and Valuation,* Open Court, La Salle, Illinois, 1946, pp. 182–183. (Lewis's emphasis.)
13. The first sentences of "On medical experience," translated by Richard Walzer, in *Three Treatises on the Nature of Science,* Hackett, Indianapolis, 1985, p. 49.
14. Like "experience," "reason" has a different sense (comprehending *theory*) in Galen's formula from what it has in C. I. Lewis's; see pp. xx–xxxi of Michael Frede's introduction to the Galen book (note 10).