Sins of the Epistemic Probabilist
Exchanges with Peter Achinstein

Deborah G. Mayo

1 Achinstein's Sins

As Achinstein notes, he and I agree on several key requirements for an adequate account of evidence: it should be objective and not subjective, it should include considerations of flaws in data, and it should be empirical and not a priori. Where we differ, at least at the moment, concerns the role of probability in inductive inference. He takes the position that probability is necessary for assigning degrees of objective warrant or "rational" belief to hypotheses, whereas, for the error statistician, probability arises to characterize how well probed or tested hypotheses are (highly probable vs. highly probed). Questions to be considered are the following:

1. *Experimental Reasoning*: How should probability enter into inductive inference – by assigning degrees of belief or by characterizing the reliability of test procedures?
2. *Objectivity and Rationality*: How should degrees of objective warrant be assigned to scientific hypotheses?
3. *Metaphilosophy*: What influences do philosophy-laden assumptions have on interpretations of historical episodes: Mill?

I allude to our disagreeing "at the moment" because I hope that the product of this latest installment in our lengthy exchange may convince him to shift his stance, at least slightly.

In making his case, Achinstein calls upon Mill (also Newton, but I largely focus on Mill). Achinstein wishes "to give Mill a better run for his money" by portraying him as an epistemic probabilist of the sort he endorses. His question is whether Mill's inductive account contains some features "that are or seem inimical to [the] error-statistical philosophy." A question for us is whether Achinstein commits some fundamental errors, or "sins" in his terminology, that are common to epistemic (Bayesian) probabilists in philosophy of science, and are or seem to be inimical to the spirit of an "objective" epistemic account. What are these sins?

1. *First*, there are *sins of "confirmation bias"* or "reading one's preferred view into the accounts of historical figures" or found in scientific episodes, even where this reconstruction is at odds with the apparent historical evidence.
2. *Second*, there are *sins of omission*, whereby we are neither told how to obtain objective epistemic probabilities nor given criteria to evaluate the success of proposed assignments.

These shortcomings have undermined the Bayesian philosophers' attempts to elucidate inductive inference; considering Achinstein's discussion from this perspective is thus of general relevance for the popular contemporary project of Bayesian epistemologists.

1.1 Mill's Innocence: Mill as Error Statistician

Before considering whether to absolve Mill from the sins Achinstein considers, we should note that Achinstein omits the one sin discussed in Mayo (1996) that arises most directly in association with Mill: Mill denies that novel predictions count more than nonnovel ones (pp. 252, 255). My reference to Mill comes from Musgrave's seminal paper on the issue of novel prediction:

> According to modern logical empiricist orthodoxy, in deciding whether hypothesis *h* is confirmed by evidence *e*... we must consider only the statements *h* and *e*, and the logical relations between them. It is quite irrelevant whether *e* was known first and *h* proposed to explain it, or whether *e* resulted from testing predictions drawn from *h*. (Musgrave, 1974, p. 2)

"We find" some variant of the logicist approach, Musgrave tells us, "in Mill, who was amazed at Whewell's view" that successfully predicting novel facts gives a hypothesis special weight. "In Mill's view, 'such predictions and their fulfillment are... well calculated to impress the uninformed' " (Mill, 1888, Book 3, p. 356). This logicism puts Mill's straight rule account at odds with the goal of severity.

Nevertheless, Achinstein gives evidence that Mill's conception is sufficiently in sync with error statistics so that he is prepared to say, "I would take Mill to be espousing at least some version of the idea of 'severe testing'" (p. 174). Here is why: Although Mill regards the form of induction to be "induction by simple enumeration," whether any particular instantiation is valid also "depends on nonformal empirical facts regarding the sample, the sampling, and the properties being generalized" (Achinstein, p. 174) – for example, Mill would eschew inadequately varied data, failure to look for negative instances, and generalizing beyond what the evidence allows. These features could well set the stage for expiating the sins of the straight ruler: The "fit" requirement, the first clause of severity, is met because observed cases of As that are Bs "fit" the hypothesis that all As are Bs; and the additional considerations Achinstein lists would help to satisfy severity's second condition – that the ways such a fit could occur despite the generalization being false have been checked and found absent.

Achinstein finds additional intriguing evidence of Mill's error-statistical leanings. Neither Mill's inductions (nor those of Newton), Achinstein points out, are "guilty" of assigning probabilities to an inductive conclusion:

> Neither in their abstract formulations of inductive generalizations . . . nor in their examples of particular inductions . . . does the term "probability" occur. . . . From the inductive premises we simply conclude that the generalization is true, or as Newton allows in rule 4, "very nearly true," by which he appears to mean not "probably true" but "approximately true." (p. 176)

Neither exemplar, apparently, requires the assignment of a probability to inductively inferred generalizations. One might have expected Achinstein to take this as at least casting doubt on his insistence that such posterior probabilities are necessary. He does not.

1.2 Mill's Guilt: Mill as Epistemic Probabilist

Instead, Achinstein declares that their (apparent) avoidance of any "probabilistic sin" is actually a transgression to be expiated, because it conflicts with his own account! "However, before we conclude that no probabilistic sin has been committed, we ought to look a little more closely at Mill." In particular, Achinstein thinks we ought to look at contexts where Mill *does* talk about probabilities – namely in speaking about events – and substitute what he says there to imagine he is talking about probabilities of hypotheses. By assigning to Mill what is a probabilistic sin to error statisticians, Mill is restored to good graces with Achinstein's Bayesian probabilist.

Mill offers examples from games of chance, Achinstein notes (p. 177), where probabilities of general outcomes are assigned on the basis of their observed relative frequencies of occurrence:

> In the cast of a die, the probability of ace is one-sixth . . . because we do actually know, either by reasoning or by experience, that in a hundred or a million of throws, ace is thrown about one-sixth of that number, or once in six times. (Mill, 1888, p. 355)

In more modern parlance, Mill's claim is that we may accept or infer ("by reasoning or by experience") a statistical hypothesis H that assigns probabilities to outcomes such as "ace." For the error statistician, this inductive inference takes place by passing *H* severely or with appropriately low error probabilities. Although there is an inference to a probabilistic model of experiment, which in turn assigns probabilities to outcomes, there is no probabilistic assignment to *H* itself, nor does there seem to be in Mill. So such statements from Mill do not help Achinstein show that Mill intends to assign posterior probabilities to hypotheses.[1]

The frequentist statistician has no trouble agreeing with the conditional probabilities of events stipulated by Achinstein. For example, we can assert the probability of spades given the outcome "ace" calculated under hypothesis *H* that cards are randomly selected from a normal deck, which we may write as P(spade | ace; *H*) – noting the distinction between the use of "|" and ";". However, we would not say that an event severely passes, or that one event severely passes another event. A statistical hypothesis must assign probabilities to outcomes, whereas the event "being an ace" does not.

I claim no Mill expertise, yet from Achinstein's own presentation, Mill distinguishes the assignment of probabilities to events from assigning probabilities to hypotheses, much as the error statistician. What then is Achinstein's justification for forcing Mill into a position that Mill apparently rejects? Achinstein seems guilty of the sin of "reading our preferred view into the accounts of historical episodes and/or figures."

2 The Error-Statistical Critique

2.1 Some Sleight-of-Hand Sins

Having converted Mill so that he speaks like a Bayesian, Achinstein turns to the business of critically evaluating the error statistician: "[O]bjective epistemic probabilists – again I include Mill – are committed to saying that the inference is justified only if the objective epistemic probability (the 'posterior' probability) of the inductive conclusion" is sufficiently high (p. 179). At times during our exchanges, I thought Achinstein meant this assertion as a tautology, playing on the equivocal uses of "probability" in ordinary language. If high epistemic probability is just shorthand for high inductive

[1] On the other hand, they do raise the question of how Achinstein's epistemic probabilist can come to accept a statistical model (on which probabilistic assignments to events are based). Achinstein's formal examples start out by assuming that we have accepted a statistical model of experiment, *H*, usually random sampling from a binomial distribution. Is this acceptance of *H* itself to be a matter of assigning *H* a high posterior probability through a Bayesian calculation? The alternatives would have to include all the ways the model could fail. If not, then Achinstein is inconsistent in claiming that warranting *H* must take the form of a Bayesian probability computation.

warrant, then his assertion may be uncontroversial, if also uninformative. If hypothesis H passes a highly severe test, there is no problem in saying there is a high degree of warrant for H. But there is a problem if this is to be cashed out as a posterior probability, attained from Bayes's theorem and satisfying the probability axioms. Nevertheless, Achinstein assures me that this is how he intends to cash out epistemic probability. High epistemic probability is to be understood as high posterior probability (obtainable from Bayes's theorem). I construe it that way in what follows.

Achinstein presupposes that if we speak of data warranting an inference or belief in H, then we must be talking about an epistemic posterior probability in H, but this is false. Musgrave's critical rationalist would reject this as yet another variant on "justificationism" (see Chapter 3). Achinstein's own heroes (Mill and Newton) attest to its falsity, because as he has convincingly shown, they speak of warranting hypotheses without assigning them posterior probabilities.

In the realm of formal statistical accounts of induction, it is not only the error statistician who is prepared to warrant hypotheses while eschewing the Bayesian algorithm. The "likelihoodist," for instance, might hold that data $\mathbf{x}$ warrants H to the extent that H is more likely (given $\mathbf{x}$) than rivals. High likelihood for H means $P(\mathbf{x};H)$ is high, but high likelihood for H does *not* mean high probability for H. To equate $P(\mathbf{x}|H)$ and $P(H|\mathbf{x})$ immediately leads to contradictions – often called the prosecutor's fallacy. For example, the likelihood of H and not-H do not sum to 1 (see Chapters 7 and 9, this volume).

If Achinstein were to accept this, then we would be in agreement that probabilistic concepts (including frequentist ones) may be used to qualify the evidential warrant for hypothesis H while denying that this is given by a probability assignment to H. Achinstein may grant these other uses of probability for induction yet hold that they are inferior to his idea of objective epistemic posterior probabilities. His position might be that, unless the warrant is a posterior probability, then it does not provide an adequate objective inductive account. This is how I construe his position.

2.2 Sins of Omission: How Do We Apply the Method? What Is So Good About It?

For such a sweeping claim to have any substance, it must be backed up with (1) some guidance as to how we are to arrive at objective epistemic posteriors and (2) an indication of (desirable) inferential criteria that objective epistemic probability accounts satisfy.

So how do we get to Achinstein's objective epistemic posterior probabilities? By and large this is not one of Achinstein's concerns; he considers informal examples where intuitively good evidence exists for a claim *H*, and then he captures this by assigning *H* a high degree of epistemic probability. If he remained at the informal level, the disagreements he finds between us would likely vanish, for then high probability could be seen as a shorthand for high inductive warrant, the latter attained by a severe test. When he does give a probabilist computation, he runs into trouble.

Unlike the standard Bayesian, Achinstein does not claim that high posterior probability is sufficient for warrant or evidence, but does "take [it] to be a necessary condition" (p. 180). It is not sufficient for him because he requires, in addition, what he describes as a non-Bayesian "explanatory connection" between data and hypotheses (p. 183). Ignoring the problem of how to determine the explanatory connection, and whether it demands its own account of inference (rendering his effort circular), let us ask: Are we bound to accept the necessity? My question, echoing Achinstein, is this: *Epistemically speaking, what, if anything, can one conclude about hypothesis H itself from the fact that Achinstein's method accords H a high objective epistemic probability?*

That is, we require some indication that *H*'s earning high (low) marks on Achinstein's objective probability scale corresponds to actually having strong (weak) evidence of the truth of *H*. If Achinstein could show this, his account would be superior to existing accounts of epistemic probability! It will not do to consider examples where strong evidence for *H* is intuitively sound and then to say we have high Achinstein posterior epistemic probability in *H*; he would have to demonstrate it with posteriors arrived at through Achinsteinian means.

3 Achinstein's Straight Rule for Attaining Epistemic Probabilities

To allow maximum latitude for Achinstein to make his case, we agree to consider the Achinsteinian example. The most clear-cut examples instantiate a version of a "straight rule," where we are to consider a "hypothesis" that consists of asserting that a sample possesses a characteristic such as "having a disease" or "being college-ready." He is led to this peculiar notion of a hypothesis because he needs to use techniques for probability assignments appropriate only for events. But let us grant all of the premises for Achinstein's examples.[2] We have (see p. 187):

[2] In an earlier exchange with Colin Howson, I discuss converting these inadmissible hypotheses into legitimate statistical ones to aid the criticism as much as possible, and I assume the same here (Mayo, 1997).

Achinstein's Straight Rule for Objective Epistemic Probabilities: If (we know only that) a_i is randomly selected from a population where p% have property C, then the objective epistemic probability that a_i has C equals p.

Next, Achinstein will follow a variant on a Bayesian gambit designed to show that data **x** from test *T* can pass hypothesis *H* severely, even though *H* is accorded a low posterior probability – namely by assuming the prior probability of *H* is sufficiently low. (Such examples were posed post-EGEK, first by Colin Howson (1997a).) I have argued that in any such example it is the Bayesian posterior, and not the severity assignment, that is indicted (Mayo 2005, 2006). This criticism and several like it have been discussed elsewhere, most relevantly in a collection discussing Achinstein's account of evidence (Mayo, 2005)! I focus mainly on the update to our earlier exchange, but first a quick review.

3.1 College Readiness

We may use Achinstein's summary of one of the canonical examples (pp. 186–7). We are to imagine a student Isaac who has taken the test and achieved a high score X, which is very rarely achieved by those who are not college-ready. Let

$$H(\mathrm{I})\text{: Isaac is college-ready.}$$

Let H' be the denial of H:

$$H'(\mathrm{I}) \text{ Isaac is not college-ready, i.e., he is deficient.}$$

I use $H(\mathrm{I})$ to emphasize that the claim is about Isaac}.

Let S abbreviate: Isaac gets a high score (as high as X).

In the error statistical account, S is evidence against Isaac's deficiency, and for $H(\mathrm{I})$. (Although we would consider degrees of readiness, I allow his dichotomy for the sake of the example.) Now we are to suppose that the probability that Isaac would get those test results, given that he is college-ready, is extremely high, so that

$$P(S|H(\mathrm{I})) \text{ is practically } 1,$$

Whereas the probability that Isaac would get those test results, given that he is not college-ready, is very low, say,

$$P(S|H'(\mathrm{I})) = .05.$$

Suppose that Isaac was randomly selected from a population – call it Fewready Town – in which college readiness is extremely rare, say one out of one thousand. The critic infers that

$$(*)\ \mathrm{P}(H(\mathrm{I})) = .001.$$

If so, then the posterior probability that Isaac is college-ready, given his high test results, would be very low; that is,

$$\mathrm{P}(H(\mathrm{I})\,|\,S) \text{ is very low,}$$

even though in this case the posterior probability has increased from the prior in (*).

The critic – for example, Achinstein – regards the conclusion as problematic for the severity account because, or so the critic assumes, the frequentist would also accept (*) $\mathrm{P}(H) = .001$. Here is the critic's flaw. Although the probability of college readiness in a randomly selected student from high schoolers from Fewready Town is .001, it does not follow that Isaac, the one we happened to select, has a probability of .001 of being college-ready (Mayo, 1997a, 2005, p. 117). To suppose it does is to commit what may be called a fallacy of probabilistic instantiation.

3.2 Fallacy of Probabilistic Instantiation

We may abbreviate by $\mathrm{P}(H(x))$ the probability that a randomly selected member of Fewready has the property C. Then the probabilistic instantiation argues from the first two premises,

$$\mathrm{P}(H(x)) = .001.$$

The randomly selected student is I.

To the inference:

$$(*)\ \mathrm{P}(H(\mathrm{I})) = .001.$$

This is fallacious. We need not preclude that $H(\mathrm{I})$ has a legitimate frequentist prior; the frequentist probability that Isaac is college-ready might refer to generic and environmental factors that determine the chance of his deficiency – although I do not have a clue how one might compute it. But this is not what the probability in (*) gives us. Now consider Achinstein's new update to our exchanges on Isaac.

3.3 Achinstein's Update

Achinstein now accepts that this assignment in (*) is a sin for a frequentist:

> [T]he probabilistic fallacy charge would be true if the probabilities in question were construed as relative frequencies. However . . . I am concerned with epistemic

probability. If all we know is that Isaac was chosen at random from a very disadvantaged population, very few of whose members are college ready, say one out of one thousand, then we would be justified in believing that it is very [improbable] that Isaac is college-ready. (Achinstein, p. 187)

Hence, (*) gives a legitimate objective epistemic frequentist prior.

Therefore, even confronted with Isaac's high test scores, Achinstein's probabilist is justified in denying that the scores are good evidence for $H(\mathrm{I})$. His high scores are instead grounds for believing $H'(\mathrm{I})$, that Isaac is not college-ready. It is given that the posterior for $H'(\mathrm{I})$ is high, and certainly an explanatory connection between the test score and readiness exists. Although the posterior probability of readiness has increased, thanks to his passing scores, for Achinstein this does not suffice to provide epistemic warrant for H (he rejects the common Bayesian distinction between increased support and confirmation). Unless the posterior reaches a threshold of a fairly high number, he claims, the evidence is "lousy." The example considers only two outcomes: reaching the high scores or not, i.e., S or $\sim S$. Clearly a lower grade gives even less evidence of readiness; that is, $\mathrm{P}(H'(\mathrm{I})|\sim S) > \mathrm{P}(H'(I)|S)$. Therefore, whether Isaac scored a high score or not, Achinstein's epistemic probabilist reports justified high belief that Isaac is not ready. The probability of Achinstein finding evidence of Isaac's readiness even if in fact he is ready (H is true) is low if not zero. Therefore, Achinstein's account violates what we have been calling the most minimal principle for evidence!

The *weak severity principle*: Data **x** fail to provide good evidence for the truth of H if the inferential procedure had very little chance of providing evidence against H, even if H is false.

Here the relevant H would be $H'(\mathrm{I})$ – Isaac is not ready.

If Achinstein allows that there is high objective epistemic probability for H even though the procedure used was practically guaranteed to arrive at such a high posterior despite H being false, then (to echo him again) the error statistician, and I presume most of us, would have a problem understanding what high epistemic probability has to do with something we regard as crucial in induction, namely ensuring that an inference to H is based on genuine evidence – on data that actually discriminate between the truth and falsity of H.

This fallacious argument also highlights the flaw in trying to glean reasons for epistemic belief by means of just any conception of "low frequency of error." If we declared "unready" for any member of Fewready, we would rarely be wrong, but in each case the "test" has failed to discriminate the particular student's readiness from his unreadiness. We can imagine a context where we are made to bet on the generic event – the next student randomly selected from the population has property C. But this is very different from having probed whether this student, Isaac, is ready or not – the job our test needs to perform for scientific inference.

I cannot resist turning Achinstein's needling of me back on him: Is Achinstein really prepared to claim there is high epistemic warrant for $H'(\mathrm{I})$ even though the procedure had little or no probability of producing evidence against $H'(\mathrm{I})$ and for $H(\mathrm{I})$ even if $H(\mathrm{I})$ is true? If he is, then the error statistician, and perhaps Mill, would have a hard time understanding what his concept has to do with giving an objective warrant for belief.

Let us take this example a bit further to explain my ironic allegation regarding "reverse discrimination." Suppose, after arriving at high belief in Issac's unreadiness, Achinstein receives a report of an error: in fact Isaac was selected randomly, not from Fewready Town, but from a population where college readiness is common, Fewdeficient Town. The same score now warrants Achinstein's assignment of a strong objective epistemic belief in Isaac's readiness (i.e., $H(\mathrm{I})$). A high school student from Fewready Town would need to have scored quite a bit higher on these same tests than one selected from Fewdeficient Town for his scores to be considered evidence of his readiness. So I find it surprising that Achinstein is content to allow this kind of instantiation to give an objective epistemic warrant.

3.4 The Case of General Hypotheses

When we move from hypotheses like "Isaac is college-ready" (which are really events) to generalizations – which Achinstein makes clear he regards as mandatory if an inductive account is not to be "puerile" – the difficulty for the epistemic probabilist becomes far worse if, like Achinstein, we are to obtain epistemic probabilities via his frequentist straight rule.

To take an example discussed earlier, we may infer with severity that the relativistic light deflection is within $\pm\,\varepsilon$ units from the GTR prediction, by finding that we fail to reject a null hypothesis with a powerful test. But how can a frequentist prior be assigned to such a null hypothesis to obtain the epistemic posterior?

The epistemic probabilist would seem right at home with a suggestion some Bayesians put forward – that we can apply a version of Achinstein's straight rule to hypotheses. We can imagine that the null hypothesis is

H_0: There are no increased risks (or benefits) associated with hormone replacement therapy (HRT) in women who have taken HRT for 10 years.

Suppose we are testing for discrepancies from zero in both positive and negative directions (Mayo, 2003). In particular, to construe the truth of a general hypothesis as a kind of "event," it is imagined that we sample randomly from a population of hypotheses, some proportion of which are assumed true. The proportion of these hypotheses that have been found to be true in the past serves as the prior epistemic probability for H_0.

Therefore, if H_0 has been randomly selected from a pool of null hypotheses, 50% of which are true, we have

$$(^*)\ \mathrm{P}(H_0) = .5.$$

Although (*) is fallacious for a frequentist, once again Achinstein condones it as licensing an objective epistemic probability. But which pool of hypotheses should we use? The percentages "initially true" will vary considerably, and each would license a distinct "objective epistemic" prior. Moreover, it is hard to see that we would ever know the proportion of true nulls rather than merely the proportion that have thus far not been rejected by other statistical tests!

The result is a kind of "innocence by association," wherein a given H_0, asserting no change in risk, gets the benefit of having been drawn from a pool of true or not-yet-rejected nulls, much as the member from Fewready Town is "deficient by association." Perhaps the tests have been insufficiently sensitive to detect risks of interest. Why should that be grounds for denying evidence of a genuine risk with respect to a treatment (e.g., HRT) that *does* show statistically significant risks?

To conclude, here is our answer to Achinstein's question:

Does Mayo allow that H may pass with high severity when the posterior probability of H is not high?

In no case would H pass severely when the grounds for warranting H are weak. But high posteriors need not correspond to high evidential warrant. Whether the priors come from frequencies or from "objective" Bayesian priors, there are claims that we would want to say had passed severely that do not get a high posterior (see Chapter 7, this volume). In fact, statistically significant results that we would regard as passing the nonnull hypothesis severely can show a decrease in probability from the prior (.5) to the posterior (see Mayo, 2003, 2005, 2006).

4 Some Futuristic Suggestions for Epistemic Probabilists

In the introductory chapter of this volume, we mentioned Achinstein's concession that "standard philosophical theories about evidence are (and ought to be) ignored by scientists" (2001, p. 3) because they view the question of whether data **x** provide evidence for H as a matter of purely logical computation, whereas whether data provide evidence for hypotheses is not an a priori but rather an empirical matter. He appears to take those failures to show that philosophers can best see their job as delineating the concepts of evidence that scientists seem to use, perhaps based on rational reconstructions of figures from the historical record.

Some may deny it is sinful that *epistemic probabilists* omit the task of how to obtain, interpret, and justify their objective epistemic probabilities, and claim I confuse the job of logic with that of methodology (Buldt 2000). Colin Howson, who had already denied that it was part of the job of Bayesian logic to supply the elements for his (subjective) Bayesian computation, declared in 1997 that he was moving away from philosophy of statistics to focus on an even purer brand of "inductive logic"; and he has clearly galvanized the large contemporary movement under the banner of "Bayesian epistemology" (see Glymour, Chapter 9, this volume). Two points: First, it is far from clear that Bayesian logics provide normative guidance about "rational" inference[3]; after all, error statistical inference embodies its own logic, and it would be good to explore which provides a better tool for understanding scientific reasoning and inductive evidence. Second, there is a host of new foundational problems (of logic and method) that have arisen in Bayesian statistical practice and in Bayesian-frequentist "unifications" in the past decade that are omitted in the Bayesian epistemological literature (see Chapter 7.2). It is to be hoped that days of atonement will soon be upon us.

References

Achinstein, P. (2001), *The Book of Evidence*, Oxford University Press, Oxford.

Buldt, B. (2000), "Inductive Logic and the Growth of Methodological Knowledge, Comment on Mayo," pp. 345–54 in M. Carrier, G. Massey, and L. Ruetsche (eds.), *Science at Century's End, Philosophical Questions on the Progress and Limits of Science*, University of Pittsburgh Press, Pittsburgh.

Howson, C. (1997a), "A Logic of Induction," *Philosophy of Science*, 64: 268–90.

Mayo, D.G. (1996), *Error and the Growth of Experimental Knowledge*, The University of Chicago Press, Chicago.

Mayo, D.G. (1997a), "Response to Howson and Laudan," *Philosophy of Science*, 64: 323–33.

Mayo, D.G. (2003), "Could Fisher, Jeffreys and Neyman Have Agreed on Testing? Commentary on J. Berger's Fisher Address," *Statistical Science*, 18: 19–24.

Mayo, D.G. (2005), "Evidence as Passing Severe Tests: Highly Probed vs. Highly Proved," pp. 95–127 in P. Achinstein (ed.), *Scientific Evidence*, Johns Hopkins University Press, Baltimore, MD.

[3] Even for the task of analytic epistemology, I suggest that philosophers investigate whether appealing to error-statistical logic offers a better tool for characterizing probabilistic knowledge than the Bayesian model. Given Achinstein's threshold view of evidential warrant, it is hard to see why he would object to using a severity assessment to provide the degree of epistemic warrant he seeks.

Mayo, D.G. (2006), "Philosophy of Statistics," pp. 802–15 in S. Sarkar and J. Pfeifer (eds.), *Philosophy of Science: An Encyclopedia*, Routledge, London.
Mill, J.S. (1888), *A System of Logic*, 8th edition, Harper and Bros., New York.
Musgrave, A. (1974), "Logical versus Historical Theories of Confirmation," *British Journal for the Philosophy of Science*, 25: 1–23.

Related Exchanges

Achinstein, P. (2000), "Why Philosophical Theories of Evidence Are (and Ought to Be) Ignored By Scientists, pp. S180–S192 in D. Howard (ed.), " *Philosophy of Science*, 67 (Symposia Proceedings).
Giere, R.N. (1997b), "Scientific Inference: Two Points of View," pp. S180–S184 in L. Darden (ed.), *Philosophy of Science*, 64 (PSA 1996: Symposia Proceedings).
Howson, C. (1997b), "Error Probabilities in Error," pp. S185–S194 in L. Darden (ed.), *Philosophy of Science*, 64 (PSA 1996: Symposia Proceedings).
Kelly, K., Schulte, O., Juhl, C., (1997), "Learning Theory and the Philosophy of Science," *Philosophy of Science*, 64: 245–67.
Laudan, L. (1997) "How about Bust? Factoring Explanatory Power Back into Theory Evaluation," *Philosophy of Science*, 64(2): 306–16.
Mayo, D.G. (1997a), "Duhem's Problem, the Bayesian Way, and Error Statistics, or "What's Belief Got to Do with It?" 64: 222–44.
Mayo, D.G. (1997b), "Error Statistics and Learning From Error: Making a Virtue of Necessity," pp. S195–S212 in L. Darden (ed.), *Philosophy of Science*, 64 (PSA 1996: Symposia Proceedings).
Mayo, D.G. (2000a), "Experimental Practice and an Error Statistical Account of Evidence," pp. S193–S207 in D. Howard (ed.), *Philosophy of Science*, 67 (Symposia Proceedings).
Mayo, D.G. (2000b), "Models of Error and the Limits of Experimental Testing," pp. 317–44 in M. Carrier, G. Massey and L. Ruetsche (eds.), *Science at Century's End, Philosophical Questions on the Progress and Limits of Science*, University of Pittsburgh Press, Pittsburgh.
Woodward, J. (1997), "Data, Phenomena, and Reliability," pp. S163–S179 in L. Darden (ed.), *Philosophy of Science*. 64 (PSA 1996: Symposia Proceedings).