

Philosophy and the practice of Bayesian statistics

Andrew Gelman

Department of Statistics and Department of Political Science, Columbia University

Cosma Rohilla Shalizi

Statistics Department, Carnegie Mellon University

Santa Fe Institute

18 June 2010

Abstract

A substantial school in the philosophy of science identifies Bayesian inference with inductive inference and even rationality as such, and seems to be strengthened by the rise and practical success of Bayesian statistics. We argue that the most successful forms of Bayesian statistics do not actually support that particular philosophy but rather accord much better with sophisticated forms of hypothetico-deductivism. We examine the actual role played by prior distributions in Bayesian models, and the crucial aspects of model checking and model revision, which fall outside the scope of Bayesian confirmation theory. We draw on the literature on the consistency of Bayesian updating and also on our experience of applied work in social science.

Clarity about these matters should benefit not just philosophy of science, but also statistical practice. At best, the inductivist view has encouraged researchers to fit and compare models without checking them; at worst, theorists have actively discouraged practitioners from performing model checking because it does not fit into their framework.

1 The usual story—which we don’t like

In so far as I have a coherent philosophy of statistics, I hope it is “robust” enough to cope in principle with the whole of statistics, and sufficiently undogmatic not to imply that all those who may think rather differently from me are necessarily stupid. If at times I do seem dogmatic, it is because it is convenient to give my own views as unequivocally as possible. (Bartlett, 1967, p. 458)

Schools of statistical inference are sometimes linked to approaches to the philosophy of science. “Classical” statistics—as exemplified by Fisher’s p -values, Neyman-Pearson hypothesis tests, and Neyman’s confidence intervals—is associated with the hypothetico-deductive and falsificationist view of science. Scientists devise hypotheses, deduce implications for observations from them, and test those implications. Scientific hypotheses can be rejected (that is, falsified), but never really established or accepted in the same way. Mayo (1996) presents the leading contemporary statement of this view.

In contrast, Bayesian statistics or “inverse probability”—starting with a prior distribution, getting data, and moving to the posterior distribution—is associated with an inductive approach of learning about the general from particulars. Rather than testing and attempted falsification, learning proceeds more smoothly: an accretion of evidence is summarized by a posterior distribution, and scientific process is associated with the rise and fall in the posterior probabilities of various models; see Figure 1 for a schematic illustration. In this view, the expression $p(\theta|y)$ says it all, and the central goal of Bayesian inference is computing the posterior probabilities of hypotheses. Anything not contained in the posterior distribution $p(\theta|y)$ is simply irrelevant, and it would be irrational (or incoherent) to attempt falsification, unless that somehow shows up in the posterior. The goal is to learn about general laws, as expressed in the probability that one model or another is correct. This view, strongly influenced by Savage (1954), is widespread and influential in the philosophy of science (especially in the form of Bayesian confirmation theory; see Howson and Urbach 1989; Earman 1992) and among Bayesian statisticians (Bernardo and Smith, 1994). Many people see support for this view in the rising use of Bayesian methods in applied statistical work over the last few decades.¹

We think most of this received view of Bayesian inference is wrong. Bayesian methods are no more inductive than any other mode of statistical inference, which is, not inductive in any strong sense. Bayesian data analysis is much better understood from a hypothetico-deductive perspective.² Implicit in the best Bayesian practice is a stance that has much in common with the error-statistical approach of Mayo (1996), despite the latter’s frequentist orientation. Indeed, crucial parts of Bayesian data analysis, such as model checking, can be understood as “error probes” in Mayo’s sense.

We proceed by a combination of examining concrete cases of Bayesian data analysis in empirical social science research, and theoretical results on the consistency and convergence of Bayesian updating. Social-scientific data analysis is especially salient for our purposes

¹Consider the current (9 June 2010) state of the Wikipedia article on Bayesian inference, which *begins* as follows:

Bayesian inference is statistical inference in which evidence or observations are used to update or to newly infer the probability that a hypothesis may be true.

It then continues with:

Bayesian inference uses aspects of the scientific method, which involves collecting evidence that is meant to be consistent or inconsistent with a given hypothesis. As evidence accumulates, the degree of belief in a hypothesis ought to change. With enough evidence, it should become very high or very low.... Bayesian inference uses a numerical estimate of the degree of belief in a hypothesis before evidence has been observed and calculates a numerical estimate of the degree of belief in the hypothesis after evidence has been observed.... Bayesian inference usually relies on degrees of belief, or subjective probabilities, in the induction process and does not necessarily claim to provide an objective method of induction. Nonetheless, some Bayesian statisticians believe probabilities can have an objective value and therefore Bayesian inference can provide an objective method of induction.

These views differ from those of, e.g., Bernardo and Smith (1994) or Howson and Urbach (1989) only in the omission of technical details.

²We are not interested in the hypothetico-deductive “confirmation theory” prominent in philosophy of science from the 1950s through the 1970s, and linked to the name of Hempel (1965). The hypothetico-deductive account of scientific method to which we appeal is distinct from, and much older than, this particular sub-branch of confirmation theory.




Figure 1: *Hypothetical picture of idealized Bayesian inference under the conventional inductive philosophy. The posterior probability of different models changes over time with the expansion of the likelihood as more data are entered into the analysis. Depending on the context of the problem, the time scale on the x-axis might be hours, years, or decades, in any case long enough for information to be gathered and analyzed that first knocks out hypothesis 1 in favor of hypothesis 2, which in turn is dethroned in favor of the current champion, model 3.*

because there is general agreement that, in this domain, all models in use are wrong—not merely falsifiable, but actually false. With enough data—and often only a fairly moderate amount—any analyst could reject any model now in use to any desired level of confidence. Model fitting is nonetheless a valuable activity, and indeed the crux of data analysis. To understand why this is so, we need to examine how models are built, fitted, used, and checked, and the effects of misspecification on models.

2 The data-analysis cycle

We begin with a very brief reminder of how statistical models are built and used in data analysis, following Gelman et al. (2003), or, from a frequentist perspective, Guttorm (1995).

The statistician begins with a model that stochastically generates all the data y , whose joint distribution is specified as a function of a vector of parameters θ from a space Θ (which may, in the case of some so-called non-parametric models, be infinite dimensional). This joint distribution is the likelihood function. The stochastic model may involve other, unmeasured but potentially observable variables \tilde{y} —that is, missing or latent data—and more-or-less fixed aspects of the data-generating process as covariates. For both Bayesians and frequentists, the joint distribution of (y, \tilde{y}) depends on θ . Bayesians insist on a full joint distribution, embracing observables, latent variables, and parameters, so that the likelihood function becomes a conditional probability density, $p(y|\theta)$. In designing the stochastic process for (y, \tilde{y}) , the goal is to represent the systematic relationships between the variables and between the variables and the parameters, and as well as to represent the noisy (contingent, accidental, irreproducible) aspects of the data stochastically. Against the desire

for accurate representation one must balance conceptual, mathematical and computational tractability. Some parameters thus have fairly concrete real-world referents, such as the famous (in statistics) survey of the rat population of Baltimore (Brown et al., 1955). Others, however, will reflect the specification as a mathematical object more than the reality being modeled— t -distributions are sometimes used to model heavy-tailed observational noise, with the number of degrees of freedom for the t representing the shape of the distribution; few statisticians would take this as realistically as the number of rats.

Bayesian modeling, as mentioned, requires a joint distribution for (y, \tilde{y}, θ) , which is conveniently factored (without loss of generality) into a prior distribution for the parameters, $p(\theta)$, and the complete-data likelihood, $p(y, \tilde{y}|\theta)$, so that $p(y|\theta) = \int p(y, \tilde{y}|\theta)d\tilde{y}$. The prior distribution is, as we will see, really part of the model. In practice, the various parts of the model have functional forms picked by a mix of substantive knowledge, scientific conjectures, statistical properties, analytical convenience, and computational tractability.

Having completed the specification, the Bayesian analyst calculates the posterior distribution $p(\theta|y)$; it is so that this quantity makes sense that the observed y and the parameters θ must have a joint distribution. The rise of Bayesian methods in applications has rested on finding new ways of to actually carry through this calculation, even if only approximately, notably by adopting Markov chain Monte Carlo methods, originally developed in statistical physics to evaluate high-dimensional integrals (Metropolis et al., 1953; Newman and Barkema, 1999), to sample from the posterior distribution. The natural counterpart of this stage for non-Bayesian analyses are various forms of point and interval estimation to identify the set of values of θ that are consistent with the data y .

According to the view we sketched above, data analysis basically ends with the calculation of the posterior $p(\theta|y)$. At most, this might be elaborated by partitioning Θ into a set of models or hypotheses, $\Theta_1, \dots, \Theta_K$, each with a prior probability $p(\Theta_k)$ and its own set of parameters θ_k . One would then compute the posterior parameter distribution within each model, $p(\theta_k|y, \Theta_k)$, and the posterior probabilities of the models,

$$\begin{aligned} p(\Theta_k|y) &= \frac{p(\Theta_k)p(y|\Theta_k)}{\sum_{k'} (p(\Theta_{k'})p(y|\Theta_{k'}))} \\ &= \frac{p(\Theta_k) \int p(y, \theta_k|\Theta_k)d\theta_k}{\sum_{k'} (p(\Theta_{k'}) \int p(y, \theta_k|\Theta_{k'})d\theta_{k'})}. \end{aligned}$$

These posterior probabilities of hypotheses can be used for Bayesian model selection or Bayesian model averaging (topics to which we return below). Scientific progress, in this view, consists of gathering data—perhaps through well-designed experiments, designed to distinguish among interesting competing scientific hypotheses (cf. Atkinson and Donev, 1992; Paninski, 2005)—and then plotting the $p(\Theta_k|y)$'s over time and watching the system learn (as sketched in Figure 1).

In our view, the account of the last paragraph is crucially mistaken. The data-analysis process—Bayesian or otherwise—does not end with calculating parameter estimates or posterior distribution. Rather, the model can then be *checked*, by comparing the implications of the fitted model to the empirical evidence. One asks questions like, Do simulations from the fitted model resemble the original data? Is the fitted model consistent with other data not used in the fitting of the model? Do variables that the model says are noise (“error terms”) in fact display readily-detectable patterns? Discrepancies between the model and

data can be used to learn about the ways in which the model is inadequate for the scientific purposes at hand, and thus to motivate expansions and changes to the model (§4).

2.1 Example: Estimating voting patterns in subsets of the population

We demonstrate the hypothetico-deductive Bayesian modeling process with an example from our recent applied research (Gelman et al., 2010). In recent years, American political scientists have been increasingly interested in the connections between politics and income inequality (see, e.g., McCarty et al. 2006). In our own contribution to this literature, we estimated the attitudes of rich, middle-income, and poor voters in each of the fifty states (Gelman et al., 2008b). As we described in our article on the topic (Gelman et al., 2008c), we began by fitting a varying-intercept logistic regression: modeling votes (coded as $y = 1$ for votes for the Republican presidential candidate or $y = 0$ for Democratic votes) given family income (coded in five categories from low to high as $x = -2, -1, 0, 1, 2$), using a model of the form $\Pr(y = 1) = \text{logit}^{-1}(a_s + bx)$, where s indexes state of residence—the model is fit to survey responses—and the varying intercepts a_s correspond to some states being more Republican-leaning than others. Thus, for example a_s has a positive value in a conservative state such as Utah and a negative value in a liberal state such as California. The coefficient b represents the “slope” of income, and its positive value indicates that, within any state, richer voters are more likely to vote Republican.

It turned out that this varying-intercept model did not fit our data, as we learned by making graphs of the average survey response and fitted curves for the different income categories within each state. We had to expand to a varying-intercept, varying-slope model, $\Pr(y = 1) = \text{logit}^{-1}(a_s + b_s x)$, in which the slopes b_s varied by state as well. This model expansion led to a corresponding expansion in our understanding: we learned that the gap in voting between rich and poor is much greater in poor states such as Mississippi than in rich states such as Connecticut. Thus, the polarization between rich and poor voters varied in important ways geographically.

We found this not through any process of Bayesian induction but rather through model checking. Bayesian inference was crucial, not for computing the posterior probability that any particular model was true—we never actually did that—but in allowing us to fit rich enough models in the first place that we could study state-to-state variation, incorporating in our analysis relatively small states such as Mississippi and Connecticut that did not have large samples in our survey. (Gelman and Hill (2006) review the hierarchical models that allow such partial pooling.)

Life continues, though, and so do our statistical struggles. After the 2008 election, we wanted to make similar plots, but this time we found that even our more complicated logistic regression model did not fit the data—especially when we wanted to expand our model to estimate voting patterns for different ethnic groups. Comparison of data to fit led to further model expansions, leading to our current specification, which uses a varying-intercept, varying-slope logistic regression as a baseline but allows for nonlinear and even non-monotonic patterns on top of that. Figure 2 shows some of our inferences in map form, while Figure 3 shows one of our diagnostics of data and model fit.

The power of Bayesian inference here is *deductive*: given the data and some model assumptions, it allows us to make lots of inferences, many of which can be checked and




Figure 2: Based on a model fitted to survey data: states won by John McCain and Barack Obama among different ethnic and income categories. States colored deep red and deep blue indicate clear McCain and Obama wins; pink and light blue represent wins by narrower margins, with a continuous range of shades going to gray for states estimated at exactly 50/50. The estimates shown here represent the culmination of months of effort, in which we fit increasingly complex models, at each stage checking the fit by comparing to data and then modifying aspects of the prior distribution and likelihood as appropriate.

2008 election: McCain share of the two-party vote in each income category within each state among all voters (black) and non-Hispanic whites (green)




Figure 3: Some of the data and fitted model used to make the maps shown in Figure 2. Dots are weighted averages from pooled June-November Pew surveys; error bars show ± 1 standard error bounds. Curves are estimated using multilevel models and have a standard error of about 3% at each point. States are ordered in decreasing order of McCain vote (Alaska, Hawaii, and D.C. excluded). We fit a series of models to these data; only this last model fit the data well enough that we were satisfied. In working with larger datasets and studying more complex questions, we encounter increasing opportunities to check model fit and thus falsify in a way that is helpful for our research goals.

potentially falsified. For example, look at New York state (in the bottom row of Figure 3): apparently, voters in the second income category supported John McCain much more than did voters in neighboring income groups in that state. This pattern is theoretically possible but it arouses suspicion. A careful look at the graph reveals that this is a pattern in the raw data which was moderated but not entirely smoothed away by our model. The natural next step would be to examine data from other surveys. We may have exhausted what we can learn from this particular dataset, and Bayesian inference was a key tool in allowing us to do so.

3 The Bayesian principal-agent problem

Before returning to discussions of induction and falsification, we briefly discuss some findings relating to Bayesian inference under misspecified models. The key idea is that Bayesian inference for model selection—statements about the posterior probabilities of candidate models—does not solve the problem of learning from data about problems with existing models.

In economics, the “principal-agent problem” refers to the difficulty of designing institutions which ensure that one selfish actor, the “agent,” will act in the interests of another, the “principal,” who cannot monitor and sanction their agent without cost or error. The problem is one of aligning incentives, so that the agent serves itself by serving the principal. There is, one might say, a Bayesian principal-agent problem as well. The Bayesian agent is the methodological fiction (now often approximated in software) of a creature with a prior distribution over a well-defined hypothesis space Θ , a likelihood function $p(y|\theta)$, and conditioning as its sole mechanism of learning and belief revision. The principal is the actual statistician or scientist.

The Bayesian agent’s ideas are much more precise than the actual scientist’s; in particular, the Bayesian (in this formulation, with which we disagree) is certain that *some* θ is the exact and complete truth, whereas the scientist is not. At some point in history, a statistician may well write down a model which he or she believes contains all the systematic influences among properly-defined variables for the system of interest, with correct functional forms and distributions of noise terms. This could happen, but we have never seen it, and in social science we’ve never seen anything that comes close, either. If nothing else, our own experience suggests that however many different specifications we think of, there are always others which had not occurred to us, but cannot be immediately dismissed *a priori*, if only because they can be seen as alternative approximations to the ones we made. Yet the Bayesian agent is required to start with a prior distribution whose support covers *all* alternatives that could be considered.³

This is not a small technical problem to be handled by adding a special value of θ , say θ^∞ standing for “none of the above”; even if one could calculate $p(y|\theta^\infty)$, the likelihood of the data under this catch-all hypothesis, this in general would *not* lead to just a small correction to the posterior, but rather would have substantial effects (Fitelson and Thomason, 2008). Fundamentally, the Bayesian agent is limited by the fact that its beliefs always remain

³It is also not at all clear that Savage and other founders of Bayesian decision theory ever thought that this principle should apply outside of the small worlds of artificially simplified and stylized problems—see Binmore (2007). But as scientists we care about the real, large world.

within the support of its prior. For the Bayesian agent, the truth must, so to speak, be always already partially believed before it can become known. This point is less than clear in the usual treatments of Bayesian convergence, and so worth some attention.

Classical results (Doob, 1949; Schervish, 1995; Lijoi et al., 2007) show that the Bayesian agent's posterior distribution will concentrate on the truth with *prior* probability 1, provided some regularity conditions are met. Without diving into the measure-theoretic technicalities, the conditions amount to (i) the truth is in the support of the prior, and (ii) the information set is rich enough that some consistent estimator exists. (See the discussion in Schervish (1995, §7.4.1).) When the truth is *not* in the support of the prior, the Bayesian agent still thinks that Doob's theorem applies and assigns zero prior probability to the set of data under which it does not converge on the truth.

The convergence behavior of Bayesian updating with a misspecified model can be understood as follows (Berk, 1966; Kleijn and van der Vaart, 2006; Shalizi, 2009). If the data are actually coming from a distribution q , then the Kullback-Leibler divergence rate, or relative entropy rate, of the parameter value θ is

$$d(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbf{E} \left[\log \frac{p(y_1, y_2, \dots, y_n | \theta)}{q(y_1, y_2, \dots, y_n)} \right],$$

with the expectation being taken under q . (For details on when the limit exists, see Gray 1990.) Then, under not-too-onerous regularity conditions, one can show (Shalizi, 2009) that

$$p(\theta | y_1, y_2, \dots, y_n) \approx p(\theta) \exp \{ -n(d(\theta) - d^*) \},$$

with d^* being the essential infimum of the divergence rate. More exactly,

$$-\frac{1}{n} \log p(\theta | y_1, y_2, \dots, y_n) \rightarrow d(\theta) - d^*,$$

q -almost-surely. Thus the posterior distribution comes to concentrate on the parts of the prior support which have the lowest values of $d(\theta)$ and the highest expected likelihood.⁴ There is a geometric sense in which these parts of the parameter space are closest approaches to the truth within the support of the prior (Kass and Vos, 1997), but they may or may not be close to the truth in the sense of giving accurate values for parameters of scientific interest. They may not even be the parameter values which give the best predictions (Grünwald and Langford, 2007; Müller, 2010). In fact, one cannot even guarantee that the posterior will concentrate on a single value of θ at all; if $d(\theta)$ has multiple global minima, the posterior can alternate between (concentrating around) them forever (Berk, 1966).

To sum up, what Bayesian updating does when the model is false (that is, in reality, always) is to try to concentrate the posterior on the best attainable approximations to the distribution of the data, “best” being measured by likelihood. But depending on *how* the model is misspecified, and how θ represents the parameters of scientific interest, the impact of misspecification on inferring the latter can range from non-existent to profound.⁵ Since

⁴More precisely, regions of Θ where $d(\theta) > d^*$ tend to have exponentially small posterior probability; this statement covers situations like $d(\theta)$ only approaching its essential infimum as $\|\theta\| \rightarrow \infty$, etc. See Shalizi (2009) for details.

⁵White (1994) gives examples of econometric models where the influence of mis-specification on the parameters of interest runs through this whole range, though only considering maximum likelihood and maximum quasi-likelihood estimation.

we are quite sure our models are wrong, we need to check whether the misspecification is so bad that inferences regarding the scientific parameters are in trouble. It is by this non-Bayesian checking of Bayesian models that we solve our principal-agent problem.

4 Model checking

In our view, a key part of Bayesian data analysis is model checking, which is where there are links to falsificationism. In particular, we emphasize the role of posterior predictive checks, creating simulations and comparing the simulated and actual data; these comparisons can often be done visually (Gelman et al., 2003, ch. 6).

Here's how this works. A Bayesian model gives us a joint distribution for the parameters θ and the observables y . This implies a marginal distribution for the data,

$$p(y) = \int p(y|\theta)p(\theta)d\theta.$$

If we have observed data y , the prior distribution $p(\theta)$ shifts to the posterior distribution $p(\theta|y)$, and so a different distribution of observables,

$$p(y^{\text{rep}}|y) = \int p(y^{\text{rep}}|\theta)p(\theta|y)d\theta,$$

where we use the y^{rep} to indicate hypothetical alternative or future data, a replicated data set of the same size and shape as the original y , generated under the assumption that the fitted model, prior and likelihood both, is true. By simulating from the posterior distribution of y^{rep} , we see what typical realizations of the fitted model are like, and in particular whether the observed dataset is the kind of thing that the fitted model produces with reasonably high probability.⁶

If we summarize the data with a test statistic $T(y)$, we can perform graphical comparisons with replicated data and calculate p -values,

$$\Pr(T(y^{\text{rep}}) > T(y)|y),$$

which can be approximated to arbitrary accuracy as soon as we can simulate y^{rep} . (This is a valid posterior probability in the model, and its interpretation is no more problematic than that of any other probability in a Bayesian model.) In practice, graphical test summaries are often more illuminating than p -values, but in considering ideas of (probabilistic) falsification, it can be helpful to think about numerical test statistics.

Under the usual understanding that T is chosen so large values indicate poor fits, these p -values work rather like classical ones (Mayo, 1996; Mayo and Cox, 2006)—they in fact are generalizations of classical p -values, merely replacing point estimates of parameters θ with averages over the posterior distribution—and their basic logic is one of falsification. A very low p -value says that it is very improbable, under the model, to get data as extreme along the T -dimension as the actual y ; we are seeing something which would be very improbable

⁶For notational simplicity, we leave out the possibility of generating new values of the hidden variables \tilde{y} and set aside choices of which parameters to vary and which to hold fixed in the replications; see Gelman et al. (1996).

if the model were true. On the other hand a high p -value merely indicates that $T(y)$ is an aspect of the data which would be unsurprising if the model is true. Whether this is evidence *for* the usefulness of the model depends how likely it is to get such a high p -value when the model is false: the “severity” of the test, in the terminology of Mayo (1996) and Mayo and Cox (2006).

Put a little more abstractly, the hypothesized model makes certain probabilistic assumptions, from which other probabilistic implications follow deductively. Simulation works out what those implications are, and tests check whether the data conform to them. Extreme p -values indicate that the data violate regularities implied by the model, or approach doing so. If these were strict violations of deterministic implications, we could just apply *modus tollens* to conclude that the model was wrong; as it is, we nonetheless have evidence and probabilities. Our view of model checking, then, is firmly in the long hypothetico-deductive tradition, running from Popper (1934/1959) back through Bernard (1865/1927) and beyond (Laudan, 1981). A more direct influence on our thinking about these matters is the work of Jaynes (2003), who illustrated how we may learn the most when we find that our model does not fit the data—that is, when it is falsified—because then we have found a problem with our model’s assumptions.⁷ And the better our probability model encodes our *scientific* or *substantive* assumptions, the more we learn from specific falsification.

In this connection, the prior distribution $p(\theta)$ is one of the assumptions of the model and does not need to represent the statistician’s personal degree of belief in alternative parameter values. The prior is connected to the data, and so is potentially testable, via the posterior predictive distribution of future data y^{rep} :

$$\begin{aligned} p(y^{\text{rep}}|y) &= \int p(y^{\text{rep}}|\theta)p(\theta|y)d\theta \\ &= \int p(y^{\text{rep}}|\theta) \frac{p(y|\theta)p(\theta)}{\int p(y|\theta')p(\theta')d\theta'} d\theta. \end{aligned}$$

The prior distribution thus has implications for the distribution of replicated data, and so can be checked using the type of tests we have described, and illustrated above.⁸ When it makes sense to think of further data coming from the same source, as in certain kinds of sampling, time-series or longitudinal problems, the prior also has implications for these new data (through the same formula as above, changing the interpretation of y^{rep}), and so becomes testable in a second way. There is thus a connection between the model-checking aspect of Bayesian data analysis and “prequentialism” (Dawid and Vovk, 1999; Grünwald, 2007), but exploring that would take us too far afield.

One advantage of recognizing that the prior distribution is a testable part of a Bayesian model is that it clarifies the role of the prior in inference, and where it comes from. To

⁷A similar point was expressed by the sociologist and social historian Charles Tilly, writing from a very different disciplinary background: “Most social researchers learn more from being wrong than from being right—provided they then recognize that they were wrong, see why they were wrong, and go on to improve their arguments. Post hoc interpretation of data minimizes the opportunity to recognize contradictions between arguments and evidence, while adoption of formalisms increases that opportunity. Formalisms blindly followed induce blindness. Intelligently adopted, however, they improve vision. Being obliged to spell out the argument, check its logical implications, and examine whether the evidence conforms to the argument promotes both visual acuity and intellectual responsibility.” (Tilly, 2004, p. 597)

⁸Admittedly, the prior only has observable implications in conjunction with the likelihood, but for a Bayesian the reverse is also true.

reiterate, it is hard to claim that the prior distributions used in applied work represent statisticians' states of knowledge and belief before examining their data, if only because most statisticians do not believe their models are true, so their prior degree of belief in all of Θ is not 1 but 0. The prior distribution is more like a regularization device, akin to the penalization terms added to the sum of squared errors when doing ridge regression and the lasso (Hastie et al., 2001) or spline smoothing (Wahba, 1990). All such devices exploit a sensitivity-stability tradeoff: they stabilize estimates and predictions by making fitted models less sensitive to certain details of the data. Using an informative prior distribution (even if only weakly informative, as in Gelman et al. (2008a)) makes our estimates less sensitive to the data than, say, maximum-likelihood estimates would be, which can be a net gain.⁹

Because we see the prior distribution as a testable part of the Bayesian model, we do not need to follow Jaynes in trying to devise unique, objectively-correct prior distribution for each situation—an enterprise with an uninspiring track record (Kass and Wasserman, 1996), even leaving aside doubts about Jaynes's specific proposal (Seidenfeld, 1979, 1987; Csiszár, 1995; Uffink, 1995, 1996). To put it even more succinctly, “the model,” for a Bayesian, is the combination of the prior distribution and the likelihood, each of which represents some compromise among scientific knowledge, mathematical convenience, and computational tractability.

This gives us a lot of flexibility in modeling. We do not have to worry about making our prior distributions match our subjective beliefs, still less about our model containing all possible truths. Instead we make some assumptions, state them clearly, see what they imply, and check the implications. This applies just much to the prior distribution as it does to the parts of the model showing up in the likelihood function.

4.1 Testing to reveal problems with a model

We are not interested in falsifying our model for its own sake—among other things, having built it ourselves, we know all the shortcuts taken in doing so, and can already be morally certain it is false. With enough data, we can certainly detect departures from the model—this is why, for example, statistical folklore says that the chi-squared statistic is ultimately a measure of sample size (cf. Lindsay and Liu 2009). As writers such as Giere (1988, ch. 3) explain, the hypothesis linking mathematical models to empirical data is not that the data-generating process is exactly isomorphic to the model, but that the data source resembles the model closely enough, in the respects which matter to us, that reasoning based on the model will be reliable. Such reliability does not require complete fidelity to the model.

The goal of model checking, then, is not to demonstrate the foregone conclusion of falsity as such, but rather to learn how, in particular, this model fails (Gelman, 2003). When we find such particular failures, they tell us how the model must be improved; when severe tests cannot find them, the inferences we draw about those aspects of the real world from

⁹A further advantage to using a prior in conjunction with misspecified models can be improved prediction; see Page (2007). The posterior predictive distribution averages over all values of θ , so its expected error equals the average of the expected errors of the individual $p(y|\theta)$, *minus* the variance of the predictions over Θ (Krogh and Vedelsby, 1995). Thus the predictions resulting from Bayesian model averaging can be more accurate than even the best individual prediction possible with the model. However, since our interest here is mainly in scientific inference and not in prediction, we will say no more about this here.

our fitted model become more credible. In designing a *good* test for model checking, we are interested in finding particular errors which, if present, would mess up particular inferences, and devise a test statistic which is sensitive to this sort of mis-specification.

All models will have errors of approximation. Statistical models, however, typically assert that their errors of approximation will be unsystematic and patternless—“noise” (Spanos, 2007). Testing this can be valuable in revising the model. In looking at the red-state/blue-state example, for instance, we concluded that the varying slopes mattered not just because of the magnitudes of departures from the equal-slope assumption, but also because there was a pattern, with richer states tending to have shallower slopes.

What we are advocating, then, is what Cox and Hinkley (1974) call “pure significance testing,” in which certain of the model’s implications are compared directly to the data, rather than entering into a contest with some alternative model. This is, we think, more in line with what actually happens in science, where it can become clear that even large-scale theories are in serious trouble and cannot be accepted unmodified even if there is no alternative available yet. A classical instance is the status of Newtonian physics at the beginning of the 20th century, where there were enough difficulties—the Michelson-Morley effect, anomalies in the orbit of Mercury, the photoelectric effect, the black-body paradox, the stability of charged matter, etc.—that it was clear, even before relativity and quantum mechanics, that something would have to give. Even today, our current best theories of fundamental physics, namely general relativity and the standard model of particle physics, an instance of quantum field theory, are universally agreed to be ultimately wrong, not least because they are mutually incompatible, and recognizing this does not require that one have a replacement theory (Weinberg, 1999).

4.2 Connection to non-Bayesian model checking

Many of these ideas about model checking are not unique to Bayesian data analysis and are used more or less explicitly by many communities of practitioners working with complex stochastic models (Ripley 1988; Guttorp 1995). The reasoning is the same: a model is a story of how the data could have been generated; the fitted model should therefore be able to generate synthetic data that look like the real data; failures to do so in important ways indicate faults in the model.

For instance, simulation-based model checking is now widely accepted for assessing the goodness of fit of statistical models of social networks (Hunter et al., 2008). That community was pushed toward predictive model checking by the observation that many model specifications were “degenerate” in various ways (Handcock, 2003). For example, under certain exponential-family network models, the maximum likelihood estimate gave a distribution over networks which was bimodal, with both modes being very different from observed networks, but located so that the expected value of the sufficient statistics matched observations. It was thus clear that these specifications could not be right even before more adequate specifications were developed (Snijders et al., 2006).

At a more philosophical level, the idea that a central task of statistical analysis is the search for specific, consequential errors has been forcefully advocated by Mayo (1996), Mayo and Cox (2006); Mayo and Spanos (2004), and Mayo and Spanos (2006). Mayo has placed a special emphasis on the idea of *severe* testing—a model being severely tested if it passes a

probe which had a high probability of detecting an error if it is present. (The exact definition of a test's severity is related to, but not quite, that of its power; see Mayo 1996 or Mayo and Spanos 2006 for extensive discussions.) Something like this is implicit in discussions about the relative merits of particular posterior predictive checks (which can also be framed non-Bayesian as graphical hypothesis tests based on the parametric bootstrap).

Our contribution here is to connect this hypothetico-deductive philosophy to Bayesian data analysis, going beyond the evaluation of Bayesian methods based on their frequency properties (as recommended by Rubin (1984), Wasserman (2006), among others) to emphasize the learning that comes from the discovery of systematic differences between model and data. At the very least, we hope this paper will motivate philosophers of hypothetico-deductive inference to take a more serious look at Bayesian data analysis (as distinct from Bayesian theory) and, conversely, to motivate philosophically-minded Bayesian statisticians to consider alternatives to the inductive interpretation of Bayesian learning.

4.3 Why not just compare the posterior probabilities of different models?

As mentioned above, the standard view of scientific learning in the Bayesian community is, roughly, that posterior odds of the models under consideration are compared, given the current data.¹⁰ When Bayesian data analysis is understood as simply getting the posterior distribution, it is held that “pure significance tests have no role to play in the Bayesian framework” (Schervish, 1995, p. 218). The dismissal rests on the idea that the prior distribution can accurately reflect our actual knowledge and beliefs.¹¹ At the risk of boring the reader by repetition, there is just no way we can ever have any hope of making Θ include all the probability distributions which might be correct, let alone getting $p(\theta|y)$ if we did so, so this is deeply unhelpful advice. The main point where we disagree with many Bayesians is that we do not see Bayesian methods as generally useful for giving the posterior probability that a model is true, or the probability for preferring model A over model B, or whatever.¹² Beyond the philosophical difficulties, there are technical problems with methods that purport to determine the posterior probability of models, most notably

¹⁰Some would prefer to compare the modification of those odds called the Bayes factor (Kass and Raftery, 1995). Everything we have to say about posterior odds carries over to Bayes factors with few changes.

¹¹As Schervish (1995) continues: “If the [parameter space Θ] describes all of the probability distributions one is willing to entertain, then one cannot reject $[\Theta]$ without rejecting probability models altogether. If one is willing to entertain models not in $[\Theta]$, then one needs to take them into account” by enlarging Θ , and computing the posterior distribution over the enlarged space.

¹²There is a vast literature on Bayes factors, model comparison, model averaging, and the evaluation of posterior probabilities of models, and although we believe most of this work to be philosophically unsound (to the extent it is designed to be a direct vehicle for scientific learning), we recognize that these can be useful techniques. Like all statistical methods, Bayesian and otherwise, these methods are summaries of available information that can be important data-analytic tools. Even if none of a class of models is plausible as truth, and even if we aren't comfortable accepting posterior model probabilities as degrees of belief in alternative models, these probabilities can still be useful as tools for prediction and for understanding structure in data, as long as these probabilities are not taken too seriously. See Raftery (1995) for a discussion of the value of posterior model probabilities in social science research and Gelman and Rubin (1995) for a discussion of their limitations, and Claeskens and Hjort (2008) for a general review of model selection. (Some of the work on “model-selection tests” in econometrics (e.g., Vuong 1989; Rivers and Vuong 2002) is exempt from our strictures, as it tries to find which model is *closest* to the data-generating process, while allowing that all of the models may be mis-specified, but it would take us too far afield to discuss this work in detail.)

that in models with continuous parameters, aspects of the model that have essentially no effect on posterior inferences *within* a model can have huge effects on the comparison of posterior probability *among* models.¹³ Bayesian inference is good for deductive inference within a model we prefer to evaluate a model by comparing it to data.

In practice, if we are in a setting where model A or model B might be true, we are inclined not to do *model selection* among these specified options, or even to perform *model averaging* over them (perhaps with a statement such as, “We assign 40% of our posterior belief to A and 60% to B”) but rather to do *continuous model expansion* by forming a larger model that includes both A and B as special cases. For example, Merrill (1994) used electoral and survey data from Norway and Sweden to compare two models of political ideology and voting: the “proximity model” (in which you prefer the political party that is closest to you in some space of issues and ideology) and the “directional model” (in which you like the parties that are in the same direction as you in issue space, but with a stronger preference to parties further from the center). Rather than using the data to pick one model or the other, we would prefer to think of a model in which voters consider both proximity and directionality in forming their preferences (Gelman, 1994).

In the social sciences, it is rare for there to be an underlying theory that can provide meaningful constraints on the functional form of the expected relationships among variables, let alone the distribution of noise terms.¹⁴ Taken seriously, then, this advice would imply that social scientists should more or less give up using parametric statistical models in favor of nonparametrics (Ghosh and Ramamoorthi, 2003). And while a greater use of nonparametric models in empirical research may be desirable on its own merits (see Li and Racine, 2007), even this would not really resolve the issue, as nonparametric models themselves embody assumptions such as conditional independence which are hard to defend except as approximations. Expanding our prior distribution to embrace all the models which are actually compatible with our prior knowledge would result in a mess we simply could not work with, nor interpret if we could.

4.4 Example: Estimating the effects of legislative redistricting

We use one of our own experiences (Gelman and King, 1994) to illustrate scientific progress through model rejection. We began by fitting a model comparing treated and control units—state legislatures, immediately after redistricting or not—following the usual practice of assuming a constant treatment effect (parallel regression lines in “after” vs. “before” plots, with the treatment effect representing the difference between the lines). In this example, the outcome was a measure of partisan bias, with positive values representing state legislatures where the Democrats were overrepresented (compared to how we estimated the Republicans would have done with comparable vote shares) and negative values in states where the Republicans were overrepresented. A positive treatment effect here would correspond to a

¹³This problem has been called the Jeffreys-Lindley paradox and it is the subject of a large literature. Unfortunately (from our perspective) the problem has usually been studied by Bayesians with an eye toward “solving” it—that is, coming up with reasonable definitions that allow the computation of nondegenerate posterior probabilities for continuously-parameterized models—but we we think that this is really a problem without a solution; see Gelman et al. (2003, sec. 6.7).

¹⁴Manski (2007) criticizes the econometric practice of making modeling assumptions (such as linearity) with no support in economic theory, simply to get identifiability.




Figure 4: *Sketch of the usual statistical model for before-after data. The difference between the fitted lines for the two groups is the estimated treatment effect. The default is to regress the “after” measurement on the treatment indicator and the “before” measurement, thus implicitly assuming parallel lines.*

redrawing of the district lines that favored the Democrats.

Figure 4 shows the default model that we (and others) typically use for estimating causal effects in before-after data. We fitted such a no-interaction model in our example too, but then we made some graphs and realized that the model did not fit the data. The line for the control units actually had a much steeper slope than the treated units. We fit a new model, and it had a completely different story about what the treatment effects meant.

The graph for the new model with interactions is shown in Figure 5. The largest effect of the treatment was not to benefit the Democrats or Republicans (that is, to change the intercept in the regression, shifting the fitted line up or down) but rather to change the slope of the line, to reduce partisan bias.

Rejecting the constant-treatment-effect model and replacing by the interaction model was, in retrospect, a crucial step in this research project. This pattern of higher before-after correlation in the control group than the treated group is quite general (Gelman, 2004), but at the time we did this study we discovered it only through the graph of model and data, which falsified the original model and motivated us to think of something better. In our experience, falsification is about plots and predictive checks, not about Bayes factors or posterior probabilities of candidate models.

The relevance of this example to the philosophy of statistics is that we began by fitting the usual regression model with no interactions. Only after visually checking the model fit—and thus falsifying it in a useful way without the specification of any alternative—did we take the crucial next step of including an interaction, which changed the whole direction of our research. The shift was induced by a falsification—a bit of deductive inference from the data and the earlier version of our model. In this case the falsification came from a graph rather than a p -value, which in one way is just a technical issue, but in a larger sense is important in that the graph revealed not just a lack of fit but also a sense of the direction of the misfit, a refutation that sent us usefully in a direction of substantive model improvement.




Figure 5: *Effect of redistricting on partisan bias.* Each symbol represents a state election year, with dots indicating controls (years with no redistricting) and the other symbols corresponding to different types of redistricting. As indicated by the fitted lines, the “before” value is much more predictive of the “after” value for the control cases than for the treated (redistricting) cases. The dominant effect of the treatment is to bring the expected value of partisan bias toward zero, and this effect would not be discovered with the usual approach (pictured in Figure 4, which is to fit a model assuming parallel regression lines for treated and control cases).

5 The question of induction

As we mentioned at the beginning, Bayesian inference is often held to be inductive in a way which classical statistics (following the Fisher or Neyman-Pearson traditions) is not. We need to address this, as we are arguing that all these forms of statistical reasoning are better seen as hypothetico-deductive.

The common core of various conceptions of induction is some form of inference from particulars to the general—in the statistical context, presumably, inference from the observations y to parameters θ describing the data-generating process. But if *that* were all that was meant, then not only is “frequentist statistics a theory of inductive inference” (Mayo and Cox, 2006), but the whole range of guess-and-test behaviors engaged in by animals Holland et al. (1986) are formalized in the hypothetico-deductive method are also inductive. Even the unpromising-sounding procedure, “Pick a model at random and keep it until its accumulated error gets too big, then pick another model completely at random,” would qualify (and could work surprisingly well under some circumstances; cf. Ashby (1960); Foster and Young (2003)). So would utterly irrational procedures (“pick a new random θ when the sum of the least significant digits in y is 13”). Clearly something more is required, or at least implied, by those claiming that Bayesian updating is inductive.

One possibility for that “something more” is to generalize the truth-preserving property of valid deductive inferences: just as valid deductions from true premises are themselves true, good inductions from true observations should also be true, at least in the limit of increasing

evidence.¹⁵ This, however, is just the requirement that our inferential procedures be consistent. As discussed above, using Bayes's rule is not sufficient to ensure consistency, nor is it necessary. In fact, every proof of Bayesian consistency known to us either posits there is a consistent non-Bayesian procedure for the same problem, or makes other assumptions which entail the existence of such a procedure. In any case, theorems establishing consistency of statistical procedures make *deductively valid* guarantees about these procedures—they are theorems, after all—but do so on the basis of probabilistic assumptions linking future events to past data.

It is also no good to say that what makes Bayesian updating inductive is its conformity to some axiomatization of rationality. If one accepts the Kolmogorov axioms for probability, and the Savage axioms (or something like them) for decision-making,¹⁶ then updating by conditioning follows, and a prior belief state $p(\theta)$ plus data y *deductively* entail that the new belief state is $p(\theta|y)$. In any case, lots of learning procedures can be axiomatized (all of them which can be implemented algorithmically, to start with), and these particular axioms do not in fact guarantee good results, like approaching the truth rather than becoming convinced of falsehoods—that's just the question of consistency again.

Karl Popper, the leading advocate of hypothetico-deductivism in the last century, denied that induction was even possible; his attitude is well-paraphrased by Greenland (1998) as: “we never use any argument based on observed repetition of instances that does not also involve a hypothesis that predicts both those repetitions and the unobserved instances of interest.” This is a recent instantiation of a tradition of anti-inductive arguments that goes back to Hume, but also beyond him to al Ghazali (1100/1997) in the middle ages, and indeed to the ancient Skeptics (Kolakowski, 1968). As forcefully put by Stove (1982, 1986), many apparent arguments against this view of induction can be viewed as statements of abstract premises linking both the observed data and unobserved instances—various versions of the “uniformity of nature” thesis have been popular, sometimes resolved into a set of more detailed postulates, as in Russell (1948, part VI, ch. 9), though Stove rather maliciously crafted a parallel argument for the existence of “angels, or something very much like them.”¹⁷ As Norton (2003) argues, these highly abstract premises are both dubious and often superfluous for supporting the sort of actual inferences scientists make—“inductions” are supported not by their matching certain formal criteria (as deductions are), but rather by material facts. To generalize about the melting point of bismuth (to use one of Norton’s examples) requires very few samples, provided we accept certain facts about the homogeneity of the physical properties of elemental substances; whether nature in general is uniform is not really at issue.

Simply put, we think the anti-inductivist view is pretty much right, but that statistical models are tools that let us draw inductive inferences on a deductive background. Most directly, random sampling allows us to learn about unsampled people (unobserved balls in an urn, as it were), but such inference, however inductive it may appear, relies not any axiom

¹⁵We owe this suggestion to conversation with Kevin Kelly; cf. Kelly (1996, esp. ch. 13).

¹⁶Despite his ideas on testing, Jaynes (2003) was a prominent and emphatic advocate of the claim that Bayesian inference is the logic of inductive inference as such, but preferred to follow Cox (1946, 1961) rather than Savage. See Halpern (1999) on the formal invalidity of Cox’s proofs.

¹⁷Stove (1986) further argues that induction by simple enumeration is reliable *without* making such assumptions, at least sometimes. However, his calculations make no sense unless his data are independent and identically distributed.

of induction but rather on deductions from the statistical properties of random samples, and the ability to actually conduct such sampling. The appropriate design depends on many contingent material facts about the system we are studying, exactly as Norton argues.

Some results in statistical learning theory establish that certain procedures are “probably approximately correct” in what’s called a “distribution-free” manner (Bousquet et al., 2004; Vidyasagar, 2003); some of these results embrace Bayesian updating (McAllister, 1999). But, here, “distribution free” just means “holding uniformly over all distributions in a very large class,” for example requiring the data to be independent and identically distributed, or from a stationary, mixing stochastic process. Another branch of learning theory does avoid making any probabilistic assumptions, getting results which hold universally across all possible data sets, and again these results apply to Bayesian updating, at least over some parameter spaces (Cesa-Bianchi and Lugosi, 2006). However, these results are all of the form “in retrospect, the posterior predictive distribution will have predicted almost as well as the best individual model could have done,” speaking entirely about performance on the past training data and revealing nothing about extrapolation to so-far unobserved cases.

To sum up, one is free to describe statistical inference as a theory of inductive logic, but these would be inductions which are deductively guaranteed by the probabilistic assumptions of stochastic models. We can see no interesting and correct sense in which Bayesian statistics is a logic of induction which does not equally imply that frequentist statistics is also a theory of inductive inference (cf. Mayo and Cox, 2006), which is to say, not very inductive at all.

6 What About Popper and Kuhn?

The two most famous philosophers of science are undoubtedly Karl Popper (1934/1959) and Thomas Kuhn (1970), and if statisticians (like other non-philosophers) know about philosophy of science at all, it is generally some version of their ideas. It may therefore help readers for see how our ideas relate to theirs. We do not pretend that our sketch fully portrays these figures, let alone the literatures of exegesis and controversy they inspired, or even how the philosophy of science has moved on since 1970.

Popper’s key idea was that of “falsification,” or “conjectures and refutations.” The inspiring example, for Popper, was the replacement of classical physics, after several centuries as the core of the best-established science, by modern physics, especially the replacement of Newtonian gravitation by Einstein’s general relativity. Science, for Popper, advances by scientists advancing theories which make strong, wide-ranging predictions capable of being refuted by observations. A good experiment or observational study is one which tests a specific theory (or theories) by confronting their predictions with data in such a way that a match is not automatically assured; good studies are designed with theories in mind, to give them a chance to fail. Theories which conflict with any evidence must be rejected, since a single counter-example implies that a generalization is false. Theories which are not falsifiable by any conceivable evidence are, for Popper, simply not scientific, though they may have other virtues.¹⁸ Even those falsifiable theories which have survived contact with data

¹⁸This “demarcation criterion” has received a lot of criticism, much of it justified. The question of what makes something “scientific” is fortunately not one we have to answer; cf. Laudan (1996, chs. 11–12) and Ziman (2000).

so far must be regarded as more or less provisional, since no finite amount of data can ever establish a generalization, nor is there any non-circular principle of induction which could let us regard theories which are compatible with lots of evidence as probably true.¹⁹ Since people are fallible, and often obstinate and overly fond of their own ideas, the objectivity of the process which tests conjectures lies not in the emotional detachment and impartiality of individual scientists, but rather in the scientific community being organized in certain ways, with certain institutions, norms and traditions, so that individuals' prejudices more or less wash out (Popper, 1945, chs. 23–24).

Clearly, we find much here to agree with, especially the general hypothetico-deductive view of scientific method and the anti-inductivist stance. On the other hand, Popper's specific ideas about testing require, at the least, substantial modification. His idea of a test comes down to the rule of deduction which says that if p implies q , and q is false, then p must be false, with the roles of p and q being played by hypotheses and data, respectively. This is plainly inadequate for statistical hypotheses, yet, as critics have noted since Braithwaite (1953) at least, he oddly ignored the theory of statistical hypothesis testing.²⁰ It is possible to do better, both through standard hypothesis tests and the kind of predictive checks we have described. In particular, as Mayo (1996) has emphasized, it is vital to consider the *severity* of tests, their capacity to detect violations of hypotheses when they are present, since it is really only passing severe tests which provides evidence for hypotheses.

Popper tried to say how science *ought* to work, supplemented by arguments that his ideals could at least be approximated and often had been. Kuhn's work, by contrast, account, in contrast, was much more an attempt to describe how science had, in point of historical fact, developed, supported by arguments that alternatives were infeasible, from which some morals might be drawn. His central idea was that of a “paradigm,” a scientific problem and its solution which served as a model or exemplar, so that solutions to other problems could be developed in imitation of it.²¹ Paradigms come along with presuppositions about the terms available for describing problems and their solutions, what counts as a valid problem, what counts as a solution, background assumptions which can be taken as a matter of course, etc. Once a scientific community accepts a paradigm and all that goes with it, its members can communicate with one another, and get on with the business of “puzzle solving,” rather than arguing about what they should be doing. Such “normal science” includes a certain amount of developing and testing of hypotheses but leaves the central presuppositions of the paradigm unquestioned.

During periods of normal science, according to Kuhn, there will always be some “anomalies”—things within the domain of the paradigm which it currently cannot explain, or even seem to refute its assumptions. These are generally ignored, or at most regarded as problems which somebody ought to investigate eventually. (Is a special adjustment for odd local

¹⁹Popper tried to work out notions of “corroboration” and increasing truth content, or “verisimilitude,” that fit with these stances, but these are generally regarded as failures.

²⁰We have generally found Popper's ideas on probability and statistics to be of little use and will not discuss them here.

²¹Examples include Newton's deduction of Kepler's laws of planetary motion and other facts of astronomy from the inverse square law of gravitation, or Planck's derivation of the black-body radiation distribution from Boltzmann's statistical mechanics and the quantization of the electromagnetic field. An internal example for statistics might be the way the Neyman-Pearson lemma inspired the search for uniformly most powerful tests in a variety of complicated situations.

circumstances called for? Might there be some clever calculational trick which fixes things? How sound are those anomalous observations?) More formally, Kuhn invokes the “Quine-Duhem thesis” (Quine, 1961; Duhem, 1914/1954). A paradigm only makes predictions about observations in conjunction with “auxiliary” hypotheses about specific circumstances, measurement procedures, etc. If the predictions are wrong, Quine and Duhem claimed that one is always free to fix the blame on the auxiliary hypotheses, and preserve belief in the core assumptions of the paradigm “come what may.”²² The Quine-Duhem thesis was also used by Lakatos (1978) as part of his “methodology of scientific research programmes,” a falsificationism more historically oriented than Popper’s distinguishing between progressive development of auxiliary hypotheses and degenerate research programs where auxiliaries become *ad hoc* devices for saving core assumptions from data.

According to Kuhn, however, anomalies can accumulate, becoming so serious as to create a crisis for the paradigm, beginning a period of “revolutionary science.” It is then that a new paradigm can form, one which is generally “incommensurable” with the old: it makes different presuppositions, takes a different problem and its solution as exemplars, re-defines the meaning of terms. Kuhn insisted that scientists who retain the old paradigm are not being irrational, because (by Quine-Duhem) they can always explain away the anomalies *somehow*; but neither are the scientists who embrace and develop the new paradigm being irrational. Switching to the new paradigm is more like a bistable illusion flipping (the apparent duck becomes an obvious rabbit) than any process of ratiocination governed by sound rules of method.²³

In some way, Kuhn’s distinction between normal and revolutionary science is analogous to the distinction between learning within a Bayesian model, and checking the model as preparation to discard or expand it. Just as the work of normal science proceeds within the presuppositions of the paradigm, updating a posterior distribution by conditioning on new data takes the assumptions embodied in the prior distribution and the likelihood function as unchallengeable truths. Model checking, on the other hand, corresponds to the identification of anomalies, with a switch to a new model when they become intolerable. Even the problems with translations between paradigms have something of a counterpart in statistical practice; for example, the intercept coefficients in a varying-intercept, constant-slope regression model have a somewhat different meaning than do the intercepts in a varying-slope model. We do not want to push the analogy too far, however, since most model checking and model re-formulation would by Kuhn have been regarded as puzzle-solving within a single paradigm, and his views of how people switch between paradigms

²²This thesis can be attacked from many directions, perhaps the most vulnerable being that one can often find multiple lines of evidence which bear on either the main principles or the auxiliary hypotheses *separately*, thereby localizing the problems (Glymour, 1980; Kitcher, 1993; Laudan, 1996; Mayo, 1996).

²³Salmon (1990) proposed a connection between Kuhn and Bayesian reasoning, suggesting that the choice between paradigms could be made rationally by using Bayes’s rule to compute their posterior probabilities, with the prior probabilities for the paradigms encoding such things as preferences for parsimony. This has at least three big problems. First, all our earlier objections to using posterior probabilities to choose between theories apply, with all the more force because every paradigm is compatible with a broad range of specific theories. Second, devising priors encoding those methodological preferences—particularly a non-vacuous preference for parsimony—is hard to impossible (Kelly, 2010). Third, it implies a truly remarkable form of Platonism: for scientists to give a paradigm positive posterior probability, they must, by Bayes’s rule, have always given it strictly positive prior probability, *even before having encountered a statement of the paradigm*.

are, as we just saw, rather different.

Kuhn's ideas about scientific revolutions are famous because they raise so many disturbing questions about the scientific enterprise. For instance, there has been considerable controversy over whether Kuhn believed in any notion of scientific progress, and over whether or not he should have, given his theory. Yet detailed historical case studies (Donovan et al., 1988) have shown that Kuhn's picture of sharp breaks between normal and revolutionary science is hard to sustain. (Arguably this is true even of Kuhn, 1957.) The leads to a tendency, already remarked by Toulmin (1972, pp. 112–17), to either expand paradigms or to shrink them. Expanding paradigms into persistent and all-embracing, because abstract and vague, bodies of ideas lets one preserve the idea of abrupt breaks in thought, but makes them rare and leaves almost everything to puzzle-solving normal science. (In the limit, there has only been one paradigm in astronomy since the Mesopotamians, something like “Many lights in the night sky are objects which are very large but very far away, and they move in interrelated, mathematically-describable, discernible patterns.”) This corresponds, we might say, to relentlessly enlarging the support of the prior. The other alternative is to shrink paradigms into increasingly concrete, specific theories and even models, making the standard for a “revolutionary” change very small indeed, in the limit reaching any kind of conceptual change whatsoever.

We suggest that there is actually some validity to both moves, that there is a sort of (weak) self-similarity involved in scientific change. Every scale of size and complexity, from local problem solving to big-picture science, features progress of the “normal science” type, punctuated by occasional revolutions. For example, in working on an applied research or consulting problem, one typically will start in a certain direction, then suddenly realize one was thinking about it wrong, then move forward, and so forth. In a consulting setting, this reevaluation can happen several times in a couple of hours. At a slightly longer time scale, we commonly reassess any approach to an applied problem after a few months, realizing there was some key feature of the problem we were misunderstanding, and so forth. There is a link between the size and the typical time scales of these changes, with small revolutions occurring fairly frequently (every few minutes for an exam-type problem), up to every few decades for a major scientific consensus. (This is related to but somewhat different from the recursive subject-matter divisions discussed by Abbott 2001.) The big changes are more exciting, even glamorous, but they rest on the hard work of extending the implications of theories far enough that they can be decisively refuted.

To sum up, our views are much closer to Popper's than to Kuhn's. The latter encouraged a close attention to the history of science and to explaining the process of scientific change, as well as putting on the agenda many genuinely deep questions, such as when and how scientific fields achieve consensus. There are even analogies between Kuhn's ideas and what happens in good data-analytic practice. Fundamentally, however, we feel that deductive model checking is central to statistical and scientific progress, and that it is the threat of such checks that motivates us to perform inferences within complex models that we know ahead of time to be false.

7 Why does this matter?

Philosophy matters to practitioners because they use philosophy to guide their practice; even those who believe themselves quite exempt from any philosophical influences are usually the slaves of some defunct methodologist. The idea of Bayesian inference as inductive, culminating in the computation of the posterior probability of scientific hypotheses, has had malign effects on statistical practice. At best, the inductivist view has encouraged researchers to fit and compare models without checking them; at worst, theorists have actively discouraged practitioners from performing model checking because it does not fit into their framework.

In our hypothetico-deductive view of data analysis, we build a statistical model out of available parts and drive it as far as it can take us, and then a little farther. When the model breaks down, we dissect it and figure out what went wrong. For Bayesian models, the most useful way of figuring out how the model breaks down is through posterior predictive checks, creating simulations of the data and comparing them to the actual data. The comparison can often be done visually; see Gelman et al. (2003, ch. 6) for a range of examples. Once we have an idea about where the problem lies, we can tinker with the model, or perhaps try a radically new design. Either way, we are using deductive reasoning as a tool to get the most out of a model, and we test the model—it is falsifiable, and when it is consequentially falsified, we alter or abandon it. None of this is especially subjective, or at least no more so than any other kind of scientific inquiry, which likewise requires choices as to the problem to study, the data to use, the models to employ, etc.—but these choices are by no means arbitrary whims, uncontrolled by objective conditions.

Conversely, a problem with the inductive philosophy of Bayesian statistics—in which science “learns” by updating the probabilities that various competing models are true—is that it assumes that the true model (or, at least, the models among which we will choose or average over) is one of the possibilities being considered. This does not fit our own experiences of learning by finding that a model doesn’t fit and needing to expand beyond the existing class of models to fix the problem.

We fear that a philosophy of Bayesian statistics as subjective, inductive inference can encourage a complacency about picking or averaging over existing models rather than trying to falsify and go further.²⁴ Likelihood and Bayesian inference are powerful, and with great power comes great responsibility. Complex models can and should be checked and falsified. This is how we can learn from our mistakes.

Acknowledgments

We thank the National Institutes of Health, the National Security Agency, and the Department of Energy for partial support of this work and Wolfgang Beirl, Chris Genovese, Clark Glymour, Mark Handcock, Jay Kadane, Rob Kass, Kevin Kelly, Kristina Klinkner, Deborah Mayo, Martina Morris, Scott Page, Aris Spanos, Erik van Nimwegen, Larry Wasserman, and Chris Wiggins for helpful conversations over the years.

²⁴Ghosh and Ramamoorthi (2003, p. 112) see a similar attitude as discouraging enquiries into consistency: “the prior and the posterior given by Bayes theorem [sic] are imperatives arising out of axioms of rational behavior—and since we are already rational why worry about one more” criterion, namely convergence to the truth?

References

Andrew Abbott. *Chaos of Disciplines*. University of Chicago Press, Chicago, 2001.

Abu Hamid Muhammad ibn Muhammad at-Tusi al Ghazali. *The Incoherence of the Philosophers = Tahafut al-Falasifah: A Parallel English-Arabic Text*. Brigham Young University Press, Provo, Utah, 1100/1997. Translated by Michael E. Marmura.

W. Ross Ashby. *Design for a Brain: The Origins of Adaptive Behavior*. Chapman and Hall, London, 2nd edition, 1960. First edition, 1956.

A. C. Atkinson and A. N. Donev. *Optimum Experimental Designs*. Clarendon Press, Oxford, 1992.

M. S. Bartlett. Inference and stochastic processes. *Journal of the Royal Statistical Society A*, 130:457–478, 1967. URL <http://www.jstor.org/stable/2982519>.

Robert H. Berk. Limiting behavior of posterior distributions when the model is incorrect. *Annals of Mathematical Statistics*, 37:51–58, 1966. doi: 10.1214/aoms/1177699597. URL <http://projecteuclid.org/euclid.aoms/1177699597>. See also correction, volume 37 (1966), pp. 745–746.

Claude Bernard. *Introduction to the Study of Experimental Medicine*. Macmillan, New York, 1865/1927. Translated by Henry Copley Green. First published as *Introduction a l'étude de la medecine experimentale*, Paris: J. B. Bailliere. Reprinted New York: Dover, 1957.

Jose M. Bernardo and Adrian F. M. Smith. *Bayesian Theory*. Wiley, New York, 1994.

Ken Binmore. Making decisions in large worlds. Technical Report 266, ESRC Centre for Economic Learning and Social Evolution, University College of London, 2007. URL <http://else.econ.ucl.ac.uk/papers/uploaded/266.pdf>.

Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to statistical learning theory. In Olivier Bousquet, Ulrike von Luxburg, and Gunnar Rätsch, editors, *Advanced Lectures in Machine Learning*, pages 169–207, Berlin, 2004. Springer-Verlag. URL http://www.econ.upf.edu/~lugosi/mlss_slt.pdf.

R. B. Braithwaite. *Scientific Explanation: A Study of the Function of Theory, Probability and Law in Science*. Cambridge University Press, Cambridge, England, 1953.

R. Z. Brown, W. Sallow, D. E. Davis, and W. G. Cochran. The rat population of Baltimore, 1952. *American Journal of Epidemiology*, 61:89–102, 1955.

Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, England, 2006.

Gerda Claeskens and Nils Lid Hjort. *Model Selection and Model Averaging*. Cambridge University Press, Cambridge, England, 2008.

D. R. Cox and D. V. Hinkley. *Theoretical Statistics*. Chapman and Hall, London, 1974.

Richard T. Cox. Probability, frequency, and reasonable expectation. *American Journal of Physics*, 14:1–13, 1946.

Richard T. Cox. *The Algebra of Probable Inference*. Johns Hopkins University Press, Baltimore, 1961.

Imre Csiszár. Maxent, mathematics, and information theory. In Kenneth M. Hanson and Richard N. Silver, editors, *Maximum Entropy and Bayesian Methods: Proceedings of the Fifteenth International Workshop on Maximum Entropy and Bayesian Methods*, pages 35–50, Dordrecht, 1995. Kluwer Academic.

A. Philip Dawid and Vladimir G. Vovk. Prequential probability: principles and properties. *Bernoulli*, 5:125–162, 1999. URL <http://projecteuclid.org/euclid.bj/1173707098>.

Arthur Donovan, Larry Laudan, and Rachel Laudan, editors. *Scrutinizing Science: Empirical Studies of Scientific Change*. Kluwer Academic, Dordrecht, 1988. Reprinted 1992 (Baltimore: Johns Hopkins University Press) with a new introduction.

Joseph L. Doob. Application of the theory of martingales. In *Colloques Internationaux du Centre National de la Recherche Scientifique*, volume 13, pages 23–27, Paris, 1949. Centre National de la Recherche Scientifique.

Pierre Duhem. *The Aim and Structure of Physical Theory*. Princeton University Press, Princeton, New Jersey, 1914/1954. Translated by Philip P. Wiener from the second edition *La théorie physique, son objet et sa structure*, Paris: Chevalier et Rivière.

John Earman. *Bayes or Bust? A Critical Account of Bayesian Confirmation Theory*. MIT Press, Cambridge, Massachusetts, 1992.

Branden Fitelson and Neil Thomason. Bayesians sometimes cannot ignore even very implausible theories (even ones that have not yet been thought of). *Australasian Journal of Logic*, 6:25–36, 2008. URL <http://fitelson.org/hiti.pdf>.

Dean P. Foster and H. Peyton Young. Learning, hypothesis testing and Nash equilibrium. *Games and Economic Behavior*, 45:73–96, 2003. URL <http://www.econ.jhu.edu/people/young/nash.pdf>.

Andrew Gelman. Treatment effects in before-after data. In Andrew Gelman and X. L. Meng, editors, *Applied Bayesian Modeling and Causal Inference from an Incomplete Data Perspective*, chapter 18, pages 191–198. Wiley, London, 2004. URL <http://www.stat.columbia.edu/~gelman/research/published/gelman.pdf>.

Andrew Gelman. A Bayesian formulation of exploratory data analysis and goodness-of-fit testing. *International Statistical Review*, 71:369–382, 2003. URL <http://www.stat.columbia.edu/~gelman/research/published/isr.pdf>.

Andrew Gelman. Discussion of “a probabilistic model for the spatial distribution of party support in multiparty elections” by s. merrill. *Journal of the American Statistical Association*, 89:1198, 1994.

Andrew Gelman and Jennifer Hill. *Data Analysis Using Regression and Multi-level/Hierarchical Models*. Cambridge University Press, Cambridge, England, 2006.

Andrew Gelman and Gary King. Enhancing democracy through legislative redistricting. *American Political Science Review*, 88:541–559, 1994.

Andrew Gelman and Donald B. Rubin. Avoiding model selection in Bayesian social research. *Sociological Methodology*, 25:165–173, 1995. URL <http://www.stat.columbia.edu/~gelman/research/published/avoiding.pdf>.

Andrew Gelman, Xiao-Li Meng, and Hal S. Stern. Posterior predictive assessment of model fitness via realized discrepancies (with discussion). *Statistica Sinica*, 6:733–807, 1996.

Andrew Gelman, John B. Carlin, Hal S Stern, and Donald B. Rubin. *Bayesian Data Analysis*. CRC Press, London, second edition, 2003.

Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, and Yu-Sung Su. A weakly informative default prior distribution for logistic and other regression models. *Annals of Applied Statistics*, 2:1360–1383, 2008a. URL <http://arxiv.org/abs/0901.4011>.

Andrew Gelman, David Park, Boris Shor, Joseph Bafumi, and Jeronimo Cortina. *Red State, Blue State, Rich State, Poor State: Why Americans Vote the Way They Do*. Princeton University Press, Princeton, New Jersey, 2008b.

Andrew Gelman, Boris Shor, David Park, and Joseph Bafumi. Rich state, poor state, red state, blue state: What’s the matter with Connecticut? *Quarterly Journal of Political Science*, 2:345–367, 2008c. URL http://redbluerichpoor.com/media/red_state_blue_state.pdf.

Andrew Gelman, Daniel Lee, and Yair Ghitza. Public opinion on health care reform. *The Forum*, 8(1), 2010. doi: 10.2202/1540-8884.1355. URL <http://www.bepress.com/forum/vol8/iss1/art8>.

J. K. Ghosh and R. V. Ramamoorthi. *Bayesian Nonparametrics*. Springer Verlag, New York, 2003.

Ronald N. Giere. *Explaining Science: A Cognitive Approach*. University of Chicago Press, Chicago, 1988.

Clark Glymour. *Theory and Evidence*. Princeton University Press, Princeton, New Jersey, 1980.

Robert M. Gray. *Entropy and Information Theory*. Springer-Verlag, New York, 1990. URL <http://ee.stanford.edu/~gray/it.html>.

S. Greenland. Induction versus Popper: substance versus semantics. *International Journal of Epidemiology*, 27:543–548, 1998.

Peter D. Grünwald. *The Minimum Description Length Principle*. MIT Press, Cambridge, Massachusetts, 2007.

Peter D. Grünwald and John Langford. Suboptimal behavior of Bayes and MDL in classification under misspecification. *Machine Learning*, 66:119–149, 2007. doi: 10.1007/s10994-007-0716-7. URL <http://www.cwi.nl/~pdg/ftp/inconsistency.pdf>.

Peter Guttorp. *Stochastic Modeling of Scientific Data*. Chapman and Hall, London, 1995.

Joseph Y. Halpern. Cox’s theorem revisited. *Journal of Artificial Intelligence Research*, 11: 429–435, 1999.

Mark S. Handcock. Assessing degeneracy in statistical models of social networks. Technical Report 39, Center for Statistics and the Social Sciences, University of Washington, 2003. URL <http://csde.washington.edu/statnet/www.csss.washington.edu/Papers/wp39.pdf>.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, New York, 2001.

Carl G. Hempel. *Aspects of Scientific Explanation*. The Free Press, Glencoe, Illinois, 1965.

John H. Holland, Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, Massachusetts, 1986.

Colin Howson and Peter Urbach. *Scientific Reasoning: The Bayesian Approach*. Open Court, La Salle, Illinois, 1989.

David R. Hunter, Steven M. Goodreau, and Mark S. Handcock. Goodness of fit of social network models. *Journal of the American Statistical Association*, 103:248–258, 2008. doi: 10.1198/016214507000000446. URL <http://www.csss.washington.edu/Papers/wp47.pdf>.

E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, England, 2003.

Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90:773–795, 1995. URL <http://www.stat.cmu.edu/~kass/papers/bayesfactors.pdf>.

Robert E. Kass and Paul W. Vos. *Geometrical Foundations of Asymptotic Inference*. Wiley, New York, 1997.

Robert E. Kass and Larry Wasserman. The selection of prior distributions by formal rules. *Journal of the American Statistical Association*, 91:1343–1370, 1996. URL <http://www.stat.cmu.edu/~kass/papers/rules.pdf>.

Kevin T. Kelly. *The Logic of Reliable Inquiry*. Oxford University Press, Oxford, 1996.

Kevin T. Kelly. Simplicity, truth, and probability. In Prasanta Bandyopadhyay and Malcolm Forster, editors, *Handbook on the Philosophy of Statistics*. Elsevier, Dordrecht, 2010. URL <http://www.andrew.cmu.edu/user/kk3n/ockham/prasanta-submit-final.pdf>.

Philip Kitcher. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford University Press, Oxford, 1993.

B. J. K. Kleijn and A. W. van der Vaart. Misspecification in infinite-dimensional Bayesian statistics. *Annals of Statistics*, 34:837–877, 2006. URL <http://arxiv.org/math.ST/0607023>.

Leszek Kolakowski. *The Alienation of Reason: A History of Positivist Thought*. Doubleday, Garden City, New York, 1968. Translated by Norbert Guterman from the Polish *Filozofia Pozytywistyczna (od Hume 'a do Kola Wiedenskiego)*, Panstwowe Wydawnictwo Naukowe, 1966.

Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation, and active learning. In Gerald Tesauro, David Touretzky, and Todd Leen, editors, *Advances in Neural Information Processing 7 [NIPS 1994]*, pages 231–238, Cambridge, Massachusetts, 1995. MIT Press. URL <http://books.nips.cc/papers/files/nips07/0231.pdf>.

Thomas S. Kuhn. *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought*. Harvard University Press, Cambridge, Massachusetts, 1957.

Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, second edition, 1970.

Imre Lakatos. *Philosophical Papers*. Cambridge University Press, Cambridge, England, 1978.

Larry Laudan. *Beyond Positivism and Relativism: Theory, Method and Evidence*. Westview Press, Boulder, Colorado, 1996.

Larry Laudan. *Science and Hypothesis*. D. Reidel, Dordrecht, 1981.

Qi Li and Jeffrey Scott Racine. *Nonparametric Econometrics: Theory and Practice*. Princeton University Press, Princeton, New Jersey, 2007.

Antonio Lijoi, Igor Prünster, and Stephen G. Walker. Bayesian consistency for stationary models. *Econometric Theory*, 23:749–759, 2007. doi: 10.1017/S0266466607070314.

Bruce Lindsay and Liwei Liu. Model assessment tools for a model false world. *Statistical Science*, 24:303–318, 2009. URL <http://projecteuclid.org/euclid.ss/1270041257>.

Charles F. Manski. *Identification for Prediction and Decision*. Harvard University Press, Cambridge, Massachusetts, 2007.

Deborah G. Mayo. *Error and the Growth of Experimental Knowledge*. University of Chicago Press, Chicago, 1996.

Deborah G. Mayo and D. R. Cox. Frequentist statistics as a theory of inductive inference. In Javier Rojo, editor, *Optimality: The Second Erich L. Lehmann Symposium*, pages 77–97, Bethesda, Maryland, 2006. Institute of Mathematical Statistics. URL <http://arxiv.org/abs/math.ST/0610846>.

Deborah G. Mayo and Aris Spanos. Methodology in practice: Statistical misspecification testing. *Philosophy of Science*, 71:1007–1025, 2004. URL <http://www.error06.econ.vt.edu/MayoSpanos2004.pdf>.

Deborah G. Mayo and Aris Spanos. Severe testing as a basic concept in a neyman-pearson philosophy of induction. *The British Journal for the Philosophy of Science*, 57:323–357, 2006. doi: 10.1093/bjps/axl003.

David A. McAllister. Some PAC-Bayesian theorems. *Machine Learning*, 37:355–363, 1999.

Nolan McCarty, Keith T. Poole, and Howard Rosenthal. *Polarized America: The Dance of Ideology and Unequal Riches*. Walras-Pareto Lectures. MIT Press, Cambridge, Massachusetts, 2006.

N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

Ulrich K. Müller. Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix. Electronic pre-print, 2010. URL <http://www.princeton.edu/~umueller/sandwich.pdf>.

Mark E. J. Newman and G. T. Barkema. *Monte Carlo Methods in Statistical Physics*. Clarendon Press, Oxford, 1999.

John D. Norton. A material theory of induction. *Philosophy of Science*, 70:647–670, 2003. URL <http://www.pitt.edu/~jdnorton/papers/material.pdf>.

Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, Princeton, New Jersey, 2007.

Liam Paninski. Asymptotic theory of information-theoretic experimental design. *Neural Computation*, 17:1480–1507, 2005. URL <http://www.stat.columbia.edu/~liam/research/abstracts/doe-nc-abs.html>.

Karl R. Popper. *The Logic of Scientific Discovery*. Hutchinson, London, 1934/1959. Translated by the author from *Logik der Forschung* (Vienna: Julius Springer Verlag).

Karl R. Popper. *The Open Society and Its Enemies*. Routledge, London, 1945.

Willard Van Orman Quine. *From a Logical Point of View: Logico-Philosophical Essays*. Harvard University Press, Cambridge, Mass., second edition, 1961. First edition, 1953.

Adrian E. Raftery. Bayesian model selection in social research. *Sociological Methodology*, 25:111–196, 1995. URL <http://www.stat.washington.edu/raftery/Research/PDF/socmeth1995.pdf>.

Brian D. Ripley. *Statistical Inference for Spatial Processes*. Cambridge University Press, Cambridge, England, 1988.

Douglas Rivers and Quang H. Vuong. Model selection tests for nonlinear dynamic models. *The Econometrics Journal*, 5:1–39, 2002.

Donald B. Rubin. Bayesianly justifiable and relevant frequency calculations for the applied statistician. *Annals of Statistics*, 12:1151–1172, 1984. URL <http://projecteuclid.org/euclid-aos/1176346785>.

Bertrand Russell. *Human Knowledge: Its Scope and Limits*. Simon and Schuster, New York, 1948.

Wesley C. Salmon. The appraisal of theories: Kuhn meets Bayes. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1990:325–332, 1990. URL <http://www.jstor.org/stable/193077>.

Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

Mark J. Schervish. *Theory of Statistics*. Springer-Verlag, Berlin, 1995.

Teddy Seidenfeld. Entropy and uncertainty. In I. B. MacNeill and G. J. Umphrey, editors, *Foundations of Statistical Inference*, pages 259–287, Dordrecht, 1987. D. Reidel. URL [http://www.hss.cmu.edu/philosophy/seidenfeld/relating%20to%20other%20probability%20and%20statistical%20issues/Entropy%20and%20Uncertainty%20\(revised\).pdf](http://www.hss.cmu.edu/philosophy/seidenfeld/relating%20to%20other%20probability%20and%20statistical%20issues/Entropy%20and%20Uncertainty%20(revised).pdf).

Teddy Seidenfeld. Why I am not an objective Bayesian: Some reflections prompted by Rosenkrantz. *Theory and Decision*, 11:413–440, 1979. URL <http://www.hss.cmu.edu/philosophy/seidenfeld/relating%20to%20other%20probability%20and%20statistical%20issues/Why%20I%20Am%20Not%20an%20Objective%20B.pdf>.

Cosma Rohilla Shalizi. Dynamics of Bayesian updating with dependent data and misspecified models. *Electronic Journal of Statistics*, 3:1039–1074, 2009. doi: 10.1214/09-EJS485. URL <http://arxiv.org/abs/0901.1342>.

Tom A. B. Snijders, Philippa E. Pattison, Garry L. Robins, and Mark S. Handcock. New specifications for exponential random graph models. *Sociological Methodology*, 36:99–153, 2006. doi: 10.1111/j.1467-9531.2006.00176.x. URL <http://www.csss.washington.edu/Papers/wp42.pdf>.

Aris Spanos. Curve fitting, the reliability of inductive inference, and the error-statistical approach. *Philosophy of Science*, 74:1046–1066, 2007. doi: 10.1086/525643.

David C. Stove. *The Rationality of Induction*. Clarendon Press, Oxford, 1986.

David C. Stove. *Popper and After: Four Modern Irrationalists*. Pergamon Press, Oxford, 1982.

Charles Tilly. *Explaining Social Processes*. Paradigm Publishers, Boulder, Colorado, 2008.

Charles Tilly. Observations of social processes and their formal representations. *Sociological Theory*, 22:595–602, 2004. doi: 10.1111/j.0735-2751.2004.00235.x. URL http://professor-murmann.info/tilly/2004_Obs_of_soc_proc.pdf. Reprinted in Tilly (2008).

Stephen Toulmin. *Human Understanding: The Collective Use and Evolution of Concepts*. Princeton University Press, Princeton, New Jersey, 1972.

Jos Uffink. The constraint rule of the maximum entropy principle. *Studies in History and Philosophy of Modern Physics*, 27:47–79, 1996. URL <http://www.phys.uu.nl/~wwwgrnsl/jos/mep2def/mep2def.html>.

Jos Uffink. Can the maximum entropy principle be explained as a consistency requirement? *Studies in History and Philosophy of Modern Physics*, 26B:223–261, 1995. URL <http://www.phys.uu.nl/~wwwgrnsl/jos/mepabst/mepabst.html>.

M. Vidyasagar. *Learning and Generalization: With Applications to Neural Networks*. Springer-Verlag, Berlin, second edition, 2003.

Quang H. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57:307–333, 1989. URL <http://www.jstor.org/pss/1912557>.

Grace Wahba. *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics, Philadelphia, 1990.

Larry Wasserman. Frequentist Bayes is objective. *Bayesian Analysis*, 1:451–456, 2006. URL <http://ba.stat.cmu.edu/journal/2006/vol01/issue03/wasserman.pdf>.

Steven Weinberg. What is quantum field theory, and what did we think it was? In Tian Yu Cao, editor, *Conceptual Foundations of Quantum Field Theory*, pages 241–251, Cambridge, England, 1999. Cambridge University Press. URL <http://arxiv.org/abs/hep-th/9702027>.

Halbert White. *Estimation, Inference and Specification Analysis*. Cambridge University Press, Cambridge, England, 1994.

John Ziman. *Real Science: What It Is, and What It Means*. Cambridge University Press, Cambridge, England, 2000.