

Dynamics of Bayesian updating with dependent data and misspecified models

Cosma Rohilla Shalizi

*Statistics Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA*

15213-3890 USA e-mail: cshalizi@cmu.edu

Abstract: Much is now known about the consistency of Bayesian updating on infinite-dimensional parameter spaces with independent or Markovian data. Necessary conditions for consistency include the prior putting enough weight on the correct neighborhoods of the data-generating distribution; various sufficient conditions further restrict the prior in ways analogous to capacity control in frequentist nonparametrics. The asymptotics of Bayesian updating with mis-specified models or priors, or non-Markovian data, are far less well explored. Here I establish sufficient conditions for posterior convergence when all hypotheses are wrong, and the data have complex dependencies. The main dynamical assumption is the asymptotic equipartition (Shannon-McMillan-Breiman) property of information theory. This, along with Egorov’s Theorem on uniform convergence, lets me build a sieve-like structure for the prior. The main statistical assumption, also a form of capacity control, concerns the compatibility of the prior and the data-generating process, controlling the fluctuations in the log-likelihood when averaged over the sieve-like sets. In addition to posterior convergence, I derive a kind of large deviations principle for the posterior measure, extending in some cases to rates of convergence, and discuss the advantages of predicting using a combination of models known to be wrong. An appendix sketches connections between these results and the replicator dynamics of evolutionary theory.

AMS 2000 subject classifications: Primary 62C10, 62G20, 62M09; secondary 60F10, 62M05, 92D15, 94A17.

Keywords and phrases: Asymptotic equipartition, Bayesian consistency, Bayesian nonparametrics, Egorov’s theorem, large deviations, posterior convergence, replicator dynamics, sofic systems.

Received September 2009.

1. Introduction

The problem of the convergence and frequentist consistency of Bayesian learning goes as follows. We encounter observations X_1, X_2, \dots , which we would like to predict by means of a family Θ of models or hypotheses (indexed by θ). We begin with a prior probability distribution Π_0 over Θ , and update this using Bayes’s rule, so that our distribution after seeing $X_1, X_2, \dots, X_t \equiv X_1^t$ is Π_t . If the observations come from a stochastic process with infinite-dimensional distribution P , when does Π_t converge P -almost surely? What is the rate of

convergence? Under what conditions will Bayesian learning be consistent, so that Π_t doesn't just converge but its limit is P ?

Since the Bayesian estimate is really the whole posterior probability distribution Π_t rather than a point or set in Θ , consistency becomes concentration of Π_t around P . One defines some sufficiently strong set of neighborhoods of P in the space of probability distributions on X_1^∞ , and says that Π_t is consistent when, for each such neighborhood N , $\lim_{t \rightarrow \infty} \Pi_t N = 1$. When this holds, the posterior increasingly approximates a degenerate distribution centered at the truth.

The greatest importance of these problems, perhaps, is their bearing on the objectivity and reliability of Bayesian inference; consistency proofs and convergence rates are, as it were, frequentist licenses for Bayesian practices. Moreover, if Bayesian learners starting from different priors converge rapidly on the same posterior distribution, there is less reason to worry about the subjective or arbitrary element in the choice of priors. (Such “merger of opinion” results [7] are also important in economics and game theory [11].) Recent years have seen considerable work on these problems, especially in the non-parametric setting where the model space Θ is infinite-dimensional [30].

Pioneering work by Doob [19], using elegant martingale arguments, established that when any consistent estimator exists, and P lies in the support of Π_0 , the set of sample paths on which the Bayesian learner fails to converge to the truth has prior probability zero. (See [13] and [42] for extensions of this result to non-IID settings, and also the discussion in [55, 21].) This is not, however, totally reassuring, since P generally also has prior probability zero, and it would be unfortunate if these two measure-zero sets should happen to coincide. Indeed, Diaconis and Freedman established that the consistency of Bayesian inference depends crucially on the choice of prior, and that even very natural priors can lead to inconsistency (see [18] and references therein).

Subsequent work, following a path established by Schwartz [56], has shown that, no matter what the true data-generating distribution P , Bayesian updating converges along P -almost-all sample paths, provided that (a) P is contained in Θ , (b) every Kullback-Leibler neighborhood in the Θ has some positive prior probability (the “Kullback-Leibler property”), and (c) certain restrictions hold on the prior, amounting to versions of capacity control, as in the method of sieves or structural risk minimization. These contributions also make (d) certain dynamical assumptions about the data-generating process, most often that it is IID [4, 26, 65] (in this setting, [27] and [59] in particular consider convergence rates), independent non-identically distributed [13, 29], or, in some cases, Markovian [28, 29]; [14] and [52] work with spectral density estimation and sequential analysis, respectively, again exploiting specific dynamical properties.

For mis-specified models, that is settings where (a) above fails, important early results were obtained by Berk [5, 6] for IID data, albeit under rather strong restrictions on likelihood functions and parameter spaces, showing that the posterior distribution concentrates on an “asymptotic carrier”, consisting of the hypotheses which are the best available approximations, in the Kullback-Leibler sense, to P within the support of the prior. More recently, [38], [68] and

[41] have dealt with the convergence of non-parametric Bayesian estimation for IID data when P is not in the support of the prior, obtaining results similar to Berk's in far more general settings, extending in some situations to rates of convergence. All of this work, however, relies on the dynamical assumption of an IID data-source.

This paper gives sufficient conditions for the convergence of the posterior without assuming (a), and substantially weakening (c) and (d). Even if one uses non-parametric models, cases where one knows that the true data generating process is exactly represented by one of the hypotheses in the model class are scarce. Moreover, while IID data can be produced, with some trouble and expense, in the laboratory or in a well-conducted survey, in many applications the data are not just heterogeneous and dependent, but their heterogeneity and dependence is precisely what is of interest. This raises the question of what Bayesian updating does when the truth is not contained in the support of the prior, and observations have complicated dependencies.

To answer this question, I first weaken the dynamical assumptions to the asymptotic equipartition property (Shannon-McMillan-Breiman theorem) of information theory, i.e., for each hypothesis θ , the log-likelihood per unit time converges almost surely. This log-likelihood per unit time is basically the growth rate of the Kullback-Leibler divergence between P and θ , $h(\theta)$. As observations accumulate, areas of Θ where $h(\theta)$ exceeds its essential infimum $h(\Theta)$ tend to lose posterior probability, which concentrates in divergence-minimizing regions. Some additional conditions on the prior distribution are needed to prevent it from putting too much weight initially on hypotheses with high divergence rates but slow convergence of the log-likelihood. As the latter assumptions are strengthened, more and more can be said about the convergence of the posterior.

Using the weakest set of conditions (Assumptions 1–3), the long-run exponential growth rate of the posterior density at θ cannot exceed $h(\Theta) - h(\theta)$ (Theorem 1). Adding Assumptions 4–6 to provide better control over the integrated or marginal likelihood establishes (Theorem 2) that the long-run growth rate of the posterior density is in fact $h(\Theta) - h(\theta)$. One more assumption (7) then lets us conclude (Theorem 3) that the posterior distribution converges, in the sense that, for any set of hypotheses A , the posterior probability $\Pi_t(A) \rightarrow 0$ unless the essential infimum of $h(\theta)$ over A equals $h(\Theta)$. In fact, we then have a kind of large deviations principle for the posterior measure (Theorem 4), as well as a bound on the generalization ability of the posterior predictive distribution (Theorem 5). Convergence rates for the posterior (Theorem 6) follow from the combination of the large deviations result with an extra condition related to assumption 6. Importantly, Assumptions 4–7, and so the results following from them, involve *both* the prior distribution and the data-generating process, and require the former to be adapted to the latter. Under mis-specification, it does not seem to be possible to guarantee posterior convergence by conditions on the prior alone, at least not with the techniques used here.

For the convenience of reader, the development uses the usual statistical vocabulary and machinery. In addition to the asymptotic equipartition property, the main technical tools are on the one hand Egorov's theorem from basic mea-

sure theory, which is used to construct a sieve-like sequence of sets on which log-likelihood ratios converge uniformly, and on the other hand Assumption 6 bounding how long averages over these sets can remain far from their long-run limits. The latter assumption is crucial, novel, and, in its present form, awkward to check; I take up its relation to more familiar assumptions in the discussion. It may be of interest, however, that the results were first found via an apparently-novel analogy between Bayesian updating and the “replicator equation” of evolutionary dynamics, which is a formalization of the Darwinian idea of natural selection. Individual hypotheses play the role of distinct replicators in a population, the posterior distribution being the population distribution over replicators and fitness being proportional to likelihood. Appendix A gives details.

2. Preliminaries and notation

Let (Ω, \mathcal{F}, P) be a probability space, and X_1, X_2, \dots , for short X_1^∞ , be a sequence of random variables, taking values in the measurable space (Ξ, \mathcal{X}) , whose infinite-dimensional distribution is P . The natural filtration of this process is $\sigma(X_1^t)$. The only dynamical properties are those required for the Shannon-McMillan-Breiman theorem (Assumption 3); more specific assumptions such as P being a product measure, Markovian, exchangeable, etc., are not required. Unless otherwise noted, all probabilities are taken with respect to P , and $\mathbf{E}[\cdot]$ always means expectation under that distribution.

Statistical hypotheses, i.e., distributions of processes adapted to $\sigma(X_1^t)$, are denoted by F_θ , the index θ taking values in the hypothesis space, a measurable space (Θ, \mathcal{T}) , generally infinite-dimensional. For convenience, assume that P and all the F_θ are dominated by a common reference measure, with respective densities p and f_θ . I do not assume that $P \in \Theta$, still less that $P \in \text{supp } \Pi_0$ — i.e., quite possibly all of the available hypotheses are false.

We will study the evolution of a sequence of probability measures Π_t on (Θ, \mathcal{T}) , starting with a *non-random* prior measure Π_0 . (A filtration on Θ is not needed; the measures Π_t change but not the σ -field \mathcal{T} .) Assume all Π_t are absolutely continuous with respect to a common reference measure, with densities π_t . Expectations with respect to these measures will be written either as explicit integrals or de Finetti style, $\Pi_t(f) = \int f(\theta)d\Pi_t(\theta)$; when A is a set, $\Pi_t(fA) = \Pi_t(f\mathbf{1}_A) = \int_A f(\theta)d\Pi_t(\theta)$.

Let $L_t(\theta)$ be the conditional likelihood of x_t under θ , i.e., $L_t(\theta) \equiv f_\theta(X_t = x_t | X_1^{t-1} = x_1^{t-1})$, with $L_0 = 1$. The integrated conditional likelihood is $\Pi_t(L_t)$. Bayesian updating of course means that, for any $A \in \mathcal{T}$,

$$\Pi_{t+1}(A) = \frac{\Pi_t(L_{t+1}A)}{\Pi_t(L_{t+1})}$$

or, in terms of the density,

$$\pi_{t+1}(\theta) = \frac{L_{t+1}(\theta)\pi_t(\theta)}{\Pi_t(L_{t+1})}$$

It will also be convenient to express Bayesian updating in terms of the prior and the total likelihood:

$$\Pi_t(A) = \frac{\int_A d\Pi_0(\theta) f_\theta(x_1^t)}{\int_\Theta d\Pi_0(\theta) f_\theta(x_1^t)} = \frac{\int_A d\Pi_0(\theta) \frac{f_\theta(x_1^t)}{p(x_1^t)}}{\int_\Theta d\Pi_0(\theta) \frac{f_\theta(x_1^t)}{p(x_1^t)}} = \frac{\Pi_0(R_t A)}{\Pi_0(R_t)}$$

where $R_t(\theta) \equiv \frac{f_\theta(x_1^t)}{p(x_1^t)}$ is the ratio of model likelihood to true likelihood. (Note that $0 < p(x_1^t) < \infty$ for all t P -a.s.) Similarly,

$$\pi_t(\theta) = \pi_0(\theta) \frac{R_t(\theta)}{\Pi_0(R_t)}$$

The one-step-ahead predictive distribution of the hypothesis θ is given by $F_\theta(X_t | \sigma(X_1^{t-1}))$, with the convention that $t = 1$ gives the marginal distribution of the first observation. Abbreviate this by F_θ^t . Similarly, let $P^t \equiv P(X_t | \sigma(X_1^{t-1}))$; this is the best probabilistic prediction we could make, did we but know P [39]. The posterior predictive distribution is given by mixing the individual predictive distributions with weights given by the posterior:

$$F_\Pi^t \equiv \int_\Theta F_\theta^t d\Pi_t(\theta)$$

Remark on the topology of Θ and on \mathcal{T} The hope in studying posterior convergence is to show that, as t grows, with higher and higher (P) probability, Π_t concentrates more and more on sets which come closer and closer to P . The tricky part here is “closer and closer”: points in Θ represent infinite-dimensional stochastic process distributions, and the topology of such spaces is somewhat odd, and irritatingly abrupt, at least under the more common distances. Any two ergodic measures are either equal or have completely disjoint supports [31], so that the Kullback-Leibler divergence between distinct ergodic processes is always infinity (in both directions), and the total variation and Hellinger distances are likewise maximal. Most previous work on posterior consistency has restricted itself to models where the infinite-dimensional process distributions are formed by products of fixed-dimensional base distributions (IID, Markov, etc.), and in effect transferred the usual metrics’ topologies from these finite-dimensional distributions to the processes. It *is* possible to define metrics for general stochastic processes [31], and if readers like they may imagine that \mathcal{T} is a Borel σ -field under some such metric. This is not necessary for the results presented here, however.

2.1. Example

The following example will be used to illustrate the assumptions (§2.2.1 and Appendix B), and, later, the conclusions (§3.6).

The data-generating process P is a stationary and ergodic measure on the space of binary sequences, i.e., $\Xi = \{0, 1\}$, and the σ -field $\mathcal{X} = 2^\Xi$. The measure




FIG 1. State-transition diagram for the “even process”. The legends on the transition arrows indicate the probability of making the transition, and the observation which occurs when the transition happens. The observation $X_t = 1$ when entering or leaving state 1, otherwise it is 0. This creates blocks of 1s of even length, separated by blocks of 0s of arbitrary length. The result is a finite-state process which is not a Markov chain of any order.

is naturally represented as a function of a two-state Markov chain S_1^∞ , with $S_t \in \{1, 2\}$. The transition matrix is

$$\mathbf{T} = \begin{bmatrix} 0.0 & 1.0 \\ 0.5 & 0.5 \end{bmatrix}$$

so that the invariant distribution puts probability 1/3 on state 1 and probability 2/3 on state 2; take S_1 to be distributed accordingly. The observed process is a binary function of the latent state transitions, $X_t = 0$ if $S_t = S_{t+1} = 2$ and $X_t = 1$ otherwise. Figure 1 depicts the transition and observation structure. Qualitatively, X_1^∞ consists of blocks of 1s of even length, separated by blocks of 0s of arbitrary length. Since the joint process $\{(S_t, X_t)\}_{1 \leq t \leq \infty}$ is a stationary and ergodic Markov chain, X_1^∞ is also stationary, ergodic and mixing.

This stochastic process comes from symbolic dynamics [43, 37], where it is known as the “even process”, and serves as a basic example of the class of *sofic* processes [66], which have finite Markovian representations, as in Figure 1, but are not Markov at any finite order. (If $X_t = 1, X_{t-1} = 1, \dots, X_{t-k} = 1$ for any finite k , the corresponding S_{t-i} must have alternated between one and two, but whether S_t is one or two, and thus the distribution of X_{t+1} , cannot be determined from the length- k history alone.) More exactly [36], sofic systems or “finitary measures” are ones which are images of Markov chains under factor maps, and strictly sofic systems, such as the even process, are sofic systems which are not themselves Markov chains of any order. Despite their simplicity, these models arise naturally when studying the time series of chaotic dynamical systems [3, 15, 57, 16], as well as problems in statistical mechanics [50] and crystallography [62].

Let Θ_k be the space of all binary Markov chains of order k with strictly positive transition probabilities and their respective stationary distributions; each Θ_k has dimension 2^k . (Allowing some transition probabilities to be zero creates uninteresting technical difficulties.) Since each hypothesis is equivalent to a function $\Xi^{k+1} \mapsto (0, 1]$, we can give Θ_k the topology of pointwise convergence of functions, and the corresponding Borel σ -field. We will take $\Theta = \bigcup_{k=1}^\infty \Theta_k$,

identifying Θ_k with the appropriate subset of Θ_{k+1} . Thus Θ consists of all strictly-positive stationary binary Markov chains, of whatever order, and is infinite-dimensional.

As for the prior Π_0 , it will be specified in more detail below (§2.2.1). At the very least, however, it needs to have the “Kullback-Leibler rate property”, i.e., to give positive probability to every ϵ “neighborhood” $N_\epsilon(\theta)$ around every $\theta \in \Theta$, i.e., the set of hypotheses whose Kullback-Leibler divergence from θ grows no faster than ϵ :

$$N_\epsilon(\theta) = \left\{ \theta' : \epsilon \geq \lim_{t \rightarrow \infty} \frac{1}{t} \int dx_1^t f_\theta(x_1^t) \log \frac{f_\theta(x_1^t)}{f_{\theta'}(x_1^t)} \right\}$$

(The limit exists for all θ, θ' combinations [32].)

This example is simple, but it is also beyond the scope of existing work on Bayesian convergence in several ways. First, the data-generating process P is not even Markov. Second, $P \notin \Theta$, so all the hypotheses are wrong, and the truth is certainly not in the support of the prior. (P can however be approximated arbitrarily closely, in various process metrics, by distributions from Θ .) Third, because P is ergodic, and ergodic distributions are extreme points in the space of stationary distributions [20], it cannot be represented as a mixture of distributions in Θ . This means that the Doob-style theorem of Ref. [42] does not apply, and even the subjective certainty of convergence is not assured. The results of Refs. [38, 68, 5, 6] on mis-specified models do not hold because the data are dependent. To be as concrete and explicit as possible, the analysis here will focus on the even process, but only the constants would change if P were any other strictly sofic process. Much of it would apply even if P were a stochastic context-free language or pushdown automaton [12], where in effect the number of hidden states is infinite, though some of the details in Appendix B would change.

Ref. [47] describes a non-parametric procedure which will adaptively learn to predict a class of discrete stochastic processes which includes the even process. Ref. [58] introduces a frequentist algorithm which consistently reconstructs the hidden-state representation of sofic processes, including the even process. Ref. [61] considers Bayesian estimation of the even process, using Dirichlet priors for finite-order Markov chains, and employing Bayes factors to decide which order of chain to use for prediction.

2.2. Assumptions

The needed assumptions have to do with the dynamical properties of the data generating process P , and with how well the dynamics meshes both with the class of hypotheses Θ and with the prior distribution Π_0 over those hypotheses.

Assumption 1. *The likelihood ratio $R_t(\theta)$ is $\sigma(X_1^t) \times \mathcal{T}$ -measurable for all t .*

The next two assumptions actually need only hold for Π_0 -almost-all θ . But this adds more measure-0 caveats to the results, and it is hard to find a natural example where it would help.

Assumption 2. For every $\theta \in \Theta$, the Kullback-Leibler divergence rate from P ,

$$h(\theta) = \lim_{t \rightarrow \infty} \frac{1}{t} \mathbf{E} \left[\log \frac{p(X_1^t)}{f_\theta(X_1^t)} \right]$$

exists (possibly being infinite) and is \mathcal{T} -measurable.

As mentioned, any two distinct ergodic measures are mutually singular, so there is a consistent test which can separate them. ([53] constructs an explicit but not necessarily optimal test.) One interpretation of the divergence rate [32] is that it measures the maximum exponential rate at which the power of such tests approaches 1, with $d = 0$ and $d = \infty$ indicating sub- and supra- exponential convergence, respectively.

Assumption 3. For each $\theta \in \Theta$, the generalized or relative asymptotic equipartition property holds, and so

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log R_t(\theta) = -h(\theta) \tag{1}$$

with P -probability 1.

Refs. [1, 32] give sufficient, but not necessary, conditions sufficient for Assumption 3 to hold for a given θ . The ordinary, non-relative asymptotic equipartition property, also known as the Shannon-McMillan-Breiman theorem, is that $\lim_{t \rightarrow \infty} t^{-1} \log p(x_1^t) = -h_P$ a.s., where h_P is the entropy rate of the data-generating process. (See [32].) If this holds and h_P is finite, one could rephrase Assumption 3 as $\lim_{t \rightarrow \infty} t^{-1} \log f_\theta(X_1^t) = -h_P - h(\theta)$ a.s., and state results in terms of the likelihood rather than the likelihood ratio. (Cf. [24, ch. 5].) However, there are otherwise-well-behaved processes for which $h_P = -\infty$, at least in the usual choice of reference measure, so I will restrict myself to likelihood ratios.

The meaning of Assumption 3 is that, relative to the true distribution, the likelihood of each θ goes to zero exponentially, the rate being the Kullback-Leibler divergence rate. Roughly speaking, an integral of exponentially-shrinking quantities will tend to be dominated by the integrand with the slowest rate of decay. This suggests that the posterior probability of a set $A \subseteq \Theta$ depends on the smallest divergence rate which can be attained at a point of prior support within A . Thus, adapting notation from large deviations theory, define

$$\begin{aligned} h(A) &\equiv \operatorname{ess\,inf}_{\theta \in A} h(\theta) \\ J(\theta) &\equiv h(\theta) - h(\Theta) \\ J(A) &\equiv \operatorname{ess\,inf}_{\theta \in A} J(\theta) \end{aligned}$$

where here and throughout $\operatorname{ess\,inf}$ is the essential infimum with respect to Π_0 , i.e., the greatest lower bound which holds with Π_0 -probability 1.

Our further assumptions are those needed for the “roughly speaking” and “should” statements of the previous paragraph to be true, so that, for reasonable sets $A \in \mathcal{T}$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0(R_t A) = -h(A)$$

Let $I \equiv \{\theta : h(\theta) = \infty\}$.

Assumption 4. $\Pi_0(I) < 1$

If this assumption fails, then every hypothesis in the support of the prior doesn't just diverge from the true data-generating distribution, it diverges so rapidly that the error rate of a test against the latter distribution goes to zero faster than any exponential. (One way this can happen is if every hypothesis has a finite-dimensional distribution assigning probability zero to some event of positive P -probability.) The methods of this paper seem to be of no use in the face of such extreme mis-specification.

Our first substantial assumption is that the prior distribution does not give too much weight to parts of Θ where the log likelihood converges badly.

Assumption 5. *There exists a sequence of sets $G_t \rightarrow \Theta$ such that*

1. $\Pi_0(G_t) \geq 1 - \alpha \exp\{-t\beta\}$, for some $\alpha > 0$, $\beta > 2h(\Theta)$;
2. *The convergence of Eq. 1 is uniform in θ over $G_t \setminus I$;*
3. $h(G_t) \rightarrow h(\Theta)$.

Comment 1: An analogy with the method of sieves [25] may clarify the meaning of the assumption. If we were constrained to some fixed G , the uniform convergence in the second part of the assumption would make the convergence of the posterior distribution fairly straightforward. Now imagine that the constraint set is gradually relaxed, so that at time t the posterior is confined to G_t , which grows so slowly that convergence is preserved. (Assumption 6 below is, in essence, about the relaxation being sufficiently slow.) The theorems work by showing that the behavior of the posterior distribution on the full space Θ is dominated by its behavior on this "sieve".

Comment 2: Recall that by Egorov's theorem [35, Lemma 1.36, p. 18], if a sequence of finite, measurable functions $f_t(\theta)$ converges pointwise to a finite, measurable function $f(\theta)$ for Π_0 -almost-all $\theta \in G$, then for each $\epsilon > 0$, there is a (possibly empty) $B \subset G$ such that $\Pi_0(B) \leq \epsilon$, and the convergence is uniform on $G \setminus B$. Thus the first two parts of the assumption really follow for free from the measurability in θ of likelihoods and divergence rates. (That β needs to be at least $2h(\Theta)$ becomes apparent in the proof of Lemma 5, but that could always be arranged.) The extra content comes in the third part of the assumption, which could fail if the lowest-divergence hypotheses were also the ones where the convergence was slowest, consistently falling into the bad sets B allowed by Egorov's theorem.

For each measurable $A \subseteq \Theta$, for every $\delta > 0$, there exists a random natural number $\tau(A, \delta)$ such that

$$t^{-1} \log \Pi_0(A R_t) \leq \delta + \limsup_t t^{-1} \log \Pi_0(A R_t)$$

for all $t > \tau(A, \delta)$, provided the lim sup is finite. We need this random last-entry time $\tau(A, \delta)$ to state the next assumption.

Assumption 6. *The sets G_t of the previous assumption can be chosen so that, for every δ , the inequality $t \geq \tau(G_t, \delta)$ holds a.s. for all sufficiently large t .*

The fraction of the prior probability mass outside of G_t is exponentially small in t , with the decay rate large enough that (Lemma 5) the posterior probability mass outside G_t also goes to zero. Using the analogy to the sieve again, the meaning of the assumption is that the convergence of the log-likelihood ratio is sufficiently fast, and the relaxation of the sieve is sufficiently slow, that, at least eventually, every set G_t has δ -converged by t , the time when we start using it.

To show convergence of the posterior measure, we need to be able to control the convergence of the log-likelihood on sets smaller than the whole parameter space.

Assumption 7. *The sets G_t of the previous two assumptions can be chosen so that, for any set A with $\Pi_0(A) > 0$, $h(G_t \cap A) \rightarrow h(A)$.*

Assumption 7 could be replaced by the logically-weaker assumption that for each set A , there exist a sequence of sets $G_{t,A}$ satisfying the equivalents of Assumptions 5 and 6 for the prior measure restricted to A . Since the most straightforward way to check such an assumption would be to verify Assumption 7 as stated, the extra generality does not seem worth it.

2.2.1. Verification of assumptions for the example

Since every $\theta \in \Theta$ is a finite-order Markov chain, and P is stationary and ergodic, Assumption 1 is unproblematic, while Assumptions 2 and 3 hold by virtue of [1].

It is easy to check that $\inf_{\theta \in \Theta_k} h(\theta) > 0$ for each k . (The infimum is not in general attained by any $\theta \in \Theta_k$, though it could be if the chains were allowed to have some transition probabilities equal to zero.) The infimum over Θ as a whole, however, is zero. Also, $h(\theta) < \infty$ everywhere (because none of the hypotheses' transition probabilities are zero), so the possible set I of θ with infinite divergence rates is empty, disposing of Assumption 4.

Verifying the remaining assumptions means building a sequence G_t of increasing subsets of Θ on which the convergence of $t^{-1} \log R_t$ is uniform and sufficiently rapid, and ensuring that the prior probability of these sets grows fast enough. This will be done by exploiting some finite-sample deviation bounds for the even process, which in turn rest on its mixing properties and the method of types. Details are referred to Appendix B. The upshot is that the sets G_t consist of chains whose order is less than or equal to $\frac{\log t}{2/3+\epsilon} - 1$, for some $\epsilon > 0$, and where the absolute logarithm of all the transition probabilities is bounded by Ct^γ , where the positive constant C is arbitrary but $0 < \gamma < \frac{2/3+\epsilon/2}{2/3+\epsilon}$. (With a different strictly sofic process P , the constant $2/3$ in the preceding expressions should be replaced by h_P .) The exponential rate $\beta > 0$ for the prior probability of G_t^c can be chosen to be arbitrarily small.

3. Results

I first give the theorems here, without proof. The proofs, in §§3.1–3.5, are accompanied by re-statements of the theorems, for the reader's convenience.

There are six theorems. The first upper-bounds the growth rate of the posterior density at a given point θ in Θ . The second matches the upper bound on the posterior density with a lower bound, together providing the growth-rate for the posterior density. The third is that $\Pi_t(A) \rightarrow 0$ for any set with $J(A) > 0$, showing that the posterior concentrates on the divergence-minimizing part of the hypothesis space. The fourth is a kind of large deviations principle for the posterior measure. The fifth bounds the asymptotic Hellinger and total variation distances between the posterior predictive distribution and the actual conditional distribution of the next observation. Finally, the sixth theorem establishes rates of convergence.

The first result uses only Assumptions 1–3. (It is not very interesting, however, unless 4 is also true.) The latter three, however, all depend on finer control of the integrated likelihood, and so finer control of the prior, as embodied in Assumptions 5–6. More exactly, those additional assumptions concern the interplay between the prior and the data-generating process, restricting the amount of prior probability which can be given to hypotheses whose log-likelihoods converge excessively slowly under P . I build to the first result in the next subsection, then turn to the control of the integrated likelihood and its consequences in the next three sub-sections, and then consider how these results apply to the example.

Theorem 1. *Under Assumptions 1–3, with probability 1, for all θ where $\pi_0(\theta) > 0$,*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) \leq -J(\theta)$$

Theorem 2. *Under Assumptions 1–6, for all $\theta \in \Theta$ where $\pi_0(\theta) > 0$,*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) = -J(\theta)$$

with probability 1.

Theorem 3. *Make Assumptions 1–7. Pick any set $A \in \mathcal{T}$ where $\Pi_0(A) > 0$ and $h(A) > h(\Theta)$. Then $\Pi_t(A) \rightarrow 0$ a.s.*

Theorem 4. *Under the conditions of Theorem 3, if $A \in \mathcal{T}$ is such that*

$$-\limsup t^{-1} \log \Pi_0(A \cap G_t^c) = \beta' \geq 2h(A)$$

then

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \Pi_t(A) = -J(A)$$

In particular, this holds whenever $2h(A) < \beta$ or $A \subset \bigcap_{k=n}^{\infty} G_k$ for some n .

Theorem 5. Under Assumptions 1–7, with probability 1,

$$\begin{aligned} \limsup_{t \rightarrow \infty} \rho_H^2(P^t, F_\Pi^t) &\leq h(\Theta) \\ \limsup_{t \rightarrow \infty} \rho_{TV}^2(P^t, F_\Pi^t) &\leq 4h(\Theta) \end{aligned}$$

where ρ_H and ρ_{TV} are, respectively, the Hellinger and total variation metrics.

Theorem 6. Make assumptions 1–7, and pick a positive sequence ϵ_t where $\epsilon_t \rightarrow 0$, $t\epsilon_t \rightarrow \infty$. If, for each $\delta > 0$,

$$\tau(G_t \cap N_{\epsilon_t}^c, \delta) \leq t$$

eventually almost surely, then

$$\Pi_t(N_{\epsilon_t}) \rightarrow 1$$

with probability 1.

3.1. Upper bound on the posterior density

The primary result of this section is a pointwise upper bound on the growth rate of the posterior density. To establish it, I use some subsidiary lemmas, which also recur in later proofs. Lemma 2 extends the almost-sure convergence of the likelihood (Assumption 3) from holding pointwise in Θ to holding simultaneously for all θ on a (possibly random) set of Π_0 -measure 1. Lemma 3 shows that the prior-weighted likelihood ratio, $\Pi_0(R_t)$ tends to be at least $\exp\{-th(\Theta)\}$. (Both assertions are made more precise in the lemmas themselves.)

I begin with a proposition about exchanging the order of universal quantifiers (with almost-sure caveats).

Lemma 1. Let $Q \subset \Theta \times \Omega$ be jointly measurable, with sections $Q_\theta = \{\omega : (\omega, \theta) \in Q\}$ and $Q_\omega = \{\theta : (\omega, \theta) \in Q\}$. If, for some probability measure P on Ω ,

$$\forall \theta P(Q_\theta) = 1 \tag{2}$$

then for any probability measure Π on Θ

$$P(\{\omega : \Pi(Q_\omega) = 1\}) = 1 \tag{3}$$

In words, if, for all θ , some property holds a.s., then a.s. the property holds *simultaneously* for almost all θ .

Proof. Since Q is measurable, for all ω and θ , the sections are measurable, and the measures of the sections, $P(Q_\theta)$ and $\Pi(Q_\omega)$, are measurable functions of θ

and ω , respectively. Using Fubini's theorem,

$$\begin{aligned} \int_{\Theta} P(Q_{\theta})d\Pi(\theta) &= \int_{\Theta} \int_{\Omega} \mathbf{1}_Q(\omega, \theta)dP(\omega)d\Pi(\theta) \\ &= \int_{\Omega} \int_{\Theta} \mathbf{1}_Q(\omega, \theta)d\Pi(\theta)dP(\omega) \\ &= \int_{\Omega} \Pi(Q_{\omega})dP(\omega) \end{aligned}$$

By hypothesis, however, $P(Q_{\theta}) = 1$ for all θ . Hence it must be the case that $\Pi(Q_{\omega}) = 1$ for P -almost-all ω . (In fact, the set of ω for which this is true must be a measurable set.) \square

Lemma 2. *Under Assumptions 1–3, there exists a set $C \subseteq \Xi^{\infty}$, with $P(C) = 1$, where, for every $y \in C$, there exists a $Q_y \in \mathcal{T}$ such that, for every $\theta \in Q_y$, Eq. 1 holds. Moreover, $\Pi_0(Q_y) = 1$.*

Proof. Let the set Q consist of the θ, ω pairs where Eq. 1 holds, i.e., for which

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log R_t(\theta, \omega) = -h(\theta) ,$$

being explicit about the dependence of the likelihood ratio on ω . Assumption 3 states that $\forall \theta P(Q_{\theta}) = 1$, so applying Lemma 1 just needs the verification that Q is jointly measurable. But, by Assumptions 1 and 2, $h(\cdot)$ is \mathcal{T} -measurable, and $R_t(\theta)$ is $\sigma(X_1^t) \times \mathcal{T}$ -measurable for each t , so the set Q where the convergence holds are $\sigma(X_1^{\infty}) \times \mathcal{T}$ -measurable. Everything then follows from the preceding lemma. \square

Remark. Lemma 2 generalizes Lemma 3 in [4]. Lemma 1 is a specialization of the quantifier-reversal lemma used in [45] to prove PAC-Bayesian theorems for learning classifiers. Lemma 1 could be used to extend any of the results below which hold a.s. for each θ to ones which a.s. hold *simultaneously* almost everywhere in Θ . This may seem too good to be true, like an alchemist's recipe for turning the lead of pointwise limits into the gold of uniform convergence. Fortunately or not, however, the lemma tells us nothing about the *rate* of convergence, and is compatible with its varying across Θ from instantaneous to arbitrarily slow, so uniform laws need stronger assumptions.

Lemma 3. *Under Assumptions 1–3, for every $\epsilon > 0$, it is almost sure that the ratio between the integrated likelihood and the true probability density falls below $\exp\{-t(h(\Theta) + \epsilon)\}$ only finitely often:*

$$P\{x_1^{\infty} : \Pi_0(R_t) \leq \exp\{-t(h(\Theta) + \epsilon)\}, \text{ i.o.}\} = 0 \tag{4}$$

and as a corollary, with probability 1,

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0(R_t) \geq -h(\Theta) \tag{5}$$

Proof. It's enough to show that Eq. 4 holds for all x_1^∞ in the set C from the previous lemma, since that set has probability 1.

Let $N_{\epsilon/2}$ be the set of all θ in the support of Π_0 such that $h(\theta) \leq h(\Theta) + \epsilon/2$. Since $x_1^\infty \in C$, the previous lemma tells us there exists a set $Q_{x_1^\infty}$ of θ for which Eq. 1 holds under the sequence x_1^∞ .

$$\begin{aligned} \exp \{t(\epsilon + h(\Theta))\} \Pi_0 (R_t) &= \int_{\Theta} R_t(\theta) \exp \{t(\epsilon + h(\Theta))\} d\Pi_0(\theta) \\ &\geq \int_{N_{\epsilon/2} \cap Q_{x_1^\infty}} R_t(\theta) \exp \{t(\epsilon + h(\Theta))\} d\Pi_0(\theta) \\ &= \int_{N_{\epsilon/2} \cap Q_{x_1^\infty}} \exp \left\{ t \left[\epsilon + h(\Theta) + \frac{\log R_t(\theta)}{t} \right] \right\} d\Pi_0(\theta) \end{aligned}$$

By Assumption 3,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log R_t(\theta) = -h(\theta)$$

and for all $\theta \in N_{\epsilon/2}$, $h(\theta) \leq h(\Theta) + \epsilon/2$, so

$$\liminf_{t \rightarrow \infty} \exp \left\{ t \left[\epsilon + h(\Theta) + \frac{1}{t} \log R_t(\theta) \right] \right\} = \infty$$

a.s., for all $\theta \in N_{\epsilon/2} \cap Q_{x_1^\infty}$. We must have $\Pi_0(N_{\epsilon/2}) > 0$, otherwise $h(\Theta)$ would not be the essential infimum, and we know from the previous lemma that $\Pi_0(Q_{x_1^\infty}) = 1$, so $\Pi_0(N_{\epsilon/2} \cap Q_{x_1^\infty}) > 0$. Thus, Fatou's lemma gives

$$\lim_{t \rightarrow \infty} \int_{N_{\epsilon/2} \cap Q_{x_1^\infty}} \exp \left\{ t \left[\epsilon + h(\Theta) + \frac{1}{t} \log R_t(\theta) \right] \right\} d\Pi_0(\theta) = \infty$$

so

$$\lim_{t \rightarrow \infty} \exp \{t(\epsilon + h(\Theta))\} \Pi_0 (R_t) = \infty$$

and hence

$$\Pi_0 (R_t) > \exp \{-t(\epsilon + h(\Theta))\} \tag{6}$$

for all but finitely many t . Since this holds for all $x_1^\infty \in C$, and $P(C) = 1$, Equation 6 holds a.s., as was to be shown. The corollary statement follows immediately. \square

Theorem 1. Under Assumptions 1-3, with probability 1, for all θ where $\pi_0(\theta) > 0$,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) \leq -J(\theta) \tag{7}$$

Proof. As remarked,

$$\pi_t(\theta) = \pi_0(\theta) \frac{R_t(\theta)}{\Pi_0 (R_t)}$$

so

$$\frac{1}{t} \log \pi_t(\theta) = \frac{1}{t} \log \pi_0(\theta) + \frac{1}{t} \log R_t(\theta) - \frac{1}{t} \log \Pi_0(R_t)$$

By Assumption 3, for each $\epsilon > 0$, it's almost sure that

$$\frac{1}{t} \log R_t(\theta) \leq -h(\theta) + \epsilon/2$$

for all sufficiently large t , while by Lemma 3, it's almost sure that

$$\frac{1}{t} \log \Pi_0(R_t) \geq -h(\Theta) - \epsilon/2$$

for all sufficiently large t . Hence, with probability 1,

$$\frac{1}{t} \log \pi_t(\theta) \leq h(\Theta) - h(\theta) + \epsilon + \frac{1}{t} \log \pi_0(\theta)$$

for all sufficiently large t . Hence

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) \leq h(\Theta) - h(\theta) = -J(\theta)$$

□

Lemma 3 gives a lower bound on the integrated likelihood ratio, showing that in the long run it has to be at least as big as $\exp\{-th(\Theta)\}$. (More precisely, it is significantly smaller than that on vanishingly few occasions.) It does not, however, rule out being larger. Ideally, we would be able to match this lower bound with an upper bound of the same form, since $h(\Theta)$ is the best attainable divergence rate, and, by Lemma 2, log likelihood ratios per unit time are converging to divergence rates for Π_0 -almost-all θ , so values of θ for which $h(\theta)$ are close to $h(\Theta)$ should come to dominate the integral in $\Pi_0(R_t)$. It would then be fairly straightforward to show convergence of the posterior distribution.

Unfortunately, additional assumptions are required for such an upper bound, because (as earlier remarked) Lemma 2 does not give *uniform* convergence, merely universal convergence; with a large enough space of hypotheses, the slowest pointwise convergence rates can be pushed arbitrarily low. For instance, let \overline{x}_1^t be the distribution on Ξ^∞ which assigns probability 1 to endless repetitions of x_1^t ; clearly, under this distribution, seeing $X_1^t = x_1^t$ is almost certain. If such measures fall within the support of Π_0 , they will dominate the likelihood, even though $h(\overline{x}_1^t) = \infty$ under all but very special circumstances (e.g., $P = \overline{x}_1^t$). Generically, then, the likelihood and the posterior weight of \overline{x}_1^t will rapidly plummet at times $T > t$. To ensure convergence of the posterior, overly-flexible measures like the family of \overline{x}_1^t 's must be either excluded from the support of Π_0 (possibly because they are excluded from Θ), or be assigned so little prior weight that they do not end up dominating the integrated likelihood, or the posterior must concentrate on them.

3.2. Convergence of posterior density via control of the integrated likelihood

The next two lemmas tell us that sets in Θ of exponentially-small prior measure make vanishingly small contributions to the integrated likelihood, and so to the posterior. They do not require assumptions beyond those used so far, but their application will.

Lemma 4. *Make Assumptions 1-3, and chose a sequence of sets $B_t \subset \Theta$ such that, for all sufficiently large t , $\Pi_0(B_t) \leq \alpha \exp \{-t\beta\}$ for some $\alpha, \beta > 0$. Then, almost surely,*

$$\Pi_0(R_t B_t) \leq \exp \{-t\beta/2\} \tag{8}$$

for all but finitely many t .

Proof. By Markov's inequality. First, use Fubini's theorem and the chain rule for Radon-Nikodym derivatives to calculate the expectation value of the ratio.

$$\begin{aligned} \mathbf{E}[\Pi_0(R_t B_t)] &= \int_{\mathcal{X}^t} dP(x_1^t) \int_{B_t} d\Pi_0(\theta) R_t(\theta) \\ &= \int_{B_t} d\Pi_0(\theta) \int_{\mathcal{X}^n} dP(x_1^t) \frac{dF_\theta}{dP}(x_1^t) \\ &= \int_{B_t} d\Pi_0(\theta) \int_{\mathcal{X}^t} dF_\theta(x_1^t) \\ &= \int_{B_t} d\Pi_0(\theta) \\ &= \Pi_0(B_t) \end{aligned}$$

Now apply Markov's inequality:

$$\begin{aligned} P\{x_1^t : \Pi_0(R_t B_t) > \exp \{-t\beta/2\}\} &\leq \exp \{t\beta/2\} \mathbf{E}[\Pi_0(R_t B_t)] \\ &= \exp \{t\beta/2\} \Pi_0(B_t) \\ &\leq \alpha \exp \{-t\beta/2\} \end{aligned}$$

for all sufficiently large t . Since these probabilities are summable, the Borel-Cantelli lemma implies that, with probability 1, Eq. 8 holds for all but finitely many t . □

The next lemma asserts a sequence of exponentially-small sets makes a (logarithmically) negligible contribution to the posterior distribution, provided the exponent is large enough compared to $h(\Theta)$.

Lemma 5. *Let B_t be as in the previous lemma. If $\beta > 2h(\Theta)$, then*

$$\frac{\Pi_0(R_t B_t^c)}{\Pi_0(R_t)} \rightarrow 1 \tag{9}$$

Proof. Begin with the likelihood integrated over B_t rather than its complement, and apply Lemmas 3 and 4: for any $\epsilon > 0$

$$\frac{\Pi_0(B_t R_t)}{\Pi_0(R_n)} \leq \frac{\exp\{-t\beta/2\}}{\exp\{-t[h(\Theta) + \epsilon]\}} \tag{10}$$

$$= \exp\{t[\epsilon + h(\Theta) - \beta/2]\} \tag{11}$$

provided t is sufficiently large. If $\beta > 2h(\Theta)$, this bound can be made to go to zero as $t \rightarrow \infty$ by taking ϵ to be sufficiently small. Since

$$\Pi_0(R_t) = \Pi_0(B_t^c R_t) + \Pi_0(B_t R_t)$$

it follows that

$$\frac{\Pi_0(B_t^c R_t)}{\Pi_0(R_t)} \rightarrow 1$$

□

Lemma 6. *Make Assumptions 1-3, and take any set G on which the convergence in Eq. 1 is uniform and where $\Pi_0(G) > 0$. Then, P-a.s.,*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0(GR_t) \leq -h(G) \tag{12}$$

Proof. Pick any $\epsilon > 0$. By the hypothesis of uniform convergence, there almost surely exists a $T(\epsilon)$ such that, for all $t \geq T(\epsilon)$ and for all $\theta \in G$, $t^{-1} \log R_t(\theta) \leq -h(\theta) + \epsilon$. Hence

$$t^{-1} \log \Pi_0(GR_t) = t^{-1} \log \Pi_0(G \exp\{\log R_t\}) \tag{13}$$

$$\leq t^{-1} \log \Pi_0(G \exp\{t[-h + \epsilon]\}) \tag{14}$$

$$= \epsilon + t^{-1} \log \Pi_0(G \exp\{-th\}) \tag{15}$$

Let $\Pi_{0|G}$ denote the probability measure formed by conditioning Π_0 to be in the set G . Then

$$\Pi_0(Gz) = \Pi_0(G) \int_G d\Pi_{0|G}(\theta) z(\theta)$$

for any integrable function z . Apply this to the last term from Eq. 15.

$$\log \Pi_0(G \exp\{-th\}) = \log \Pi_0(G) + \log \int_G d\Pi_{0|G}(\theta) \exp\{-th(\theta)\}$$

The second term on the right-hand side is the cumulant generating function of $-h(\theta)$ with respect to $\Pi_{0|G}$, which turns out (cf. 6) to have exactly the right behavior as $t \rightarrow \infty$.

$$\begin{aligned}
 \frac{1}{t} \log \int_G d\Pi_{0|G}(\theta) \exp \{-th(\theta)\} &= \frac{1}{t} \log \int_G d\Pi_{0|G}(\theta) |\exp \{-h(\theta)\}|^t \\
 &= \frac{1}{t} \log \left(\left(\int_G d\Pi_{0|G}(\theta) |\exp \{-h(\theta)\}|^t \right)^{1/t} \right)^t \\
 &= \frac{1}{t} \left(t \log \|\exp \{-h(\theta)\}\|_{t, \Pi_{0|G}} \right) \\
 &= \log \|\exp \{-h(\theta)\}\|_{t, \Pi_{0|G}} \tag{16}
 \end{aligned}$$

Since $h(\theta) \geq 0$, $\exp \{-h(\theta)\} \leq 1$, and the L_p norm of the latter will grow towards its L_∞ norm as p grows. Hence, for sufficiently large t ,

$$\begin{aligned}
 \log \|\exp \{-h(\theta)\}\|_{t, \Pi_{0|G}} &\leq \log \|\exp \{-h(\theta)\}\|_{\infty, \Pi_{0|G}} + \epsilon \\
 &= -\operatorname{ess\,inf}_{\theta \in G} h(\theta) + \epsilon \\
 &= -h(G) + \epsilon \tag{17}
 \end{aligned}$$

where the next-to-last step uses the monotonicity of \log and \exp .

Putting everything together, we have that, for any $\epsilon > 0$ and all sufficiently large t ,

$$t^{-1} \log \Pi_0(GR_t) \leq -h(G) + 2\epsilon + \frac{\log \Pi_0(G)}{t} \tag{18}$$

Hence the limit superior of the left-hand side is at most $-h(G)$. □

Lemma 7. *Under Assumption 1-6,*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0(R_t) \leq -h(\Theta) \tag{19}$$

Proof. By Lemma 5,

$$\lim_{t \rightarrow \infty} \frac{\Pi_0(R_t)}{\Pi_0(G_t R_t)} = 1$$

implying that

$$\lim_{t \rightarrow \infty} \log \Pi_0(R_t) - \log \Pi_0(G_t R_t) = 0$$

so for every $\epsilon > 0$, for t large enough

$$\log \Pi_0(R_t) \leq \epsilon/3 + \log \Pi_0(G_t R_t)$$

Consequently, again for large enough t ,

$$\frac{1}{t} \log \Pi_0(R_t) \leq \epsilon/3t + \frac{1}{t} \log \Pi_0(G_t R_t)$$

Now, for each set G , for every $\epsilon > 0$, if $t \geq \tau(G, \epsilon/3)$ then

$$\frac{1}{t} \log \Pi_0(GR_t) \leq -h(G) + \epsilon/3$$

by Lemma 6. By Assumption 6, $t \geq \tau(G_t, \epsilon/3)$ for all sufficiently large t . Hence

$$\frac{1}{t} \log \Pi_0 (R_t) \leq -h(G_t) + \epsilon/3t + \epsilon/3$$

for all $\epsilon > 0$ and all t sufficiently large. Since, by Assumption 5, $h(G_t) \rightarrow h(\Theta)$, for every $\epsilon > 0$, $h(G_t)$ is within $\epsilon/3$ of $h(\Theta)$ for large enough t , so

$$\frac{1}{t} \log \Pi_0 (R_t) \leq -h(\Theta) + \epsilon/3t + \epsilon/3 + \epsilon/3$$

Thus, for every $\epsilon > 0$, then we have that

$$\frac{1}{t} \log \Pi_0 (R_t) \leq -h(\Theta) + \epsilon$$

for large enough t , or, in short,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0 (R_t) \leq -h(\Theta)$$

□

Lemma 8. *Under Assumptions 1–6, if $\Pi_0(I) = 0$, then*

$$\frac{1}{t} \log \Pi_0 (R_t) \rightarrow -h(\Theta) \tag{20}$$

almost surely.

Proof. Combining Lemmas 3 and 7,

$$-h(\Theta) \leq \liminf_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0 (R_t) \leq \limsup_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0 (R_t) \leq -h(\Theta)$$

□

The standard version of Egorov’s theorem concerns sequences of finite measurable functions converging pointwise to a finite measurable limiting function. However, the proof is easily adapted to an infinite limiting function.

Lemma 9. *Let $f_t(\theta)$ be a sequence of finite, measurable functions, converging to ∞ almost everywhere (Π_0) on I . Then for each $\epsilon > 0$, there exists a possibly-empty $B \subset I$ such that $\Pi_0(B) < \epsilon$, and the convergence is uniform on $I \setminus B$.*

Proof. Parallel to the usual proof of Egorov’s theorem. Begin by removing the measure-zero set of points on which pointwise convergence fails; for simplicity, keep the name I for the remaining set. For each natural number t and k , let $B_{t,k} \equiv \{\theta \in I : f_t(\theta) < k\}$ — the points where the function fails to be at least k by step t . Since the limit of f_t is ∞ everywhere on I , each θ has a last t such that $f_t(\theta) < k$, no matter how big k is. Hence $\bigcap_{t=1}^{\infty} B_{t,k} = \emptyset$. By continuity of measure, for any $\delta > 0$, there exists an n such that $\Pi_0(B_{t,k}) < \delta$ if $t \geq n$. Fix ϵ as in the statement of the lemma, and set $\delta = \epsilon 2^{-k}$. Finally, set $B = \bigcup_{k=1}^{\infty} B_{n,k}$. By the union bound, $\Pi_0(B) \leq \epsilon$, and by construction, the rate of convergence to ∞ is uniform on $I \setminus B$. □

Lemma 10. *The conclusion of Lemma 8 is unchanged if $\Pi_0(I) > 0$.*

Proof. The integrated likelihood ratio can be divided into two parts, one from integrating over I and one from integrating over its complement. Previous lemmas have established that the latter is upper bounded, in the long run, by a quantity which is $O(\exp\{-h(\Theta)t\})$. We can use Lemma 9 to divide I into a sequence of sub-sets, on which the convergence is uniform, and hence on which the integrated likelihood shrinks faster than any exponential function, and remainder sets, of prior measure no more than $\alpha \exp\{-n\beta\}$, on which the convergence is less than uniform (i.e., slow). If we ensure that $\beta > 2h(\Theta)$, however, by Lemma 5 the remainder sets' contributions to the integrated likelihood is negligible in comparison to that of $\Theta \setminus I$. Said another way, if there are alternatives which a consistent test would rule out at a merely exponential rate, those which would be rejected at a supra-exponential rate end up making vanishingly small contributions to the integrated likelihood. \square

Theorem 2. *Under Assumptions 1-6, for all $\theta \in \Theta$ where $\pi_0(\theta) > 0$,*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) = -J(\theta) \quad (21)$$

with probability 1.

Proof. Theorem 1 says that, for all θ ,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) \leq -J(\theta)$$

a.s., so there just needs to be a matching lim inf. Pick any $\epsilon > 0$. By Assumption 3, it's almost certain that, for all sufficiently large t ,

$$\frac{1}{t} \log R_t(\theta) \geq -h(\theta) - \epsilon/2$$

while by Lemma 10, it's almost certain that for all sufficiently large t ,

$$\frac{1}{t} \log \Pi_0(R_t) \leq -h(\Theta) + \epsilon/2$$

Combining these as in the proof of Theorem 1, it's almost certain that for all sufficiently large t

$$\frac{1}{t} \log \pi_t(\theta) \geq h(\Theta) - h(\theta) - \epsilon$$

so

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \pi_t(\theta) \geq h(\Theta) - h(\theta) = -J(\theta)$$

\square

3.3. Convergence and large deviations of the posterior measure

Adding Assumption 7 to those before it implies that the posterior measure concentrates on sets $A \subset \Theta$ where $h(A) = h(\Theta)$.

Theorem 3. *Make Assumptions 1–7. Pick any set $A \in \mathcal{T}$ where $\Pi_0(A) > 0$ and $h(A) > h(\Theta)$. Then $\Pi_t(A) \rightarrow 0$ a.s.*

Proof.

$$\begin{aligned} \Pi_t(A) &= \Pi_t(A \cap G_t) + \Pi_t(A \cap G_t^c) \\ &\leq \Pi_t(A \cap G_t) + \Pi_t(G_t^c) \end{aligned}$$

The last term is easy to bound. From Eq. 11 in the proof of Lemma 5,

$$\begin{aligned} \Pi_t(G_t^c) &= \frac{\Pi_0(R_t G_t^c)}{\Pi_0(R_t)} \\ &\leq \exp\{t[\epsilon + h(\Theta) - \beta/2]\} \end{aligned} \tag{22}$$

for any $\epsilon > 0$, for all sufficiently large t , almost surely. Since $\beta > 2h(\Theta)$, the whole expression $\rightarrow 0$ as $t \rightarrow \infty$.

To bound $\Pi_t(A \cap G_t)$, reasoning as in the proof of Lemma 7, but invoking Assumption 7, leads to the conclusion that, for any $\epsilon > 0$, with probability 1,

$$\frac{1}{t} \log \Pi_0(R_t(A \cap G_t)) \leq -h(A) + \epsilon$$

for all sufficiently large t . Recall that by Lemma 3, for all $\epsilon > 0$ it's almost sure that

$$\frac{1}{t} \log \Pi_0(R_t) \geq -h(\Theta) - \epsilon$$

for all sufficiently large t . Hence for every $\epsilon > 0$, it's almost certain that for all sufficiently large t ,

$$\Pi_t(A \cap G_t) \leq \exp\{t[h(\Theta) - h(A) + 2\epsilon]\} \tag{23}$$

Since $h(A) > h(\Theta)$, by picking ϵ small enough the right hand side goes to zero. □

The proof of the theorem provides an exponential *upper* bound on the posterior measure of sets where $h(A) > h(\Theta)$. In fact, even without the final assumption needed for the theorem, there is an exponential *lower* bound on that posterior measure.

Lemma 11. *Make Assumption 1–6, and pick a set $A \in \mathcal{T}$ with $\Pi_0(A) > 0$. Then*

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \Pi_t(A) \geq -J(A) \tag{24}$$

Proof. Reasoning as in the proof of Lemma 3, it is easy to see that

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0 (R_t A) \geq -h(A)$$

and by Lemma 7,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \Pi_0 (R_t) \leq -h(\Theta)$$

hence

$$\begin{aligned} \liminf_{t \rightarrow \infty} \frac{1}{t} \log \Pi_t(A) &= \liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{\Pi_0 (R_t A)}{\Pi_0 (R_t)} \\ &\geq -h(A) + h(\Theta) \end{aligned}$$

□

Theorem 4. Under the conditions of Theorem 3, if $A \in \mathcal{T}$ is such that

$$-\limsup t^{-1} \log \Pi_0(A \cap G_t^c) = \beta' \geq 2h(A) \tag{25}$$

then

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \Pi_t(A) = h(\Theta) - h(A) \tag{26}$$

In particular, this holds whenever $2h(A) < \beta$ or $A \subset \bigcap_{k=n}^{\infty} G_k$ for some n .

Proof. Trivially,

$$\frac{1}{t} \log \Pi_t(A) = \frac{1}{t} \log \Pi_t(A \cap G_t) + \frac{1}{t} \log \Pi_t(A \cap G_t^c)$$

From Eq. 23 from the proof of Theorem 3, we know that, for any $\epsilon > 0$,

$$\Pi_t(A \cap G_t) \leq \exp \{t[h(\Theta) - h(A) + \epsilon]\}$$

a.s. for sufficiently large t . On the other hand, under the hypothesis of the theorem, the proof of Eq. 22 can be imitated for $\Pi_t(A \cap G_t^c)$, with the conclusion that, for all $\epsilon > 0$,

$$\Pi_t(A \cap G_t^c) \leq \exp \{t[h(\Theta) - \beta'/2 + \epsilon]\}$$

again a.s. for sufficiently large t . Since $\beta'/2 > h(A)$, Eq. 26 follows.

Finally, to see that this holds for any A where $h(A) < \beta/2$, observe that we can always upper bound $\Pi_t(A \cap G_t^c)$ by $\Pi_t(G_t^c)$, and the latter goes to zero with rate at least $-\beta/2$. □

Remarks. Because $h(A)$ is the essential infimum of $h(\theta)$ on the set A , as the set shrinks $h(A)$ grows. Sets where $h(A)$ is much larger than $h(\Theta)$ tend accordingly to be small. The difficulty is that the sets G_t^c are also small, and conceivably overlap so heavily with A that the integral of the likelihood over A is dominated

by the part coming from $A \cap G_t^c$. Eventually this will shrink towards zero exponentially, but perhaps only at the comparatively slow rate $h(\Theta) - \beta/2$, rather than the faster rate $h(\Theta) - h(A)$ attained on the well-behaved part $A \cap G_t$.

Theorem 4 is close to, but not quite, a large deviations principle on Θ . We have shown that the posterior probability of any arbitrary set A where $J(A) > 0$ goes to zero with an exponential rate at least equal to

$$\beta/2 \wedge \operatorname{ess\,inf}_{\theta \in A} J(\theta) = \operatorname{ess\,inf}_{\theta \in A} \beta/2 \wedge J(\theta) \quad (27)$$

But in a true LDP, the rate would have to be an infimum, not just an essential infimum, of a point-wise rate function. This deficiency could be removed by means of additional assumptions on Π_0 and $h(\theta)$.

Ref. [22] obtains proper large and even moderate deviations principles, but for the location of Π_t in the space $\mathcal{M}_1(\Theta)$ of all distributions on Θ , rather than on Θ itself. Essentially, they use the assumption of IID sampling, which makes the posterior a function of the empirical distribution, to leverage the LDP for the latter into an LDP for the former. This strategy may be more widely applicable but goes beyond the scope of this paper. Papangelou [49], assuming that Θ consists of discrete-valued Markov chains of arbitrary order and P is in the support of the prior, and using methods similar to those in Appendix B, derives a result which is closely related to Theorem 4. In fact, fixing the sets G_t as in Appendix B, Theorem 4 implies the theorem of [49].

3.4. Generalization performance

Lemma 10 shows that, in hindsight, the Bayesian learner does a good job of matching the data: the log integrated likelihood ratio per time-step approaches $-h(\Theta)$, the limit of values attainable by individual hypotheses within the support of the prior. This leaves open, however, the question of the prospective or generalization performance.

What we want is for the posterior predictive distribution F_{Π}^t to approach the true conditional distribution of future events P^t , but we cannot in general hope for the convergence to be complete, since our models are mis-specified. The next theorem uses $h(\Theta)$ to put an upper bound on how far the posterior predictive distribution can remain from the true predictive distribution.

Theorem 5. *Under Assumptions 1–7, with probability 1,*

$$\limsup_{t \rightarrow \infty} \rho_H^2(P^t, F_{\Pi}^t) \leq h(\Theta) \quad (28)$$

$$\limsup_{t \rightarrow \infty} \rho_{TV}^2(P^t, F_{\Pi}^t) \leq 4h(\Theta) \quad (29)$$

where ρ_H and ρ_{TV} are, respectively, the Hellinger and total variation metrics.

Proof. Recall the well-known inequalities relating Hellinger distance to Kullback-Leibler divergence on the one side and to total variation distance on the other

[30]: for any two distributions P and Q ,

$$\rho_H^2(P, Q) \leq D(P\|Q) \tag{30}$$

$$\rho_{TV}(P, Q) \leq 2\rho_H(P, Q) \tag{31}$$

It's enough to prove Eq. 28, and Eq. 29 then follows from Eq. 31.

Abbreviate $\rho_H(P^t, F_\theta^t)$ by $\rho_H(t, \theta)$. Pick any $\epsilon > h(\Theta)$, and say that $A_\epsilon = \{\theta : \rho_H^2(t, \theta) > \epsilon\}$. By convexity and Jensen's inequality,

$$\begin{aligned} \rho_H^2(P^t, F_\Pi^t) &\leq \int_{\Theta} \rho_H^2(t, \theta) d\Pi_n(\theta) \\ &= \int_{A_\epsilon} \rho_H^2(t, \theta) d\Pi_n(\theta) + \int_{A_\epsilon^c} \rho_H^2(t, \theta) d\Pi_n(\theta) \\ &= \epsilon \Pi_t(A_\epsilon^c) + \sqrt{2} \Pi_t(A_\epsilon) \end{aligned}$$

By Eq. 30, $d(\theta) > \rho_H^2(t, \theta)$. Thus $h(A_\epsilon) \geq \epsilon$, and $\epsilon > h(\Theta)$ so, by Theorem 3, $\Pi_t(A_\epsilon) \rightarrow 0$ a.s. Hence

$$\rho_H^2(P^t, F_\Pi^t) \leq \epsilon$$

eventually almost surely. Since this holds for any $\epsilon > h(\Theta)$, Eq. 28 follows. \square

Remark. It seems like it should be possible to prove a similar result for the divergence rate of the predictive distribution, namely that

$$\limsup_{t \rightarrow \infty} h(\Pi_t) \leq h(\Theta)$$

but it would take a different approach, because $h(\cdot)$ has no upper bound, and the posterior weight of the high-divergence regions might decay too slowly to compensate for this.

3.5. Rate of convergence

Recall that N_ϵ was defined as the set of all θ such that $h(\theta) \leq h(\Theta) + \epsilon$. (This is measurable by Assumption 2.) The set N_ϵ^c thus consists of all hypotheses whose divergence rate is more than ϵ above the essential infimum $h(\Theta)$. For any $\epsilon > 0$, $\Pi_t(N_\epsilon^c) \rightarrow 0$ a.s., by Theorem 3, and for sufficiently small ϵ , $\lim_{t \rightarrow \infty} t^{-1} \log \Pi_t(N_\epsilon^c) = -\epsilon$ a.s., by Theorem 4. For such sets, in other words, for any $\delta > 0$, it's almost certain that for all sufficiently large t ,

$$\Pi_t(N_\epsilon^c) \leq \exp\{-t(\epsilon - \delta)\} \tag{32}$$

Now consider a non-increasing positive sequence $\epsilon_t \rightarrow 0$. Presumably if ϵ_t decays slowly enough, $\Pi_t(N_{\epsilon_t}^c)$ will still go to zero, even though the sets $N_{\epsilon_t}^c$ are non-decreasing. Examination of Eq. 32 suggests, naively, that this will work if $t\epsilon_t \rightarrow \infty$, i.e., if the decay of ϵ_t is strictly sublinear. This is correct under an additional condition, similar to Assumption 6.

Theorem 6. *Make Assumptions 1–7, and pick a positive sequence ϵ_t where $\epsilon_t \rightarrow 0$, $t\epsilon_t \rightarrow \infty$. If, for each $\delta > 0$,*

$$\tau(G_t \cap N_{\epsilon_t}^c, \delta) \leq t \tag{33}$$

eventually almost surely, then

$$\Pi_t(N_{\epsilon_t}) \rightarrow 1 \tag{34}$$

with probability 1.

Proof. By showing that $\Pi_t(N_{\epsilon_t}^c) \rightarrow 0$ a.s. Begin by splitting the sets into the parts inside the G_t , say U_t , and the parts outside:

$$\begin{aligned} \Pi_t(N_{\epsilon_t}^c) &= \Pi_t(N_{\epsilon_t}^c \cap G_t) + \Pi_t(N_{\epsilon_t}^c \cap G_t^c) \\ &\leq \Pi_t(U_t) + \Pi_t(G_t^c) \end{aligned}$$

From Lemma 4, the second term $\rightarrow 0$ with probability 1, so for any $\eta_1 > 0$, it is $\leq \eta_1$ eventually a.s.

Turning to the other term, Theorem 4 applies to U_k for any fixed k , so

$$\lim_{t \rightarrow \infty} t^{-1} \log \Pi_t(U_k) = h(\Theta) - h(U_k)$$

(a.s.), implying, with Lemma 10, that

$$\lim_{t \rightarrow \infty} t^{-1} \log \Pi_0(U_k R_t) = -h(U_k)$$

(a.s.). By Eq. 33, for any $\eta_2 > 0$,

$$t^{-1} \log \Pi_0(U_t R_t) \leq -h(U_t) + \eta_2$$

eventually almost surely. By Lemma 10 and Bayes’s rule, then,

$$t^{-1} \log \Pi_t(U_t) \leq h(\Theta) - h(U_t) + \eta_3$$

eventually a.s., for any $\eta_3 > 0$. Putting things back together, eventually a.s.,

$$\begin{aligned} \Pi_t(N_{\epsilon_t}^c) &\leq \exp \{t(h(\Theta) - h(U_t) + \eta_2)\} + \eta_1 \\ &\leq \exp \{t(-\epsilon_t + \eta_3)\} + \eta_1 \end{aligned}$$

Since $t\epsilon_t \rightarrow \infty$, the first term goes to zero, and since η_1 can be as small as desired,

$$\Pi_t(N_{\epsilon_t}^c) \rightarrow 0$$

almost surely. □

The theorem lets us attain rates of convergence just slower than t^{-1} (so that $t\epsilon_t \rightarrow \infty$). This matches existing results on rates of posterior convergence for mis-specified models with IID data in [68, Corollary 5.2] (t^{-1} in the Renyi divergence) and in [38] ($t^{-1/2}$ in the Hellinger distance; recall Eq. 30), and for correctly-specified non-IID models in [29] ($t^{-\alpha}$ for suitable $\alpha < 1/2$, again in the Hellinger distance).

3.6. Application of the results to the example

Because $h(\Theta) = 0$, while $h(\theta) > 0$ everywhere, the behavior of the posterior is somewhat peculiar. Every compact set $K \subset \Theta$ has $J(K) > 0$, so by Theorem 3, $\Pi_t(K) \rightarrow 0$. On the other hand, $\Pi_t(G_t) \rightarrow 1$ — the sequence of good sets contains models of increasingly high order, with increasingly weak constraints on the transition probabilities, and this lets its posterior weight grow, even though every individual compact set within it ultimately loses all weight.

In fact, each G_t is a convex set, and $h(\cdot)$ is a convex function, so there is a unique minimizer of the divergence rate within each good set. Conditional on being within G_t , the posterior probability becomes increasingly concentrated on neighborhoods of this minimizer, but the minimizer itself keeps moving, since it can always be improved upon by increasing the order of the chain and reducing some transition probabilities. (Recall that P gives probability 0 to sequences 010, 01110, etc., where the block of 1's is of odd length, but Θ contains only chains with strictly positive transition probabilities.)

Outside of the good sets, the likelihood is peaked around hypotheses which provide stationary and smooth approximations to the \bar{x}_1^t distribution that endlessly repeats the observed sequence to date. The divergence rates of these hypotheses are however extremely high, so none of them retains its high likelihood for very long. (\bar{x}_1^t is a Markov chain of order t , but it is not in Θ , since it's neither stationary nor does it have strictly positive transition probabilities. It can be made stationary, however, by assigning equal probability to each of its t states; this gives the data likelihood $1/t$ rather than 1, but that still is vastly larger than the $O(-ct)$ log-likelihoods of better models. (Recall that even the log-likelihood of the true distribution is only $O(-\frac{2}{3}t)$.) Allowing each of the t states to have a probability $0 < \iota \ll 1$ of not proceeding to the next state in the periodic sequence is easy and leads to only an $O(\iota t)$ reduction in the likelihood up to time t . In the long run, however, it means that the log-likelihood will be $O(t \log \iota)$.) In any case, the total posterior probability of G_t^c is going to zero exponentially.

Despite — or rather, because of — the fact that no point in Θ is the *ne plus ultra* around which the posterior concentrates, the conditions of Theorem 5 are met, and since $h(\Theta) = 0$, the posterior predictive distribution converges to the true predictive distribution in the Hellinger and total variation metrics. That is, the weird gyrations of the posterior do not prevent us from attaining *predictive* consistency. This is so even though the posterior always gives the wrong answer to such basic questions as “Is $P(X_t^{t+2} = 010) > 0$?” — inferences which in this case can be made correctly through non-Bayesian methods [47, 58].

4. Discussion

The crucial assumptions were 3, 5 and 6. Together, these amount to assuming that the time-averaged log likelihood ratio converges universally; to fashioning a sieve, successively embracing regions of Θ where the convergence is increasingly

ill-behaved; and the hope that the prior weight of the remaining bad sets can be bounded exponentially.

Using asymptotic equipartition in place of the law of large numbers is fairly straightforward. Both results belong to the general family of ergodic theorems, which allow us to take sufficiently long sample paths as representative of entire processes. The unique a.s. limit in Eq. 1 can be replaced with a.s. convergence to a distinct limit in each ergodic component of P . However, the notation gets ugly, so the reader should regard $h(\theta)$ as that random limit, and treat all subsequent results as relative to the ergodic decomposition of P . (Cf. [31, 17].) It may be possible to weaken this assumption yet further, but it is hard to see how Bayesian updating can succeed if the past performance of the likelihood is not a guide to future results.

A bigger departure from the usual approach to posterior convergence may be allowing $h(\Theta) > 0$; this rules out posterior consistency, to begin with. More subtly, it requires $\beta > 2h(\Theta)$. This means that a prior distribution which satisfies the assumptions for one value of P may not satisfy them for another, depending, naturally enough, on just how mis-specified the hypotheses are, and how much weight the prior puts on very bad hypotheses. On the other hand, when $h(\Theta) = 0$, Theorem 5 implies predictive consistency, as in the example.

Assumption 6 is frankly annoying. It ensures that the log likelihood ratio converges fast enough, at least on the good sets, that we can be confident that integrated likelihood of G_t has converged well by the time we want G_t to start dominating the prior. It was shaped, however, to fill a hole in the proof of Lemma 7 rather than by more natural considerations. The result is that verifying the assumption in its present form means proving the sub-linear growth rate of sequences of random last entry times, and these times are not generally convenient to work with. (Cf. Appendix B.) It would be nice to replace it with a bracketing or metric entropy condition, as in [4, 68] or similar forms of capacity control, as used in [46, 63]. Alternately, the uniformly consistent test conditions widely employed in Bayesian nonparametrics [30, 67] have been adapted the mis-specified setting by [38], where the tests become reminiscent of the “model selection tests” used in econometrics [64]. Since the latter can work for dynamical models [51], this approach may also work here. In any event, replacing Assumption 6 with more primitive, comprehensible and easily-verified conditions seems a promising direction for future work.

These results go some way toward providing a frequentist explanation of the success of Bayesian methods in many practical problems. Under these conditions, the posterior is increasingly weighted towards the parts of Θ which are closest (in the Kullback-Leibler sense) to the data-generating process P . For a $\Pi_t(A)$ to persistently be much more or much less than $\approx \exp\{-tJ(A)\}$, $R(\theta)$ must be persistently far from $\exp\{-th(\theta)\}$, not just for isolated $\theta \in A$, but a whole positive-measure subset of them. With a reasonably smooth prior, this requires a run of bad luck amounting almost to a conspiracy. From this point of view, Bayesian inference amounts to introducing bias so as to reduce variance, and then relaxing the bias. Experience with frequentist non-parametric methods shows this can work if the bias is relaxed sufficiently slowly, which is

basically what the assumptions here do. As the example shows, this can succeed as a predictive tactic without supporting substantive inferences about the data-generating process. However, 4–7 involve both the prior and the data-generating process, and so cannot be verified using the prior alone. For empirical applications, it would be nice to have ways of checking them using sample data.

When $h(\Theta) > 0$ and all the models are more or less wrong, there is an additional advantage to averaging the models, as is done in the predictive distribution. (I owe the argument which follows to Scott Page; cf. [48].) With a convex loss function ℓ , such as squared error, Kullback-Leibler divergence, Hellinger distance, etc., the loss of the predictive distribution $\ell(\Pi_t)$ will be no larger than the posterior-mean loss of the individual models $\Pi_t(\ell(\theta))$. For squared error loss, the difference is equal to the variance of the models' predictions [40]. For divergence, some algebra shows that

$$h(\Pi_t) = \Pi_t(h(\theta)) + \Pi_t\left(\mathbf{E}\left[\log\frac{dF_\theta}{dF_\Pi}\right]\right) \quad (35)$$

where the second term on the RHS is again an indication of the diversity of the models; the more different their predictions are, on the kind of data generated by P , the smaller the error of made by the mixture. Having a diversity of wrong answers can be as important as reducing the average error itself. The way to accomplish this is to give more weight to models which make mostly good predictions, but make different mistakes. This suggests that there may actually be predictive benefits to having the posterior concentrate on a set containing multiple hypotheses.

Finally, it is worth remarking on the connection between these results and prediction with “mixtures of experts” [2, 10]. Formally, the role of the negative log-likelihood and of Bayes's rule in this paper was to provide a loss function and a multiplicative scheme for updating the weights. All but one of the main results (Theorem 5, which bounds Hellinger distance by Kullback-Leibler divergence) would carry over to multiplicative weight training using a different loss function, provided the accumulated loss per unit time converged.

Appendix A: Bayesian updating as replicator dynamics

Replicator dynamics are one of the fundamental models of evolutionary biology; they represent the effects of natural selection in large populations, without (in their simplest form) mutation, sex, or other sources of variation. [34] provides a thorough discussion. They also arise as approximations to many other adaptive processes, such as reinforcement learning [8, 9, 54]. In this appendix, I show that Bayesian updating also follows the replicator equation.

We have a set of *replicators* — phenotypes, species, reproductive strategies, etc. — indexed by $\theta \in \Theta$. The population density at type θ is $\pi(\theta)$. We denote by $\phi_t(\theta)$ the *fitness* of θ at time t , i.e., the average number of descendants left by each individual of type θ . The fitness function ϕ_t may in fact be a function of π_t , in which case it is said to be *frequency-dependent*. Many applications assume

the fitness function to be deterministic, rather than random, and further assume that it is not an explicit function of t , but these restrictions are inessential.

The discrete-time *replicator dynamic* [34] is the dynamical system given by the map

$$\pi_t(\theta) = \pi_{t-1}(\theta) \frac{\phi_t(\theta)}{\Pi_t(\phi_t)} \tag{36}$$

where $\Pi_t(\phi_t)$ is the population mean fitness at t , i.e.,

$$\Pi_t(\phi_t) \equiv \int_{\Theta} \phi_t(\theta) d\pi_t(\theta)$$

The effect of these dynamics is to re-weight the population towards replicators with above-average fitness.

It is immediate that Bayesian updating has the same form as Eq. 36, as soon as we identify the distribution of replicators with the posterior distribution, and the fitness with the conditional likelihood. In fact, Bayesian updating is an extra simple case of the replicator equation, since the fitness function is frequency-independent, though stochastic. Updating corresponds to the action of natural selection, without variation, in a fluctuating environment. The results in the main text assume (Assumption 3) that, despite the fluctuations, the long-run fitness is nonetheless a determinate function of θ . The theorems assert that selection can then be relied upon to drive the population to the peaks of the long-run fitness function, at the cost of reducing the diversity of the population, rather as in Fisher’s fundamental theorem of natural selection [23, 34].

Corollary 1. *Define the relative fitness $\tilde{\phi}_t(\theta) \equiv L_t(\theta)/\Pi_t(L_t)$. Under the conditions of Theorem 2, the time average of the log relative fitness converges a.s.*

$$\frac{1}{t} \sum_{n=1}^t \log \tilde{\phi}_n(\theta) \rightarrow -J(\theta) + o(1) \tag{37}$$

Proof. Unrolling Bayes’s rule over multiple observations,

$$\pi_t(\theta) = \pi_0(\theta) \prod_{n=1}^t \tilde{\phi}_n(\theta)$$

Take log of both sides, divide through by t , and invoke Theorem 2. □

Remark. Theorem 2 implies that

$$H_t \equiv |\log \pi_t(\theta) + tJ(\theta)|$$

is a.s. $o(t)$. To strengthen Eq. 37 from convergence of the time average or Cesàro mean to plain convergence requires forcing $H_t - H_{t-1}$ to be $o(1)$, which it generally isn’t.

It is worth noting that Haldane [33] defined the *intensity of selection* on a population as, in the present notation,

$$\log \frac{\pi_t(\hat{\theta})}{\pi_0(\hat{\theta})}$$

where $\hat{\theta}$ is the “optimal” (i.e., most selected-for) value of θ . For us, this intensity of selection is just $R_t(\hat{\theta})/\Pi_0(R_t)$ where $\hat{\theta}$ is the (or a) MLE.

Appendix B: Verification of Assumptions 5–7 for the example

Since the X_1^∞ process is a function of the S_1^∞ process, and the latter is an aperiodic Markov chain, both are ψ -mixing (see [44, 60] for the definition of ψ -mixing and demonstrations that aperiodic Markov chains and their functions are ψ -mixing). Let $\widehat{P}_t^{(k)}$ be the empirical distribution of sequences of length k obtained from x_1^t . For a Markov chain of order k , the likelihood is a function of $\widehat{P}_t^{(k+1)}$ alone; we will use this and the ergodic properties of the data-generating process to construct sets on which the time-averaged log-likelihood converges uniformly. Doing this will involve constraining both the order of the Markov chains and their transition probabilities, and gradually relaxing the constraints.

It will simplify notation if from here on all logarithms are taken to base 2.

Pick $\epsilon > 0$ and let $k(t)$ be an increasing positive-integer-valued function of t , $k(t) \rightarrow \infty$, subject to the limit $k(t) \leq \frac{\log t}{h_P + \epsilon}$, where h_P is the Shannon entropy rate of P , which direct calculation shows is $2/3$. The ψ -mixing property of X_1^∞ implies [60, p. 179] that

$$P(\rho_{TV}(\widehat{P}_t^{(k(t))}, P^{(k(t))}) > \delta) \leq \frac{\log t}{h + \epsilon} 2(n + 1)^{t^{\gamma_1}} 2^{-nC_1\delta^2} \tag{38}$$

where ρ_{TV} is total variation distance, $P^{(k(t))}$ is P 's restriction to sequences of length $k(t)$, $n = \lfloor t/k(t) \rfloor - 1$, $\gamma_1 = (h_P + \epsilon/2)/(h_P + \epsilon)$ and C_1 is a positive constant specific to P (the exact value of which is not important).

The log-likelihood per observation of a Markov chain $\theta \in \Theta_k$ is

$$t^{-1} \log f_\theta(x_1^t) = t^{-1} \log f_\theta(x_1^k) + \sum_{w \in \Xi^k} \sum_{a \in \Xi} \widehat{P}_t^{(k+1)}(wa) \log f_\theta(a|w)$$

where $f_\theta(a|w)$ is of course the probability, according to θ , of producing a after seeing w . By asymptotic equipartition, this is converging a.s. to its expected value, $-h_P - h(\theta)$.

Let $z(\theta) = \max_{w,a} |\log f_\theta(a|w)|$. If $z(\theta) \leq z_0$ and $\rho_{TV}(\widehat{P}_t^{(k+1)}, P^{(k+1)}) \leq \delta$, then $t^{-1} \log f_\theta(x_1^t)$ is within $z_0\delta$ of $-h_P - h(\theta)$. Meanwhile, $t^{-1} \log p(x_1^t)$ is converging a.s. to $-h_P$, and again [60]

$$P(|t^{-1} \log p(X_1^t) - h_P| > \delta) \leq q(t, \delta) 2^{-tC_2\delta} \tag{39}$$

for some $C_2 > 0$ and sub-exponential $q(t, \delta)$. (The details are unilluminating in the present context and thus skipped.)

Define $G(t, z_0)$ as the set of all Markov models whose order is less than or equal to $k(t) - 1$ and whose log transition probabilities do not exceed z_0 , in

symbols

$$G(t, z_0) = \{\theta : z(\theta) \leq z_0\} \cap \left(\bigcup_{j=1}^{k(t)-1} \Theta_j \right)$$

Combining the deviation-probability bounds 38 and 39, for all $\theta \in G(t, z_0)$

$$P \left(\left| \frac{\log R_t(\theta)}{t} - h(\theta) \right| > \delta \right) \leq \frac{\log t}{h + \epsilon} 2(n + 1)^{t\gamma_1} 2^{-\frac{nC_1\delta^2}{4z_0}} + q(t, \delta) 2^{-\frac{tC_2\delta}{2}} \quad (40)$$

These probabilities are clearly summable as $t \rightarrow \infty$, so by the Borel-Cantelli lemma, we have uniform almost-sure convergence of $t^{-1} \log R_t(\theta)$ to $-h(\theta)$ for all $\theta \in G(t, z_0)$.

The sets $G(t, z_0)$ eventually expand to include Markov models of arbitrarily high order, but maintain a constant bound on the transition probabilities. To relax this, let z_t be an increasing function of t , $z(t) \rightarrow \infty$, subject to $z_t \leq C_3 t^{\gamma_2}$ for positive $\gamma_2 < \gamma_1$. Then the deviation probabilities remain summable, and for each t , the convergence of $t^{-1} \log R_t(\theta)$ is still uniform on $G(t, z_t)$. Set $G_t = G(t, z_t)$, and turn to verifying the remaining assumptions.

Start with Assumption 5; take its items in reverse order. So far, the only restriction on the prior Π_0 has been that its support should be the whole of Θ , and that it should have the ‘‘Kullback-Leibler rate property’’, giving positive weight to every set $N_\epsilon = \{\theta : d(\theta) < \epsilon\}$. This, together with the fact that $\lim_t G_t = \Theta$, means that $h(G_t) \rightarrow h(\Theta)$, which is item (3) of the assumption. The same argument also delivers Assumption 7. Item (2), uniform convergence on each G_t , is true by construction. Finally (for this assumption), since $h(\Theta) = 0$, any $\beta > 0$ will do, and there are certainly probability measures where $\Pi_0(G_t^c) \leq \alpha \exp\{-\beta t\}$ for some $\alpha, \beta > 0$. So, Assumption 5 is satisfied.

Only Assumption 6 remains. Since Assumptions 1–3 have already been checked, we can apply Eq. 18 from the proof of Lemma 6 and see that, for each fixed G from the sequence of G_t , for any $\epsilon > 0$, for all sufficiently large t ,

$$t^{-1} \log \Pi_0(GR_t) \leq -h(G) + \epsilon + t^{-1} \log \Pi_0(G) \text{ a.s.}$$

This shows that $\tau(G_t, \delta)$ is almost surely finite for all t and δ , but still leaves open the question of whether for every δ and all sufficiently large t , $t \geq \tau(G_t, \delta)$ (a.s.). Reformulating a little, the desideratum is that for each δ , with probability 1, $t < \tau(G_t, \delta)$ only finitely often. By the Borel-Cantelli lemma, this will happen if $\sum_t P(\tau(G_t, \delta) > t) < \infty$. However, if $\tau(G_t, \delta) > t$, it must be equal to some particular $n > t$, so there is a union bound:

$$\sum_t P(\tau(G_t, \delta) > t) \leq \sum_t \sum_{n=t+1}^\infty P \left(\frac{\log \Pi_0(G_t R_n)}{n} > \delta - h(G_t) \right) \quad (41)$$

From the proof of Lemma 6 (specifically from Eqs. 15, 16 and 17), we can see that by making t large enough, the only way to have the event $n^{-1} \log \Pi_0(G_t R_n) > \delta - h(G_t)$ is to have $|n^{-1} \log R_n(\theta) - h(\theta)| > \delta/2$ everywhere on a positive-measure subset of G_t . But we know from Eq. 40 not only that the inner sum

can be made arbitrarily small by taking t sufficiently large, but that the whole double sum is finite. So $\tau(G_t, \delta) > t$ only finitely often (a.s.).

Acknowledgments

Portions of this work were done while at the Center for the Study of Complex Systems at the University of Michigan, supported by a grant from the James S. McDonnell Foundation, and while at the “Science et Gastronomie 2003” workshop organized by the Institut des Systèmes Complexes and the Laboratoire de l’Informatique du Parallélisme of the École Normale Supérieure de Lyon. I thank Michel Morvan and Christopher Moore for organizing that workshop; Carl Bergstrom, Anthony Brockwell, Chris Genovese, Peter Hoff, Giles Hooker, Martin Nilsson Jacobi, Jay Kadane, Kevin T. Kelly, David Krakauer, Michael Lachmann, Scott E. Page, Mark Schervish, Eric Smith, Surya Tokdar, and Larry Wasserman for discussion; the anonymous referees for valuable suggestions; and above all, Kristina Klinkner.

References

- [1] ALGOET, P. H. AND COVER, T. M. (1988). A sandwich proof of the Shannon-McMillan-Breiman theorem. *Annals of Probability* 16, 899–909. <http://projecteuclid.org/euclid.aop/1176991794>. MR0929085
- [2] ARORA, S., HAZAN, E., AND KALE, S. (2005). The multiplicative weights update method: a meta algorithm and applications. <http://www.cs.princeton.edu/~arora/pubs/MWsurvey.pdf>.
- [3] BADI, R. AND POLITI, A. (1997). *Complexity: Hierarchical Structures and Scaling in Physics*. Cambridge University Press, Cambridge, England. MR1445998
- [4] BARRON, A., SCHERVISH, M. J., AND WASSERMAN, L. (1999). The consistency of posterior distributions in nonparametric problems. *The Annals of Statistics* 27, 536–561. <http://projecteuclid.org/euclid.aos/1018031206>. MR1714718
- [5] BERK, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *Annals of Mathematical Statistics* 37, 51–58. See also correction, volume 37 (1966), pp. 745–746, <http://projecteuclid.org/euclid.aoms/1177699597>. MR0189176
- [6] BERK, R. H. (1970). Consistency a posteriori. *Annals of Mathematical Statistics* 41, 894–906. <http://projecteuclid.org/euclid.aoms/1177696967>. MR0266356
- [7] BLACKWELL, D. AND DUBINS, L. (1962). Merging of opinion with increasing information. *Annals of Mathematical Statistics* 33, 882–886. <http://projecteuclid.org/euclid.aoms/1177704456>. MR0149577
- [8] BÖRGERS, T. AND SARIN, R. (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77, 1–14. MR1484291

- [9] BORKAR, V. S. (2002). Reinforcement learning in Markovian evolutionary games. *Advances in Complex Systems* 5, 55–72. [MR1899835](#)
- [10] CESA-BIANCHI, N. AND LUGOSI, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, England. [MR2409394](#)
- [11] CHAMLEY, C. (2004). *Rational Herds: Economic Models of Social Learning*. Cambridge University Press, Cambridge, England.
- [12] CHARNIAK, E. (1993). *Statistical Language Learning*. MIT Press, Cambridge, Massachusetts.
- [13] CHOI, T. AND RAMAMOORTHI, R. V. (2008). Remarks on consistency of posterior distributions. In *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, B. Clarke and S. Ghosal, Eds. Institute of Mathematical Statistics, Beechwood, Ohio, 170–186. <http://arxiv.org/abs/0805.3248>. [MR2459224](#)
- [14] CHOUDHURI, N., GHOSAL, S., AND ROY, A. (2004). Bayesian estimation of the spectral density of a time series. *Journal of the American Statistical Association* 99, 1050–1059. <http://www4.stat.ncsu.edu/~sghosal/papers/specden.pdf>. [MR2109494](#)
- [15] CRUTCHFIELD, J. P. (1992). Semantics and thermodynamics. In *Nonlinear Modeling and Forecasting*, M. Casdagli and S. Eubank, Eds. Addison-Wesley, Reading, Massachusetts, 317–359.
- [16] DAW, C. S., FINNEY, C. E. A., AND TRACY, E. R. (2003). A review of symbolic analysis of experimental data. *Review of Scientific Instruments* 74, 916–930. <http://www-chaos.engr.utk.edu/abs/abs-rsi2002.html>.
- [17] DĘBOWSKI, Ł. (2006). Ergodic decomposition of excess entropy and conditional mutual information. Tech. Rep. 993, Institute of Computer Science, Polish Academy of Sciences (IPI PAN). http://www.ipipan.waw.pl/~ldebowsk/docs/raporty/ee_report.pdf.
- [18] DIACONIS, P. AND FREEDMAN, D. (1986). On the consistency of Bayes estimates. *The Annals of Statistics* 14, 1–26. <http://projecteuclid.org/euclid.aos/1176349830>. [MR0829555](#)
- [19] DOOB, J. L. (1949). Application of the theory of martingales. In *Colloques Internationaux du Centre National de la Recherche Scientifique*. Vol. 13. Centre National de la Recherche Scientifique, Paris, 23–27. [MR0033460](#)
- [20] DYNKIN, E. B. (1978). Sufficient statistics and extreme points. *Annals of Probability* 6, 705–730. <http://projecteuclid.org/euclid.aop/1176995424>. [MR0518321](#)
- [21] EARMAN, J. (1992). *Bayes or Bust? A Critical Account of Bayesian Confirmation Theory*. MIT Press, Cambridge, Massachusetts. [MR1170349](#)
- [22] EICHELNBACHER, P. AND GANESH, A. (2002). Moderate deviations for Bayes posteriors. *Scandinavian Journal of Statistics* 29, 153–167. [MR1894388](#)
- [23] FISHER, R. A. (1958). *The Genetical Theory of Natural Selection*, Second ed. Dover, New York. First edition published Oxford: Clarendon Press, 1930. [MR1785121](#)

- [24] FRASER, A. M. (2008). *Hidden Markov Models and Dynamical Systems*. SIAM Press, Philadelphia. MR2451733
- [25] GEMAN, S. AND HWANG, C.-R. (1982). Nonparametric maximum likelihood estimation by the method of sieves. *The Annals of Statistics* 10, 401–414. <http://projecteuclid.org/euclid.aos/1176345782>. MR0653512
- [26] GHOSAL, S., GHOSH, J. K., AND RAMAMOORTHI, R. V. (1999). Consistency issues in Bayesian nonparametrics. In *Asymptotics, Nonparametrics and Time Series: A Tribute to Madan Lal Puri*, S. Ghosh, Ed. Marcel Dekker, 639–667. <http://www4.stat.ncsu.edu/~sghosal/papers/review.pdf>. MR1724711
- [27] GHOSAL, S., GHOSH, J. K., AND VAN DER VAART, A. W. (2000). Convergence rates of posterior distributions. *Annals of Statistics* 28, 500–531. <http://projecteuclid.org/euclid.aos/1016218228>. MR1790007
- [28] GHOSAL, S. AND TANG, Y. (2006). Bayesian consistency for Markov processes. *Sankhya* 68, 227–239. http://sankhya.isical.ac.in/search/68_2/2006010.html. MR2303082
- [29] GHOSAL, S. AND VAN DER VAART, A. (2007). Convergence rates of posterior distributions for non-iid observations. *Annals of Statistics* 35, 192–223. <http://arxiv.org/abs/0708.0491>. MR2332274
- [30] GHOSH, J. K. AND RAMAMOORTHI, R. V. (2003). *Bayesian Nonparametrics*. Springer Verlag, New York. MR1992245
- [31] GRAY, R. M. (1988). *Probability, Random Processes, and Ergodic Properties*. Springer-Verlag, New York. <http://ee.stanford.edu/~gray/arp.html>. MR0918767
- [32] GRAY, R. M. (1990). *Entropy and Information Theory*. Springer-Verlag, New York. <http://ee.stanford.edu/~gray/it.html>. MR1070359
- [33] HALDANE, J. B. S. (1954). The measurement of natural selection. In *Proceedings of the 9th International Congress of Genetics*. Vol. 1. 480–487.
- [34] HOFBAUER, J. AND SIGMUND, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, England. MR1635735
- [35] KALLENBERG, O. (2002). *Foundations of Modern Probability*, Second ed. Springer-Verlag, New York. MR1876169
- [36] KITCHENS, B. AND TUNCEL, S. (1985). *Finitary Measures for Subshifts of Finite Type and Sofic Systems*. Memoirs of the American Mathematical Society, Vol. 338. American Mathematical Society, Providence, Rhode Island. MR0818917
- [37] KITCHENS, B. P. (1998). *Symbolic Dynamics: One-sided, Two-sided and Countable State Markov Shifts*. Springer-Verlag, Berlin. MR1484730
- [38] KLEIJN, B. J. K. AND VAN DER VAART, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. *Annals of Statistics* 34, 837–877. <http://arxiv.org/math.ST/0607023>. MR2283395
- [39] KNIGHT, F. B. (1975). A predictive view of continuous time processes. *Annals of Probability* 3, 573–596. <http://projecteuclid.org/euclid.aop/1176996302>. MR0383513

- [40] KROGH, A. AND VEDELSBY, J. (1995). Neural network ensembles, cross validation, and active learning. In *Advances in Neural Information Processing 7 [NIPS 1994]*, G. Tesauro, D. Touretsky, and T. Leen, Eds. MIT Press, Cambridge, Massachusetts, 231–238. <http://books.nips.cc/papers/files/nips07/0231.pdf>.
- [41] LIAN, H. (2007). On rates of convergence for posterior distributions under misspecification. E-print, arxiv.org. <http://arxiv.org/abs/math.ST/0702126>.
- [42] LIJOI, A., PRÜNSTER, I., AND WALKER, S. G. (2007). Bayesian consistency for stationary models. *Econometric Theory* 23, 749–759. [MR2364385](#)
- [43] LIND, D. AND MARCUS, B. (1995). *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, Cambridge, England. [MR1369092](#)
- [44] MARTON, K. AND SHIELDS, P. C. (1994). Entropy and the consistent estimation of joint distributions. *The Annals of Probability* 22, 960–977. Correction, *The Annals of Probability*, 24 (1996): 541–545, <http://projecteuclid.org/euclid.aop/1176988736>. [MR1288138](#)
- [45] MCALLISTER, D. A. (1999). Some PAC-Bayesian theorems. *Machine Learning* 37, 355–363.
- [46] MEIR, R. (2000). Nonparametric time series prediction through adaptive model selection. *Machine Learning* 39, 5–34. <http://www.ee.technion.ac.il/~rmeir/Publications/MeirTimeSeries00.pdf>.
- [47] ORNSTEIN, D. S. AND WEISS, B. (1990). How sampling reveals a process. *The Annals of Probability* 18, 905–930. <http://projecteuclid.org/euclid.aop/1176990729>. [MR1062052](#)
- [48] PAGE, S. E. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, Princeton, New Jersey.
- [49] PAPANGELOU, F. (1996). Large deviations and the Bayesian estimation of higher-order Markov transition functions. *Journal of Applied Probability* 33, 18–27. <http://www.jstor.org/stable/3215260>. [MR1371950](#)
- [50] PERRY, N. AND BINDER, P.-M. (1999). Finite statistical complexity for sofic systems. *Physical Review E* 60, 459–463.
- [51] RIVERS, D. AND VUONG, Q. H. (2002). Model selection tests for nonlinear dynamic models. *The Econometrics Journal* 5, 1–39. [MR1909299](#)
- [52] ROY, A., GHOSAL, S., AND ROSENBERGER, W. F. (2009). Convergence properties of sequential Bayesian d -optimal designs. *Journal of Statistical Planning and Inference* 139, 425–440. [MR2474016](#)
- [53] RYABKO, D. AND RYABKO, B. (2008). Testing statistical hypotheses about ergodic processes. E-print, arxiv.org, 0804.0510. <http://arxiv.org/abs/0804.0510>.
- [54] SATO, Y. AND CRUTCHFIELD, J. P. (2003). Coupled replicator equations for the dynamics of learning in multiagent systems. *Physical Review E* 67, 015206. <http://arxiv.org/abs/nlin.AO/0204057>.
- [55] SCHERVISH, M. J. (1995). *Theory of Statistics*. Springer Series in Statistics. Springer-Verlag, Berlin. [MR1354146](#)

- [56] SCHWARTZ, L. (1965). On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 4, 10–26. [MR0184378](#)
- [57] SHALIZI, C. R. AND CRUTCHFIELD, J. P. (2001). Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics* 104, 817–879. <http://arxiv.org/abs/cond-mat/9907176>. [MR1853995](#)
- [58] SHALIZI, C. R. AND KLINKNER, K. L. (2004). Blind construction of optimal nonlinear recursive predictors for discrete sequences. In *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference (UAI 2004)*, M. Chickering and J. Y. Halpern, Eds. AUAI Press, Arlington, Virginia, 504–511. <http://arxiv.org/abs/cs.LG/0406011>.
- [59] SHEN, X. AND WASSERMAN, L. (2001). Rates of convergence of posterior distributions. *Annals of Statistics* 29, 687–714. <http://projecteuclid.org/euclid.aos/1009210686>. [MR1865337](#)
- [60] SHIELDS, P. C. (1996). *The Ergodic Theory of Discrete Sample Paths*. American Mathematical Society, Providence, Rhode Island. [MR1400225](#)
- [61] STRELIOFF, C. C., CRUTCHFIELD, J. P., AND HÜBLER, A. W. (2007). Inferring Markov chains: Bayesian estimation, model comparison, entropy rate, and out-of-class modeling. *Physical Review E* 76, 011106. <http://arxiv.org/math.ST/0703715>. [MR2365498](#)
- [62] VARN, D. P. AND CRUTCHFIELD, J. P. (2004). From finite to infinite range order via annealing: The causal architecture of deformation faulting in annealed close-packed crystals. *Physics Letters A* 324, 299–307. <http://arxiv.org/abs/cond-mat/0307296>.
- [63] VIDYASAGAR, M. (2003). *Learning and Generalization: With Applications to Neural Networks*, Second ed. Springer-Verlag, Berlin. [MR1938842](#)
- [64] VUONG, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57, 307–333. <http://www.jstor.org/pss/1912557>. [MR0996939](#)
- [65] WALKER, S. (2004). New approaches to Bayesian consistency. *Annals of Statistics* 32, 2028–2043. <http://arxiv.org/abs/math.ST/0503672>. [MR2102501](#)
- [66] WEISS, B. (1973). Subshifts of finite type and sofic systems. *Monatshefte für Mathematik* 77, 462–474. [MR0340556](#)
- [67] XING, Y. AND RANNEYBY, B. (2008). Both necessary and sufficient conditions for Bayesian exponential consistency. <http://arxiv.org/abs/0812.1084>.
- [68] ZHANG, T. (2006). From ϵ -entropy to KL-entropy: Analysis of minimum information complexity density estimation. *Annals of Statistics* 34, 2180–2210. <http://arxiv.org/math.ST/0702653>. [MR2291497](#)