

Adaptive Independent Metropolis-Hastings by Fast Estimation of Mixtures of Normals

Paolo Giordani
Research Department, Sveriges Riksbank
paolo.giordani@riksbank.se

Robert Kohn
Australian School of Business
University of New South Wales

January 19, 2009

Abstract

Adaptive Metropolis-Hastings samplers use information obtained from previous draws to tune the proposal distribution automatically and repeatedly. Adaptation needs to be done carefully to ensure convergence to the correct target distribution because the resulting chain is not Markovian. We construct an adaptive independent Metropolis-Hastings sampler that uses a mixture of normals as a proposal distribution. To take full advantage of the potential of adaptive sampling our algorithm updates the mixture of normals frequently, starting early in the chain. The algorithm is built for speed and reliability and its sampling performance is evaluated with real and simulated examples. Our article outlines conditions for adaptive sampling to hold and gives a readily accessible proof that under these conditions the sampling scheme generates iterates that converge to the target distribution.

Keywords: Clustering; Gibbs sampling; Markov chain Monte Carlo; Semiparametric regression models; State space models.

1 Introduction

Bayesian methods using Markov chain Monte Carlo (MCMC) simulation have greatly influenced statistical practice over the past twenty years because of their ability to estimate complex models and produce finite sample inference. A key component in implementing MCMC simulation is the Metropolis-Hastings (MH) method (Metropolis et al. 1953; Hastings 1970), which requires the specification of one or more proposal distributions. The speed at which the chain converges to the posterior distribution and its ability to move efficiently across the state space depend crucially on whether the proposal distribution(s) provide good approximations to the target distributions, either locally or globally. Given the key role played by proposal distributions, it is natural to use experience from previous

draws to *adapt* the proposal to the target. Our article considers adaptive sampling that is subject to theoretical rules which ensure that the iterates converge to realizations from the correct target (posterior) distribution.

The literature on adaptive MCMC methods follows three main strands. Adaptation by *regeneration* stems from the work of Gilks et al. (1998). Adapting on the target distribution is proposed by Wang and Landau (2001) and generalized by Liang et al. (2007) and Atchade and Liu (2007). Our article focuses exclusively on *diminishing adaptation* schemes. Important theoretical advances in diminishing adaptation were made by Holden (1998), Haario et al. (2001), Andrieu and Robert (2001), Andrieu and Moulines (2006), Andrieu et al. (2005), Atchade and Rosenthal (2005), Nott and Kohn (2005) and Roberts and Rosenthal (2007). The proofs of convergence for strict adaptive sampling are more complex than for the non adaptive case as the iterates are not Markovian because the MH kernel can depend on the entire history of the draws. Although more theoretical work on adaptive sampling can be expected, the existing body of results provides sufficient justification and guidelines to build adaptive MH samplers for challenging problems.

Research is now needed on how to design efficient and reliable adaptive samplers

Our paper contributes to the development of algorithms for adaptive independent MH sampling in continuous state spaces. Increased sampling efficiency is one important goal, particularly in cases where current best practice (typically some version of random walk Metropolis or Gibbs sampling) is less than satisfactory. But equally important achievements of adaptive schemes may be to expand the set of problems that can be handled efficiently by general purpose samplers and to reduce coding effort. In particular, adaptive schemes can reduce dependence on the use of conjugate priors. Such priors make it easier to implement MCMC schemes, but are less necessary when using adaptive sampling.

Our adaptive sampling approach is built on four main ideas. The first is to combine preliminary exploration of the target distribution and adaptive sampling into one estimation procedure. The second is to estimate mixtures of Gaussians from the history of the draws and use them as proposal distributions for independent MH in all parts of the estimation. The third is to perform this estimation frequently, particularly during the early part of the estimation, a strategy that we call *intensive adaptation*. The fourth is to ensure that the theoretical conditions for the correct ergodic behavior of the sampler are respected during adaptation. To apply these ideas successfully, estimation of the mixture parameters needs to be fast, reliable, and robust. We achieve a good balance of these goals by carefully selecting and tailoring to our needs algorithms developed in the clustering literature.

We study the performance of our adaptive sampler in two examples in which commonly used Gibbs schemes can be very inefficient and compare it with an adaptive random walk Metropolis sampler proposed by Roberts and Rosenthal (2006) that builds on the work of Haario et al. (2001).

Our paper also provides conditions and outlines a proof that our adaptive sampling scheme converges to the correct target distribution and gives convergence rates.

Our working paper available at http://arxiv.org/PS_cache/arxiv/pdf/0801/0801.1864v1.pdf contains several other applications and a more extensive discussion.

2 Some theory for adaptive sampling

A *diminishing adaptation* MH sampler performs the accept/reject step like a standard MH, but updates the proposal distribution using the history of the draws. This updating is 'diminishing' in the sense that changes in the proposal distribution tend to zero asymptotically in the number of iterations. In practice, this usually means that the proposal distribution itself settles down asymptotically.

This section outlines the theoretical framework for adaptive independent Metropolis-Hastings sampling as used in our article that gives some support for our practice. The appendix outlines proofs of the theoretical results, which extend similar results in Nott and Kohn (2005) for the case of a finite state space. Our aim is to provide simple accessible proofs that will help to popularize the adaptive methodology. All densities in this section are with respect to Lebesgue measure or counting measure, which we denote as $\nu\{\cdot\}$.

Let \mathcal{Z} be a sample space and $\mathcal{H}(z)$ a target density on \mathcal{Z} . We use the following adaptive MH scheme to construct a sequence of random variables $\{Z_n; n \geq 2\}$ whose distribution converges to $\mathcal{H}(z)$. Z_0 and Z_1 are generated from $g_0(z)$ which is defined below. For $n \geq 1$, let $q_n(z; \nu_n)$ be a proposal density for generating Z_{n+1} , where ν_n is a parameter vector that is based on $Z_0 = z_0; \dots; Z_{n-1} = z_{n-1}$. Thus, given $Z_n = z$, we generate $Z_{n+1} = z'$ from q_n , and then with probability

$$\alpha_n(z; z') = \min\left(1; \frac{\mathcal{H}(z') q_n(z; \nu_n)}{\mathcal{H}(z) q_n(z'; \nu_n)}\right) \quad (1)$$

we take $Z_{n+1} = z'$; otherwise we take $Z_{n+1} = z$. Our choice of $q_n(z; \nu_n)$ is of the form

$$q_n(z; \nu_n) = \alpha_1 g_0(z) + (1 - \alpha_1) g_n(z; \nu_n) \quad (2)$$

where $0 < \alpha_1 < 1$. The density $g_0(z)$ is constant and the density $g_n(z; \nu_n)$ has parameter vector ν_n that evolves with the iterates. The form of both densities as used in our article is described more fully below.

We assume that there exists a constant $K > 0$ such that for all $z \in \mathcal{Z}$

$$\frac{q(z)}{g_0(z)} \leq K \quad \text{and} \quad \frac{g_n(z; \cdot, n)}{g_0(z)} \leq K ; \quad (3)$$

and

$$\sup_{z \in \mathcal{Z}} \left| \left(g_n(z; \cdot, n) - g_{n+1}(z; \cdot, n+1) \right) - g_0(z) \right| = a_n \quad (4)$$

where $a_n = O(n^{-r})$ for some $r > 0$ almost surely. In relation to the dominance condition (3), we note that in the non-adaptive case, that is $q_n(z; \cdot, n) = q(z)$ for all n , Mengersen and Tweedie (1996) show that $q(z) \leq K$ for all z is a necessary and sufficient condition for geometric ergodicity.

Under conditions (3) and (4), the following results are proved in Appendix 2.

Theorem 1 *For all measurable subsets A*

$$\sup_{A \subset \mathcal{Z}} \left| \Pr(Z_n \in A) - q(A) \right| \rightarrow 0 \quad \text{as} \quad n \rightarrow \infty ; \quad (5)$$

Theorem 2 *Suppose that $h(z)$ is a measurable function that is square integrable with respect to the density g_0 . Then, almost surely,*

$$\frac{1}{n} \sum_{j=1}^n h(Z_j) \rightarrow E_{\pi}(h(Z)) \quad \text{as} \quad n \rightarrow \infty ; \quad (6)$$

We now describe the construction of $g_n(z; \cdot, n)$ in the article. Let $g_n^*(z; \cdot, n^*)$ be a mixture of normals with parameters \cdot, n^* which is obtained using k-harmonic means clustering as described in section 4 and appendix 1. Let $\tilde{g}_n^*(z; \cdot, n^*)$ be a second mixture of normals having the same component weights and means as $g_n^*(z; \cdot, n^*)$, but with its component variances inflated by a factor $k > 1$. We set

$$g_n(z; \cdot, n) = \frac{1}{2} g_n^*(z; \cdot, n^*) + (1 - \frac{1}{2}) \tilde{g}_n^*(z; \cdot, n^*) ; \quad (7)$$

where $\beta_2 = \beta_2(1 - \beta_1)$ with β_1 defined in (2), $\beta_1 > 0$; $\beta_2 > 0$, and $\beta_1 + \beta_2 < 1$. Thus,

$$q_n(z; \nu_n) = \beta_1 g_0(z) + \beta_2 g_n^*(z; \nu_n^*) + (1 - \beta_1 - \beta_2) g_n^*(z; \nu_n^*) :$$

We note that $g_n(z; \nu_n)$ is also a mixture of normals with parameters ν_n , and we say that it is obtained by 'stretching and fattening' the tails of $g_n^*(z; \nu_n^*)$. This strategy for obtaining heavier tailed mixtures is used extensively in our paper.

Conditions (3) and (4) provide guidance for our algorithms and are checked informally in our work. The target $\mathcal{H}(z)$ in our applications is the posterior density of the parameters and so is proportional to $\mathcal{H}_0(z)\mathcal{H}_1(z)$, where $\mathcal{H}_0(z)$ is the prior and $\mathcal{H}_1(z)$ is the likelihood. Thus, if $\mathcal{H}_0(z)=g_0(z)$ is bounded for all n and all $z \in \mathcal{Z}$ and if the maximum likelihood estimator exists then $\mathcal{H}(z)=g_0(z)$ is also bounded. If we choose $g_0(z)$ such that $g_n^*(z; \nu_n^*)=g_0(z) \leq K$ for all n and $z \in \mathcal{Z}$ then $g_n(z; \nu_n)$ will also be bounded. Finally, we check informally that (4) holds by checking that the iterates $g_n^*(z; \nu_n^*)$ converge to a fixed proposal. In our experience, this almost always happens. However, we now show how to more formally ensure that the conditions (3) and (4) hold, while the practical performance of the algorithm remains very much as above. First, we add an extremely heavy tailed component to $g_0(z)$ with a probability that is extremely small. Second, let $\tilde{g}_0(z) = g_0(z)$ and for $n \geq 1$ define $\tilde{g}_n(z; \nu_n^1) = (1 - a_n)\tilde{g}_{n-1}(z; \nu_{n-1}^1) + a_n g_n(z; \nu_n)$, with $a_n = (1 - \epsilon)^n$ and ϵ very close to zero and positive. This means that for $n \geq 1$, ν_n^1 is defined recursively in terms of ν_{n-1}^1 . Third, we redefine the proposal density as $q_n(z; \nu_n) = \beta_1 g_0(z) + (1 - \beta_1)\tilde{g}_n(z; \nu_n^1)$. This ensures that $\mathcal{H}_0(z)=g_0(z)$ and $\tilde{g}_n(z; \nu_n^1)$ are bounded in n and $z \in \mathcal{Z}$ and that the diminishing adaptation condition holds, i.e.

$$\sup_{z \in \mathcal{Z}} \left| \left(\tilde{g}_n(z; \nu_n^1) - \tilde{g}_{n+1}(z; \nu_{n+1}^1) \right) - g_0(z) \right| = a_n :$$

Section 7 and appendix F of Andrieu and Moulines (2006) give general convergence results for adaptive independent Metropolis-Hastings and Roberts and Rosenthal (2007) give an elegant proof of the convergence of adaptive sampling schemes. However, we believe that readers may find the conditions (3) and (4) and the proofs of Theorems 1 and 2 easier to understand for the methodology proposed in our article.

3 Implementation of the adaptive sampling scheme

This section describes the implementation of the sampling scheme. We anticipate that readers will use this as a basis for their own experimentation. The adaptive sampling scheme is run in two phases, a preliminary phase where the conditions for ergodicity are not enforced and a strict adaptive stage that enforces ergodicity. To make the structure of the algorithm clearer, we first describe it in pseudo-code and then discuss it in more detail.

Pseudo-code. Let $q_n(z) = \alpha_1 g_0(z) + (1 - \alpha_1) g_n(z)$; $g_0(z) = 0.6 \hat{A}_0(z) + 0.4 \hat{A}_0^*(z)$, where $\hat{A}_0(z)$ is a mixture of normals, initialized at iteration $n = 1$ by a Laplace expansion (in which case $\hat{A}_0(z)$ is a multivariate normal), or by the prior or by a density estimated from a preliminary MCMC run if the Laplace approximation is unavailable. The density $\hat{A}_0^*(z)$ is a mixture of normals with the same parameters as $\hat{A}_0(z)$ except that the covariance matrices are multiplied by 25. Let $g_n(z) = \alpha_2 g_n^*(z) + (1 - \alpha_2) g_n^{\dagger}(z)$; where $g_n^*(z)$ is a mixture of normals and $g_n^{\dagger}(z)$ is a mixture of normals with the same parameters except that the covariance matrices are multiplied by a scalar k . For notational convenience we omit to show in the pseudo-code the dependence of the densities $q_n(z)$; $g_n(z)$ and $g_n^*(z)$ on their parameters.

Let A_n denote the number of accepted MH draws up to but not including iteration n , and $S_n(M)$ the smallest MH acceptance probability in iterations $n - M$ to $n - 1$.

1. At iteration n , $q_n(z)$ is used as a proposal to update z as in (2).
2. Define the *preliminary phase* to start at $n = 1$ and end when $S_n(M) > \epsilon_M$. The density $\hat{A}_0(z)$ is updated once at the end of the preliminary phase, and set equal to g_{last}^* , the last estimated mixture of normals in the preliminary phase.
3. The mixture of normals $g_n^*(z)$ is set equal to $g_{n-1}^*(z)$ unless $A_n > 5 \dim(z)$ and *either*
 - $n - n^*$ belongs to a predetermined set of positive updating times given below, where $n^* = n : A_n = 5 \dim(z)$, or
 - The average MH probability in the last L iterations is lower than ϵ_L (preliminary phase only).
4. Otherwise $g_n^*(z)$ is updated as in section 4 and appendix 1.

We set $k = 16; \alpha_1 = 0.05; \alpha_2 = 0.15; \alpha'_2 = \alpha_2(1 - \alpha_1); L = 10; \theta_L = 0.1; M = 20; \theta_M = 0.02$.

End of pseudo-code

We find these values of $k; \alpha_1; \alpha_2; L; \theta_L; M$ and θ_M work well, but we do not claim they are optimal. We conjecture that the speed of convergence and efficiency of our sampler can be further improved with a more careful (and possibly adaptive) choice of these parameters. Any other value of k in the range 9 to 25 and of α_1 and α_2 in the range 0.05 to 0.3 worked well for the examples given in the paper.

During the preliminary phase, we *first* estimate the k-harmonic means mixture after a given number of *accepted* draws in order to ensure that the estimated covariance matrices are positive definite and estimated with sufficient accuracy. When there are 2 to 4 unknown parameters as in the inflation example we first estimate the k-harmonic means mixture after 20 *accepted* draws. If our parameter space is bigger then we would increase that number appropriately. The estimation after a given number of accepted draws is only done *once* throughout the sampling scheme. We then re-estimate the mixture after 50, 100, . . . , 350, 400, 500, . . . , 1000 and then every 1000 draws thereafter. We also recommend updating the proposal in the preliminary stage following a period of low acceptance probabilities in the MH step. Specifically, we re-estimate the mixture parameters if the average acceptance probability in the last L iterations is lower than θ_L , where we set L and θ_L as above. Low acceptance probabilities signal a poor fit of the proposal, and it is therefore sensible to update the proposal to give it a better chance of covering the area around the current parameter value. Since it is unclear that this does not violate any of the conditions required for the ergodicity of our adaptive sampler, we limit the updating of the proposal at endogenously chosen points to the preliminary phase, after which the proposal is updated only at predetermined intervals. The end of the preliminary adaptation phase could be set ex-ante, but we prefer to determine it endogenously by requiring the smallest acceptance probability in the last M iterations to be higher than θ_M ; where M and θ_M are set as above. During the second phase (period of strict adaptation), we update the proposal every 1000 draws.

The updating schedule given above is also not optimal in any specific sense. The principle is to update more frequently in the initial phases of the chain. We would update more (less) frequently if the likelihood was very expensive (inexpensive) to compute compared

to updating the proposal. In our experience, the estimation frequency typically does not affect the performance of the algorithms dramatically. The fact that we update the proposal when it has performed poorly in recent iterates is very helpful in this regard since otherwise the MH can reject for long stretches if the updates are infrequent. Frequent updates are important if the initial proposal is poor and/or the target distribution is multimodal. Fortunately, updating the proposal when it is not performing well largely endogenizes the updating schedule.

The estimation of the mixture of normals can become slow when the number of iterations is large. To avoid this problem, after 1000 accepted draws we only add every j -th draw to the sample used to estimate the mixture parameters, where j is chosen so that the mixture is not estimated on more than 10000 observations.

4 A clustering algorithm for fast estimation of mixtures of normals in adaptive IMH

Finite mixtures of normals are an attractive option to construct the proposal density because they can approximate any continuous density arbitrarily well and are fast to sample from and evaluate. See McLachlan and Peel (2000) for an extensive treatment of finite mixture models.

However, estimating mixtures of normals is already a difficult problem when an independent and identically distributed sample from the target is given and estimation needs to be performed only once: the likelihood goes to infinity whenever a component has zero variance (an even more concrete possibility if, as unavoidable in IMH, some observations appear more than once), and its maximization, whether by the EM algorithm or directly, is plagued by local modes. Although several authors have made substantial advances in dealing with these problems (e.g. Figueredo and Jain 2002; Ueda, Nakano, Ghahramani, and Hinton 2000; Verbeek, Vlassis, and Krose 2003), in our experience the EM algorithm is not sufficiently reliable when the sample is small and contains a non-trivial share of rejected draws. The inevitable short runs of rejections give rise to small clusters with zero covariance matrix at which the EM algorithm frequently gets stuck. Moreover, EM-based algorithm

are computationally expensive and slow to converge, which makes them less attractive when the proposal is to be updated frequently.

We have therefore limited our attention to algorithms that estimate mixtures of normals quickly and without explicitly computing the covariance matrix of each component (for robustness). Within this class, the *k-means* algorithm is the most popular algorithm. We employ the *k-harmonic means* algorithm, an extension of the k-means algorithm that allows for soft membership. Degeneracies can be easily prevented, so the algorithm is remarkably robust even in the presence of a long series of rejections. The number of clusters is chosen with the BIC criterion. The increase in speed and reliability is paid for with a decreased τ to the target, as the family of k-means algorithms performs best when an optimal τ requires all components of the mixture to have the same probability and covariance matrix (see Bradley and Fayyad 1998, for a discussion). Hamerly and Elkan (2002) show that the performance of k-harmonic means deteriorates less rapidly than alternatives of similar computational cost with departures from these ideal conditions. An outline of the k-harmonic means algorithms is given in Appendix 1.

Although the k-harmonic means algorithm is less sensitive to initialization than either k-means or EM (Hamerly and Elkan 2002), in an unsupervised environment it is important to have good starting values. We have found the algorithm of Bradley and Fayyad (1998) to perform very well and at a low computational cost.

When most parameters are nearly normally distributed, fitting a mixture of normals to all the parameters is problematic in the sense that the chances of finding a local mode with all parameters normally distributed is quite high (though depending on the starting value of course). This is true of clustering algorithms and also of EM-based algorithms. To improve the performance of the sampler in these situations, we divide the parameter vector μ into two groups, μ_1 and μ_2 , where parameters in μ_1 are classified as approximately normal and parameters in μ_2 are skewed. Our rule of thumb is to place a parameter in the 'normal' group if its marginal distribution has $|s| < 0.2$; where s is the skeweness. Our fattening the tails of the mixture should handle the kurtosis, though this would optimally be done with mixtures of more flexible distributions than the normal.

A normal is then fit to the first group and a mixture of p normals to the second. For

μ_1 , we can compute the mean θ_1 and covariance matrix S_{θ_1} inexpensively from the draws. For μ_2 , we fit a mixture of normals as detailed below, estimating probabilities $\omega_1, \dots, \omega_p$, means μ_1, \dots, μ_p and covariance matrices S_1, \dots, S_p . We then need to build a mixture for $\mu = \{\mu_1, \mu_2\}$. The mean is straightforward: for the normal parameters, all components have the same mean. The diagonal blocks of the covariance matrices Σ_i corresponding to $\text{var}(\mu_1)$ and $\text{var}(\mu_2)$ for component i are also straightforward. The off-diagonal blocks of Σ_i corresponding to $\text{cov}(\mu_1, \mu_2)$ is computed as

$$\Sigma_i^{-12} = \sum_{t=1}^n \omega_{i,t}^* [(\mu_{1,t} - \theta_1)(\mu_{2,t} - \mu_i)] = \sum_{t=1}^n \omega_{i,t}^*$$

where $\omega_{i,t}^* = \Pr(K_t = i | \mu_{2,t})$ is the probability of $\mu_{2,t}$ coming from the i -th component.

If the proposal distribution is normal then it is computationally inexpensive to update it at every iteration. It is tempting to update a mixture of normals proposal with an on-line estimation procedure such as the on-line EM algorithm proposed in Andrieu and Moulines (2006). The advantage of on-line estimation is that the parameters of the mixture are updated recursively, so the proposal itself is updated at each iteration at a very small computational cost. However, on-line estimation of the mixture parameters in AIMH has a number of serious shortcomings. The estimates are inefficient compared to batch estimators because each data point is used only once, which corresponds to requiring a batch optimization to converge in one step. The loss of efficiency is more severe in small samples, that is, in the early phases of the chain. Direct estimation of the mixture component covariance matrices often leads to numerical problems in the early phases of the chain given that rejections in MH produce degenerate clusters. Finally, a limitation of on-line estimators is related to the fact that they are a form of stochastic gradient descent (see Spall (2003) for an introduction). When the function to be maximized is multimodal (as is typically the case with mixtures) on-line estimates are in general sensitive to the order of the draws, with initial draws having heavier influence than later draws in determining the mode at which estimates settle down. We have verified empirically that the quality of solutions given by several on-line algorithms deteriorates rapidly if the initial observations are not representative of the entire target distribution. This makes on-line algorithms unsuitable for use in the early, exploratory phases of the chain, though they can work well if the initial

proposal distribution already provides a reasonably good approximation of the target and the acceptance rates are sufficiently high.

Since we are opting for batch estimators, it is too costly to update the proposal at each iteration. We update it at predetermined numbers of iterations, more frequently in the earlier stages of the chain. Implementation details for the empirical examples are given in section 3.

We make two further comments on Andrieu and Moulines (2006). First, they propose to keep the number of components in the mixture constant, whereas we have found it advantageous to select the number adaptively as outlined in appendix 1. Second, they outline a proposed approach to using mixtures as proposal densities, but do not report on the empirical performance of their proposal.

5 Applications

State space models and nonparametric models are ideal initial applications for AIMH schemes. Although they can have a large number of parameters or latent variables, it often happens that conditional on a small subset of these, most parameters and latent variables can be integrated out or have known analytical form from which we can generate them. It is therefore often possible to draw all parameters in one or two blocks. Exploiting these features, it is also often inexpensive to find the posterior mode, possibly for a simplified version of the model, and therefore obtain a reasonable initialization of the proposal distribution. Finally, the standard approach based on Gibbs and Metropolis-within-Gibbs can be very inefficient, particularly for state space models (see Fruhwirth-Schnatter 2004).

For each of our applications we checked the results of the adaptive sampling scheme by re-running the sampler at a number (at least 5) of different starting points using a fixed proposal based on the last mixture of normals update in the second stage of the adaptive sampling. In all cases we got very similar results to those obtained using strict adaptation, where by 'very similar' we mean that the posterior means of the parameters were within three Monte Carlo standard deviations of those obtained using adaptive sampling.

For our examples we define the inefficiency of a sampling scheme as the factor by which the number of iterates needs to increase to give the same precision (standard error) as a

sampler generating independent draws. For two sampling schemes A and B, we define the inefficiency of scheme B relative to A as the factor by which it is necessary to increase the running time of B in order for it to obtain the same accuracy as A. It is computed as the inefficiency factor of B times its run time per iteration divided by the inefficiency factor of A times its run time per iteration.

In the examples below we compare the performance of the AIMH sampler to the following version of the Haario et al. (2001) adaptive random walk Metropolis sampler proposed on page 3 of Roberts and Rosenthal (2006). Specifically, let μ be the parameters in the model, $\hat{\mu}$ the posterior mode and V the variance covariance matrix of the Laplace approximation to the posterior. Then at iteration j the proposal distribution is given by

$$Q_j(\mu^c; \cdot) = N(\mu^c; (0.1)^2 V = d) \quad \text{if} \quad j < 5d;$$

$$Q_j(\mu^c; \cdot) = (1 - \alpha) N(\mu^c; (2.38)^2 S_j = d) + \alpha N(\mu^c; (0.1)^2 I_{d=d}) \quad \text{if} \quad j \geq 5d;$$

where $N(\mu; V)$ is the normal density with mean μ and covariance matrix V , μ^c is the current value of μ , d is the dimension of μ , $\alpha = 0.05$ and S_j is the current empirical estimate of the covariance matrix of the target distribution based on the iterates thus far. In all cases we initialized this sampler at the posterior mode.

5.1 Time-varying parameter autoregressive models

Consider the following time-varying parameter first order autoregressive (AR(1)) process (the extension to a more general autoregressive process is straightforward):

$$y_t = c_t + \frac{1}{2} y_{t-1} + \frac{3}{4} \epsilon_t^2; \quad c_t = c_{t-1} + \frac{1}{10} u_t \quad \text{and} \quad \frac{1}{2}_t = \frac{1}{2}_{t-1} + \frac{1}{10} v_t; \quad (8)$$

where $\epsilon_t; u_t; v_t$ are all $nid(0; 1)$. The model has three parameters $(\frac{3}{4}_\epsilon; \frac{2}{10}; \frac{2}{10})$; while c_0 and $\frac{1}{2}_0$ can be treated either as parameters or (our choice) as states. Given conjugate priors (inverse gamma for the parameters, and normal for c_0 and $\frac{1}{2}_0$), Gibbs sampling is straightforward (Carter and Kohn 1994). Fruhwirth-Schnatter (2004) reports that based on the autocorrelations of the iterates, Gibbs sampling can be very inefficient for these models. In the following application we also find that the Gibbs draws are highly autocorrelated and,

by comparing posterior statistics from Gibbs sampling and from our AIMH, we also find that the autocorrelations do not reveal the full extent of the problem.

5.1.1 Application: US CPI inflation

We apply the model to quarterly U.S. CPI inflation for the period 1960-2005 (184 observations). The annualized quarterly CPI inflation is defined as $400(P_t = P_{t-1} - 1)$; where P_t is aggregated from monthly data (averages) on Consumer Price Index For All Urban Consumers: All Items, seasonally adjusted, Series ID CPIAUCSL, Source: U.S. Department of Labor: Bureau of Labor Statistics. We use rather dispersed inverse gamma priors for $\mathcal{V}_\epsilon^2, \sigma_0^2, \sigma_1^2$ with a common shape parameter of 1. The scale parameters are defined by setting the modes of the priors close to maximum likelihood estimates: \mathcal{V}_{OLS}^2 for \mathcal{V}_ϵ^2 (where \mathcal{V}_{OLS}^2 is the residual variance from an AR(1) model estimated by OLS), at $0.01\mathcal{V}_{OLS}^2$ for σ_0^2 and at 0.001^2 for σ_1^2 . The modes of σ_0^2 and σ_1^2 are centered at the maximum likelihood estimates to ensure that the bimodality in the posterior distribution of the log of σ_1^2 documented in Figure 2 is not induced by the prior.

For given parameters, the likelihood is easily computed via the Kalman filter. It is therefore simple to find the posterior mode, at which the chain is initialized. Posterior mode values suggest that time variation is nearly absent.

Starting with Gibbs sampling, we draw 40 000 times after a burn-in of 5000. The recursive parameter means seem to settle down (not reported) and the posterior distributions are in line with a normal approximation taken at the mode, suggesting a persistent AR(1) with little sign of parameter variation (see Figure 1). It may therefore seem reasonable to assume that the chain has produced a sample representative of the entire posterior.

However, the AIMH scheme tells a different story. The proposal is initialized at a mixture of two normals $g_0(z) = 0.5\mathcal{A}(z; \tau; \hat{\mathcal{S}}) + 0.5\mathcal{A}(z; \tau; 16\hat{\mathcal{S}})$, where τ is the posterior mode and $-\hat{\mathcal{S}}$ is the inverse of the Hessian of the log-posterior evaluated at τ . The AIMH soon discovers that the posterior distribution of $\log(\sigma_1^2)$; not to mention σ_1^2 , is highly non-normal (see Figure 2), with substantial probability mass around a second mode corresponding to non-trivial amounts of time variation in \mathcal{V}_t and a lower \mathcal{V}_1 :

We also ran the adaptive random walk Metropolis sampler outlined at the start of the

section. The sampler settles down to an acceptance rate of 20% and obtains the correct posterior distribution, and in particular finds both modes. Table 1 gives the inefficiency factors for all three samplers as well as the inefficiency factors of the Gibbs and ARWM relative to AIMH. The table shows that the AIMH sampler is appreciably more efficient than the other two samplers.

5.2 Additive semiparametric Gaussian models

In this example we consider the additive semiparametric regression model with Gaussian errors, with some of the covariates entering linearly and the others entering more flexibly

$$y_i = \sum_{j=1}^m \beta_j z_{ji} + \sum_{h=1}^H f_h(x_{h,i}) + \mathcal{U}_\epsilon z_i; \quad (9)$$

the z_i are $nid(0;1)$ and z is a vector of regressors that enter linearly. The $x_h; h = 1; \dots; H$ are covariates that enter more flexibly by using the quadratic polynomial spline functions

$$f_h(x_{h,i}) = \beta_{0,h} x_{h,i} + \sum_{j=1}^J \beta_{h,j} (x_{h,i} - x_{h,j})_+^2 = \beta_{0,h} x_i + g_h(x_{h,i}); \quad (10)$$

where $x_+ = x$ if $x > 0$ and 0 otherwise and $\{\tilde{x}_{h,1}; \dots; \tilde{x}_{h,k}\}$ are points (or 'knots') on the abscissae of x_h such that $\min(x_h) = \tilde{x}_{h,1} < \dots < \tilde{x}_{h,J} < \max(x_h)$. In this paper we choose 30 knots so that each interval contains the same number of observed values of x_h . For a discussion of quadratic spline bases and other related bases see chapter 3 of Ruppert et al. (2003). We assume that a global intercept term is included in z in (9) and for simplicity we include the parameters $\beta_{h,0}; h = 1; \dots; H$ in the vector β and $x_h; h = 1; \dots; H$ as part of the vector z . This transforms the nonparametric model into an highly parametrized linear model

$$y = \tilde{Z}\beta + \mathcal{U}_\epsilon; \quad (11)$$

The prior for the linear parameters β is normal with a diagonal covariance matrix $\beta \sim N(0; \nu_\gamma^2 I)$, where ν_γ can be set to a large number. It is also convenient to assume a normal prior for the nonparametric part, with all parameters independent and $\beta_{h,j} \sim N(0; \zeta_h^2); j =$

$1; \dots; J; h = 1; \dots; H$: However, with this prior there is a high risk of over-fitting if we simply set ζ_h^2 to a large number. The variance ζ_h^2 is often chosen by cross-validation, but in a fully Bayesian setting we can treat ζ_h^2 as a parameter. To illustrate the advantage of AIMH in working with different priors, we experiment with two options for the prior ζ_h^2 . The first prior is log-normal and rather dispersed: $\ln(\zeta_h^2) \sim N(0; 5^2)$, the second is inverse gamma with shape parameter 1 and scale parameter implied by setting the mode at 0.1^2 . The prior for \mathcal{V}_ϵ^2 is inverse gamma with shape parameter one and scale parameter implied by setting the prior mode at the OLS residual variance estimated on (11). The prior for $\tau = (\tau_0; \tau_1; \dots; \tau_H)$ is therefore jointly normal conditional on $\zeta^2 = \{\zeta_1^2; \dots; \zeta_H^2\}$, $\tau | \zeta \sim N(\mathbf{0}; V_\tau(\zeta))$, where $V_\tau(\zeta) = \text{diag}(v_\tau^2 I; \zeta_1^2 I; \dots; \zeta_H^2 I)$ is a block diagonal matrix. One way to estimate the posterior density of the semiparametric model is to use Gibbs or Metropolis-within-Gibbs sampling as proposed by Wong and Kohn (1996). In this approach the parameters $\tau = \{\tau_0; \tau_1; \dots; \tau_H\}$ are conjugate given $\mu = \{\mathcal{V}_\epsilon^2; \zeta_1^2; \dots; \zeta_H^2\}$; and \mathcal{V}_ϵ^2 is conjugate given τ : Each variance ζ_h^2 can be updated with a Gibbs step for the inverse gamma prior, or with a Metropolis-Hastings step for the log-normal prior. In this second case, we use a Laplace approximation of $p(\ln(\zeta_h^2) | \tau_h)$, which is very fast to compute using analytical derivatives. However, the correlation between ζ_h^2 and $\{\tau_{h,1}; \dots; \tau_{h,J}\}$ could be quite high using either prior for ζ_h^2 . In addition, using a log normal prior for ζ_h^2 leads to high rejection rates in the Metropolis-Hastings step when generating the ζ_h^2 in the Boston housing data example of section 5.2.1. Both problems are elegantly solved by integrating out τ and generating μ as a block using an efficient AIMH sampler.

The next example shows how to update all parameters in one block with an efficient AIMH sampler. We first note that, conditional on μ , τ can be integrated out, making it possible to compute $p(\mu | y) \propto p(y | \mu) p(\mu)$; where $y | \mu \sim N(\mathbf{0}; \mathcal{V}_\epsilon^2 I + \tilde{Z}' V_\tau(\zeta) \tilde{Z})$.

5.2.1 Application: Boston housing data

We use the Gaussian semiparametric model to study the Boston housing data introduced by Harrison and Rubinfeld (1978) and analyzed semiparametrically by Smith and Kohn (1996). The dataset is available at www.cs.utoronto.ca/~delve/data/boston and has 506 observations. The dependent variable is the log of MV , the median value of owner-occupied

homes. We use all 13 available covariates (see Smith and Kohn or the web-site for a full description) in the linear part and the following six in the nonparametric part (Smith and Kohn use only the first five): $X_5 = NOX$, nitrogen oxide concentration, $X_6 = RM$; average number of rooms, $X_8 = DIS$; logarithm of the distance from five employment centers, $X_{10} = TAX$; property tax rate, $X_{13} = STAT$; proportion of the population that is lower status, $X_1 = CRIM$; per capita crime rate by town.

The proposal distribution for the seven parameters $\mu = \{\ln(\frac{\gamma_e^2}{\epsilon^2}); \ln(\zeta_5^2); \dots; \ln(\zeta_1^2)\}$ is initialized by fattening the tails of the Laplace approximation. To find the Laplace approximation, we simply apply Newton-Raphson optimization (with numerical derivatives) to $\ln p(y|\mu) + \ln p(\mu)$; which involves no extra coding effort since both densities are needed to compute the MH acceptance ratio. Figure 3 provides results for the case of a log-normal prior on ζ_h^2 ; $h = 1; \dots; H$ and shows that the acceptance rate quickly improves and stabilizes at around 60% when all seven parameters are updated jointly. Most parameters are approximately lognormally distributed, except those connected to the variables TAX and $CRIM$, which benefit from the added flexibility of mixtures. The correlation matrix of the smoothing parameters $\{\ln(\zeta_5^2); \dots; \ln(\zeta_1^2)\}$ is nearly diagonal. This suggests that the AIMH could handle large numbers of smoothing parameters efficiently by updating them in blocks (with a different proposal density estimated adaptively on each block), since the blocks would be nearly independent of each other.

Table 1 reports the inefficiency factors for both the Gibbs sampler and the AIMH sampler for both inverse gamma and log normal priors, as well as the inefficiency of the Gibbs sampler relative to the AIMH sampler. The table shows that in terms of relative efficiency (defined at the beginning of section 5), the AIMH is about 40% more efficient than the Gibbs sampler when both samplers use the inverse gamma prior on ζ_h^2 , and nearly seven times more efficient when both samplers use the log-normal prior. Reported results are for the average inefficiency factors (over both h and i) of $f_h(x_{h,i})$. Looking at the autocorrelation of the log-parameters gives similar inefficiency ratios.

We also applied the adaptive random walk Metropolis sampler to this data set, but could not make it work well. With the sampler initialized at the posterior mode, the acceptance rate started at over 50%, but within a few hundred iterations fell to below 1% and stayed

there indefinitely. We do not report any inefficiency factors for this sampler because we do not believe that inference is reliable with such a low acceptance rate. We conjecture that the poor behavior of the ARWM sampler in this example compared to the other two examples is because this example has 7 parameters whereas the other two have 3 and 2 parameters. In addition, the second derivatives of the log posterior in this example are far from constant, so a unique covariance matrix may do very poorly. By contrast, a mixture of normals allows for local correlations between the parameter and therefore may be less affected.

Boston	mean $\hat{f}_h(x_{h,i})$		Inflation	$\log(\frac{2}{\epsilon})$	$\log(\frac{2}{\epsilon_0})$	$\log(\frac{2}{\epsilon_1})$
AIMH, IG	2.6		AIMH	6.7	2.8	6.1
Gibbs, IG	6.3 (1.4)		Gibbs	9.4 (1.3)	113.3 (37.4)	156.4 (23.7)
AIMH, LN	1.6		ARWM	21.5 (3.1)	23.5 (8.3)	23.6 (3.8)
M-Gibbs, LN	18.4 (6.8)					

Table 1: Inefficiency factors for the semiparametric (Boston) and state space (inflation) models, together with the inefficiencies of the Gibbs sampler and the ARWM relative to the AIMH sampler in brackets. AIMH: adaptive independent Metropolis-Hastings; M-Gibbs: Metropolis-within-Gibbs; and ARWM: adaptive random walk Metropolis. IG and LN: inverse gamma and log-normal priors for the Boston data.

6 Discussion and conclusion

In order to understand the strengths and limitations of our sampler, we find it useful to consider two desirable qualities of an adaptive IMH scheme. First, given a sufficiently large sample drawn from the target, we wish to construct a proposal density which fits the target as well as possible. This is an approximating ability: we want to accurately 'map' the areas that we have already explored. Second, we wish the sampler to perform as well as possible when the initial proposal fails to cover part of the support of the target distribution. This is an exploring ability: when we propose in a region where our map is poor, we want to explore that region and quickly update our map.

For example, using a normal proposal when the target is highly non-normal results in little approximating ability. Updating the proposal only once or very rarely results in little exploring ability, since the proposal reacts slowly or not at all to the information that it is fitting poorly at a given point.

Our sampler has several characteristics designed to enhance its exploring ability. Frequent updating, particularly at early iterations, and updating following a sequence of low MH acceptance probabilities quicken the pace at which the proposal adapts to the information that it is not fitting well in a certain area. Long tails are useful not only to satisfy (3) and (4), but also to improve the chances of venturing into unexplored parts of the state space. Finally, mixtures are ideally suited for this exploration because a new component can be centered on a value causing a sequence of rejections. The long runs of rejections that can plague standard IMH are therefore much less likely using our AIMH sampling scheme because the proposal distribution is updated frequently and will accommodate the cluster of rejections by changing the mixture parameters or adding a new component. If our sampler makes a move in a region where the proposal fits poorly, it should therefore be able to explore it. Of course as the parameter dimension increases, if the initial proposal fails to cover a region we may never explore that region simply because the probability of making a proposal there is too small.

The most interesting applications arise when current best practice is inefficient or cumbersome and, in our opinion, when adaptation starts early. In work in progress we have worked successfully with 15 to 25 parameters in several real data applications, but in general, the number of parameters that it is possible to handle using our approach depends on the shape of the target distribution.

Acknowledgement

We would like to thank Luke Tierney, Christophe Andrieu, Antonietta Mira and two anonymous referees for helpful suggestions and questions that helped improve the accuracy and presentation of a previous version of the paper. Robert Kohn's research was partially supported by ARC grant DP0667069.

References

Andrieu, C. and Moulines, D. (2006), "On the ergodicity properties of some adaptive MCMC algorithms," *Annals of Applied Probability*, 16, 1462{1505.

- Andrieu, C., Moulines, D., and Doucet, A. (2005), "Stability of stochastic approximation under variable conditions," *SICON*, 44, 283{312.
- Andrieu, C. and Robert, C. P. (2001), "Controlled MCMC for optimal sampling," Technical report, University of Bristol.
- Atchade, Y. and Liu, J. (2007), "The Wang-Landau algorithm for Monte carlo computation in general state spaces," Available at <http://www.mathstat.uottawa.ca/~atch436/gwl.pdf>.
- Atchade, Y. and Rosenthal, J. (2005), "On adaptive Markov chain Monte Carlo algorithms," *Bernoulli*, 11, 815{828.
- Bradly, P. and Fayyad, U. (1998), "Refining initial points for k-means clustering," *Proceedings of the 15th International Conference on Machine Learning*, 91{99.
- Carter, C. and Kohn, R. (1994), "On Gibbs sampling for state-space models," *Biometrika*, 83, 589{601.
- Figueredo, M. and Jain, A. (2002), "Unsupervised learning of finite mixture models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24, 381{396.
- Fruhwirth-Schnatter, S. (2004), "Efficient Bayesian parameter estimation," in *State space and unobserved component models*, eds. Harvey, A., Koopman, S., and Shephard, N., Cambridge: Cambridge University Press, pp. 123{151.
- Gasemyr, J. (2003), "On an adaptive version of the Metropolis-Hastings algorithm with independent proposal distribution," *Scandinavian Journal of Statistics*, 30, 159{173.
- Gilks, W., Roberts, G., and Sahu, S. (1998), "Adaptive Markov chain Monte Carlo through regeneration," *Journal of the American Statistical Association*, 93, 1045{1054.
- Giordani, P. and Kohn, R. (2008), "Efficient Bayesian inference for multiple change-point and mixture innovation models," *Journal of Business and Economic Statistics*, 26, 66{77.
- Haario, H., Saksman, E., and Tamminen, J. (2001), "An adaptive Metropolis algorithm," *Bernoulli*, 7, 223{242.

- Hamerly, G. and Elkan, C. (2002), "Alternatives to the k-means algorithm that find better clusterings," in *Proceedings of the Eleventh International Conference on Information and Knowledge Management*, eds. Kalpakis, K., Goharian, N., and Grossmann, D., New York: Academic Press, pp. 600{607.
- Harrison, D. and Rubinfeld, D. (1978), "Hedonic prices and the demand for clean air," *Journal of Environmental Economics and Management*, 5, 81{102.
- Hastie, D. (2005), "Towards automatic reversible jump Markov chain Monte Carlo," Unpublished PhD dissertation, Department of Mathematics, University of Bristol.
- Hastings, W. (1970), "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, 57, 97{109.
- Holden, L. (1998), "Adaptive chains," Manuscript, Norwegian Computing Center, Oslo.
- Liang, F., Liu, C., and Carroll, R. J. (2007), "Stochastic approximation in Monte Carlo computation," *Journal of the American Statistical Association*, 102, 305{320.
- McLachlan, G. and Peel, D. (2000), *Finite Mixture Models*, New York: Wiley.
- Mengersen, K. L. and Tweedie, R. L. (1996), "Rates of convergence of the Hastings and Metropolis algorithms," *The Annals of Statistics*, 24, 101{21.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., H., T. A., and Teller, E. (1953), "Equation of state calculations by fast computing machines," *Journal of Chemical Physics*, 21, 1087{1092.
- Nott, D. and Kohn, R. (2005), "Adaptive sampling for Bayesian variable selection," *Biometrika*, 92, 747{763.
- Roberts, G. O. and Rosenthal, J. S. (2006), "Examples of adaptive MCMC," Preprint (<http://probability.ca/je@ftplibdir/adaptex.pdf>).
- | (2007), "Coupling and ergodicity of adaptive MCMC," *Journal of Applied Probability*, 44, 458{475.

- Ruppert, D., Wand, M., , and Carroll, R. (2003), *Semiparametric regression*, Cambridge: Cambridge University Press.
- Smith, M. and Kohn, R. (1996), "Nonparametric regression using Bayesian variable selection," *Journal of Econometrics*, 75, 317{343.
- Spall, J. (2003), *Introduction to Stochastic Search and Optimization*, New York: Wiley.
- Ueda, N., Nakano, R., Ghahramani, Z., and Hinton, G. (2000), "SMEM algorithm for mixture models," *Neural Computation*, 12, 2109{2128.
- Verbeek, J., Vlassis, N., and Krose, B. (2003), "Efficient Greedy Learning of Gaussian Mixture Models," *Neural Computation*, 15, 469{485.
- Wang, F. and Landau, D. P. (2001), "Efficient multiple-range random walk algorithm to calculate the density of states," *Physical Review*, 86, 2050{2053.
- Wong, C. and Kohn, R. (1996), "A Bayesian approach to additive semiparametric regression," *Journal of Econometrics*, 74, 209{235.

Appendix 1: k-harmonic means clustering

We estimate the mixture of normal parameters using the k-harmonic means clustering algorithm which can be described as follows (See Hamerly and Elkan 2002, for a discussion).

Let p be the number of clusters.

1. Initialize the algorithm with c_1, \dots, c_p ; the component centers. The starting values are c with the procedure of Bradley and Fayyad (1998) . We depart slightly from Bradley and Fayyad in using the harmonic k-means algorithm (rather than k-means) in the initialization procedure.
2. For each data point μ_t ; compute a weight function $w(\mu_t)$ and a membership function $m(c_i|\mu_t)$ for $t = 1, \dots, n$ as

$$w(\mu_t) = \frac{\sum_{i=1}^p \|\mu_t - c_i\|^{-p-2}}{(\sum_{i=1}^p \|\mu_t - c_i\|^{-p})^2} \quad \text{and} \quad m(c_i|\mu_t) = \frac{\|\mu_t - c_i\|^{-p-2}}{\sum_{i=1}^p \|\mu_t - c_i\|^{-p-2}};$$

where $\|\mu_t - c_i\|$ is the Euclidean or Mahalanobis distance. Following Bradley and Fayyad (1998), we put a lower boundary ϵ on $\|\mu_t - c_i\|$ (to avoid degeneracies when trying $\|c_i - c_i\|$). The membership function softens the sharp membership of the k-means algorithm, so one observation can belong to more than one cluster in differing degrees. The weight function gives more weight to observations that are currently covered poorly (i.e. that are far from the nearest center).

3. Update each center c_i

$$c_i = \frac{\sum_{t=1}^n m(c_i|\mu_t) w(\mu_t) \mu_t}{\sum_{t=1}^n m(c_i|\mu_t) w(\mu_t)}.$$

4. Repeat until convergence. This gives the cluster centers, which we take as estimates of the component means. The other mixture parameters can then be estimated for $i = 1; \dots; k$ as

$$V_i = \frac{\sum_{t=1}^n m(c_i|\mu_t) w(\mu_t) (\mu_t - c_i)(\mu_t - c_i)'}{\sum_{t=1}^n m(c_i|\mu_t) w(\mu_t)} \quad \text{and} \quad \pi_i \propto \sum_{t=1}^n m(c_i|\mu_t) w(\mu_t):$$

5. The number of clusters is chosen with the BIC criterion given a maximum number (5 in our examples).

We notice that the covariance matrices V_i are only estimated once, after convergence. k-means type algorithms also differ from the EM algorithm in that they do not evaluate the likelihood $p(\mu|c_1; \dots; \pi_1, V_1; \dots)$. This sub-optimal use of information in fact turns out to be a great advantage for our purposes. Fewer iterations than for EM are needed for convergence, and each iteration is faster. Even more importantly, the algorithm does not get stuck in the small degenerate clusters caused by rejections in the sense that, unlike for the EM algorithm with freely estimated covariances, these small clusters are not absorbing. If k-harmonic means does find a degenerate cluster, this causes no trouble for convergence, and after convergence we can use a pre-defined matrix in place of any non-positive-definite covariance matrix (for example, if V_i is not positive definite we set it to $0.5^2 \text{Var}(\mu)$). If desired, the mixture parameters can be refined with a few steps of the EM algorithm. In this case, we recommend not updating the the covariance matrices for the reasons just discussed.

Appendix 2: Proofs

The one-step transition kernel for Z_{n+1} in section 2 is given by

$$T_n(z; dz') = \mathbb{Q}_n(z; z')q_n(z') \mathbb{1}\{dz'\} + \mathbb{1}_z(dz')(1 - \mathbb{Q}_n(z)) \quad (12)$$

where $\mathbb{1}_z(dz') = 1$ if $z \in dz'$ and is 0 otherwise, and

$$\mathbb{Q}_n(z) = \int_{\mathcal{Z}} \mathbb{Q}_n(z; z')q_n(z') \mathbb{1}\{dz'\}; \quad (13)$$

By the construction of the MH transition kernel,

$$\int_{\mathcal{Z}} \mathbb{H}(z) T_n(z; dz') \mathbb{1}\{dz\} = \mathbb{H}(z') \mathbb{1}\{dz'\}; \quad (14)$$

In this section K is a generic constant, independent of $n; z$ and z' . It is convenient to write any function of z and z' of the form $h_n(z; z')$ as $h_n(z)$. Without loss of generality we assume throughout this section that \mathcal{Z} is a discrete space. Exactly the same proof goes through for the continuous case with summations replaced by integrals. We use the notation $z_{s:t}$ to mean $\{z_s; \dots; z_t\}$ for $s \leq t$, with a similar interpretation for $Z_{s:t}$.

To prove Theorem 1 we first obtain the following two lemmas.

Lemma 1 *Under the assumptions of Section 2, for any $n; k > 0$ and $z; z' \in \mathcal{Z}$,*

- (a) $q_n(z) \leq K g_0(z)$.
- (b) $\mathbb{Q}_n(z; z')q_n(z') \leq K g_0(z')$
- (c) *There exists an $\alpha_1, 0 < \alpha_1 < 1$, such that $\mathbb{Q}_n(z; z')q_n(z') > \alpha_1 \mathbb{H}(z')$ for all $z; z' \in \mathcal{Z}$.*
- (d) $\mathbb{Q}_n(z) > \alpha_1$ for all $z \in \mathcal{Z}$, where $\mathbb{Q}_n(z)$ is defined by (13).
- (e) *For $k \geq 1$, let $\Phi_n(z; z') = \mathbb{Q}_n(z; z')q_n(z') - \mathbb{Q}_{n+1}(z; z')q_{n+1}(z')$. Then,*

$$|\Phi_n(z; z')| \leq K \left(g_0(z') + \frac{\mathbb{H}(z')}{\mathbb{H}(z)} g_0(z) \right) a_n; \quad (15)$$

(f)

$$|\mathbb{Q}_n(z) - \mathbb{Q}_{n+1}(z)| \leq K \left(1 + \frac{g_0(z)}{\mathbb{H}(z)} \right) a_n \quad (16)$$

Proof. (a) $q_n(z) = g_0(z) = \lambda_1 + (1 - \lambda_1)g_n(z) = g_0(z)$

and the result follows from (e). ■

With z_1 as in Lemma 1, choose $0 < \alpha < z_1$ and let

$$R_n(z; z') = \frac{T_n(z; z') - \alpha \mathcal{H}(z')}{1 - \alpha} \quad (17)$$

Then, $R_n(z; z')$ is a one-step transition kernel with the following properties.

Lemma 2 (a)

$$\sum_z \mathcal{H}(z) R_n(z; z') = \mathcal{H}(z') :$$

(b)

$$R_n(z; z') \leq K g_0(z_n) + \alpha \mathcal{H}(z')$$

where $0 < \alpha < 1$.

(c)

$$|R_n(z; z') - R_{n+1}(z; z')| \leq K a_n \left\{ \left(g_0(z') + \frac{\mathcal{H}(z')}{\mathcal{H}(z)} g_0(z) \right) + \left(1 + \frac{g_0(z)}{\mathcal{H}(z)} \right) \mathcal{H}(z') \right\}$$

(d)

$$\sum_{z_{n-m+1}} \cdots \sum_{z_{n-1}} \prod_{k=1}^m R_{n-k}(z_{n-k}; z_{n-k+1}) \leq K g_0(z_n) + \alpha^m \mathcal{H}_{z_{n-m}}(z_n)$$

(e) For $1 \leq l \leq j-1$ and $j = 1; \dots; n$,

$$\sum_{z_{n-l-1}} \cdots \sum_{z_{n-j+1}} \mathcal{H}(z_{n-j+1}) \prod_{k=1+1}^{j-1} R_{n-j}(z_{n-k}; z_{n-k+1}) = \mathcal{H}(z_{n-l})$$

(f) For $j = 1; \dots; n$ and $l = 1; \dots; j-1$,

$$\begin{aligned} & \left| \sum_{z_{n-j+1}} \cdots \sum_{z_{n-1}} \mathcal{H}(z_{n-j+1}) \prod_{k=l+1}^{j-1} R_{n-j}(z_{n-k}; z_{n-k+1}) \right. \\ & \left. \times \left(R_{n-l}(z_{n-l}; z_{n-l+1}) - R_{n-j}(z_{n-l}; z_{n-l+1}) \right) \prod_{k=1}^{l-1} R_{n-k}(z_{n-k}; z_{n-k+1}) \right| \leq K g_0(z_n) a_{n-j} (j-l)^{-1} \end{aligned}$$

Proof. (a) follows from (17) and (14). (b) follows from (17). (c) follows from (15) and (16). (d) is true for $m = 1$ and is obtained in general by induction. (e) follows from part (a). (f) follows from parts (a) to (e). ■

Proof of Theorem 1. Let $\Phi_j; j = 1; 2; \dots$ be an independent Bernoulli process such that $\Phi_j = 1$ with probability λ and $\Phi_j = 0$ with probability $1 - \lambda$. From (17), $T_n(z; z') = (1 - \lambda)R_n(z; z') + \lambda\mathcal{K}(z')$ so that we can interpret $T_n(z; z')$ as a mixture of transition kernels, such that $T_n(z; z') = R_n(z; z')$ if $\Phi_n = 0$ and $T_n(z; z') = \mathcal{K}(z')$ if $\Phi_n = 1$. For $j = 1; \dots; n$, let $A_{n,j}$ be the event that $\Phi_{n-j+1} = 1; \Phi_k = 0; k = n - j + 2; \dots; n$. Let B_n be the event that Φ_j

We also have that

$$\begin{aligned} \Pr(Z_n = z_n \mid B_n) &= \sum_{z_0} \cdots \sum_{z_{n-1}} q_0(z_0) g_0(z_1) \prod_{k=1}^{n-1} R_k(z_k; z_{k+1}) \\ &\leq K g_0(z_n) + n^{-1} g_0(z_n) \leq K g_0(z_n) : \end{aligned}$$

using Lemma 2 (c) and (3). Hence,

$$| \Pr(Z_n = z_n) - \mathbb{1}(z_n) | \leq K g_0(z_n) \left((1 - \lambda)^n + n^{-r_1} \right) \quad (18)$$

The proof of Theorem 1 follows. ■

The proof of Theorem 2 is similar to that in Nott and Kohn (2005) if we use (18).

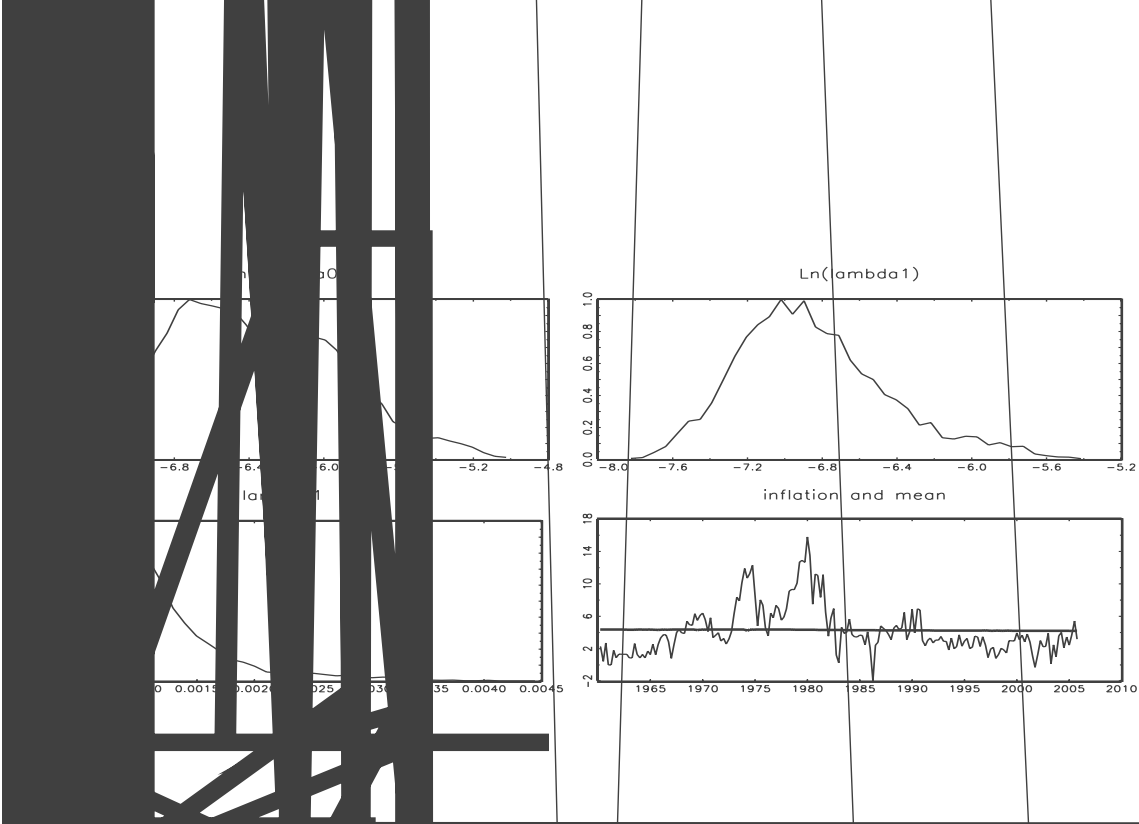


Figure 1: Inference for a time varying parameter AR(1) model for US inflation by Gibbs sampling. (a) marginal distribution of $\ln(\lambda_0)$ (b) marginal distribution of $\ln(\lambda_1)$ (c) marginal distribution λ_1 (d) inflation plot and mean, estimated as $E[(c_t=1 - \lambda_t)|y]$ (e) marginal distribution of $\lambda_0|y$ (f) $E(c_t|y)$ (bold line) and $E(\lambda_t|y)$:




Figure 2: Inference for the model of Figure 1 by adaptive IMH. The interpretation of the panels is the same as in Figure 1.

Figure 3: Inference for semiparametric model of housing prices by adaptive IMH. First row: recursive acceptance rate for the last $\min(it, 500)$ iterations, recursive means of $\ln(\zeta_i)$; marginal of \mathcal{X}_e . Second and third rows: marginals of $\ln(\zeta_i)$. Fourth and fifth rows: means of $\bar{X} + g_i(X)$.