

Informed reversible jump algorithms^{*†}

Philippe Gagnon[‡]

Department of Mathematics and Statistics, Université de Montréal, C.P. 6128, Succursale Centre-ville, Montreal, QC, H3C 3J7, Canada. e-mail: philippe.gagnon.3@umontreal.ca

Abstract: Incorporating information about the target distribution in proposal mechanisms generally produces efficient Markov chain Monte Carlo algorithms (or at least, algorithms that are more efficient than uninformed counterparts). For instance, it has proved successful to incorporate gradient information in fixed-dimensional algorithms, as seen with algorithms such as Hamiltonian Monte Carlo. In trans-dimensional algorithms, Green (2003) recommended to sample the parameter proposals during model switches from normal distributions with informative means and covariance matrices. These proposal distributions can be viewed as asymptotic approximations to the parameter distributions, where the limit is with regard to the sample size. Models are typically proposed using uninformed uniform distributions. In this paper, we build on the approach of Zanella (2020) for discrete spaces to incorporate information about neighbouring models. We rely on approximations to posterior model probabilities that are asymptotically exact. We prove that, in some scenarios, the samplers combining this approach with that of Green (2003) behave like ideal ones that use the exact model probabilities and sample from the correct parameter distributions, in the large-sample regime. We show that the implementation of the proposed samplers is straightforward in some cases. The methodology is applied to a real-data example. The code is available online.[§]

Keywords and phrases: Bayesian statistics, large-sample asymptotics, Markov chain Monte Carlo methods, model averaging, model selection, trans-dimensional samplers, variable selection, weak convergence.

Received April 2020.

Contents

1	Introduction	3952
1.1	Reversible jump algorithms	3952
1.2	Problems with RJ	3954
1.3	Objectives	3955

^{*}The author acknowledges support from NSERC (Natural Sciences and Engineering Research Council) of Canada and FRQNT (Le Fonds de recherche du Québec – Nature et technologies).

[†]The author corrected the first two sentences of the abstract. The abstract was corrected on 27 August 2021.

[‡]The author thanks an associate editor and two anonymous referees for helpful suggestions that led to an improved manuscript. Also, he acknowledges enlightening discussions with Sebastian M Schmon who introduced the author to weak convergence of Markov chains in Bayesian large-sample regimes.

[§]See ancillary files on [arXiv:1911.02089](https://arxiv.org/abs/1911.02089).

1.4	Organization of the paper	3956
2	Input design and proposed RJ	3957
2.1	Specification of the model proposal distributions	3957
2.2	Specification of the parameter proposal distributions	3959
2.3	Resulting RJ	3961
2.4	Methods to use when the approximations are not accurate	3963
3	Asymptotic result and analysis	3964
3.1	Asymptotic result	3964
3.2	Analysis of the limiting RJ	3967
4	Application: Variable selection in wholly-robust linear regression	3968
4.1	Wholly-robust linear regression	3968
4.2	Results and analysis	3970
5	Discussion	3972
A	Improving the approximations	3973
A.1	RJ with the method of Karagiannis and Andrieu (2013)	3973
A.2	RJ with additionally the method of Andrieu et al. (2018)	3975
A.3	Improving the model proposal distribution	3975
B	Details of the variable-selection example	3979
B.1	Linear regression	3979
B.2	Normal linear regression	3980
B.3	Robust linear regression	3981
C	Proofs	3983
	References	3993

1. Introduction

1.1. Reversible jump algorithms

Reversible jump (RJ, Green (1995)) algorithms are Markov chain Monte Carlo (MCMC) methods that one uses to sample from a target distribution $\pi(\cdot, \cdot \mid \mathbf{D}_n)$ defined on a union of sets $\bigcup_{k \in \mathcal{K}} \{k\} \times \mathbb{R}^{d_k}$, \mathcal{K} being a countable set and d_k positive integers. This distribution corresponds in Bayesian statistics to a joint posterior distribution of a model indicator $K \in \mathcal{K}$ and the parameters of Model K , denoted by $\mathbf{X}_K \in \mathbb{R}^{d_K}$, given \mathbf{D}_n , a data sample of size n . Such a posterior distribution allows to jointly infer about (K, \mathbf{X}_K) , or in other words, to simultaneously achieve model selection/averaging (Hoeting et al., 1999) and parameter estimation. In the following, we assume for simplicity that the parameters of all models are continuous random variables. Again for simplicity, we will abuse notation by also using $\pi(\cdot, \cdot \mid \mathbf{D}_n)$ to denote the joint posterior density with respect to a product of the counting and Lebesgue measures. We will use $\pi(\cdot \mid \mathbf{D}_n)$ and $\pi(\cdot \mid k, \mathbf{D}_n)$ to denote the marginal posterior probability mass function (PMF) of K and conditional posterior distribution/density of \mathbf{X}_K given that $K = k$, respectively. Note that in this paper, we assume that \mathcal{K} is defined such that $\pi(k \mid \mathbf{D}_n) > 0$ for all k .

The variable $K \in \mathcal{K}$ can take different forms in practice. In mixture modelling (Richardson and Green, 1997), K is the number of components and \mathcal{K} is a subset

of the positive integers and defines a sequence of nested models, i.e. Model 1 is nested in Model 2 which is nested in Model 3, and so on. Often there is no such “ordering” between the models. This is for instance the case in variable selection. In this framework, K is a vector of 0’s and 1’s indicating which covariates are included in a model, and for instance, Model $k_0 = (1, 1, 1, 0, 0, 1, 0, \dots, 0)$ is the model with covariates 1, 2, 3 and 6 (as in Figure 1 below); \mathcal{K} is thus a collection of such vectors. A RJ algorithm explores both the model space and the parameter spaces. We focus in this paper on the case where \mathcal{K} *does not* define a sequence of nested models. When \mathcal{K} *defines* a sequence of nested models, K can be recoded as an ordinal discrete variable as above and this case has been recently studied by Gagnon and Doucet (2021) who proposed efficient non-reversible trans-dimensional samplers.

The proposal mechanism in a RJ algorithm can be outlined as follows: given a current state of the Markov chain (k, \mathbf{x}_k) , the algorithm first proposes a model, say Model k' , by using a PMF $g(k, \cdot)$ and then parameter values for this model by applying a diffeomorphism $\mathcal{D}_{k \mapsto k'}$ to \mathbf{x}_k and auxiliary variables $\mathbf{u}_{k \mapsto k'} \sim q_{k \mapsto k'}$, yielding the parameter proposal $\mathbf{y}_{k'}$ and another vector of auxiliary variables $\mathbf{u}_{k' \mapsto k}$. The proposal is accepted, meaning that the next state of the Markov chain is $(k', \mathbf{y}_{k'})$, with probability (assuming that the current state has positive density under the target):

$$\alpha_{\text{RJ}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'})) := 1 \wedge \frac{\pi(k', \mathbf{y}_{k'} \mid \mathbf{D}_n) g(k', k) q_{k' \mapsto k}(\mathbf{u}_{k' \mapsto k})}{\pi(k, \mathbf{x}_k \mid \mathbf{D}_n) g(k, k') q_{k \mapsto k'}(\mathbf{u}_{k \mapsto k'}) |J_{\mathcal{D}_{k \mapsto k'}}(\mathbf{x}_k, \mathbf{u}_{k \mapsto k'})|^{-1}}, \quad (1)$$

where $x \wedge y := \min(x, y)$ and $|J_{\mathcal{D}_{k \mapsto k'}}(\mathbf{x}_k, \mathbf{u}_{k \mapsto k'})|$ is the absolute value of the determinant of the Jacobian matrix of the function $\mathcal{D}_{k \mapsto k'}$. If the proposal is rejected, the chain remains at the same state (k, \mathbf{x}_k) . Recall that a diffeomorphism is a differentiable map having a differentiable inverse and note that we abused notation by using $q_{k \mapsto k'}$ to denote both the distribution and the probability density function (PDF) of $\mathbf{u}_{k \mapsto k'}$. The notation $k \mapsto k'$ in subscript is used to highlight a dependence on the model transition that is proposed, which is from Model k to Model k' . In the particular case where $k' = k$, we say that a *parameter update* is proposed; otherwise, we say that a *model switch* is proposed. Model k' is reachable from Model k if it belongs to the support of $g(k, \cdot)$ which is considered in this paper to be the neighbourhood of Model k , denoted by $\mathbf{N}(k)$.

In trans-dimensional samplers, the neighbourhoods are usually formed of the “closest” models. The notion of “proximity” is often natural, making these closest models straightforward to identify. For instance, in Section 4 we apply the proposed methodology in a variable-selection example and the closest models are those with an additional and one less covariates, to which the current model is added. In this paper, we consider that Model k belongs to the support of $g(k, \cdot)$ for all k ; this will be seen to be useful when the posterior model PMF concentrates.

Algorithm 1 presents a general RJ.

Algorithm 1 RJ

1. Sample $k' \sim g(k, \cdot)$, $\mathbf{u}_{k \rightarrow k'} \sim q_{k \rightarrow k'}$ and $u \sim \mathcal{U}[0, 1]$.
2. Apply $\mathcal{D}_{k \rightarrow k'}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) = (\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k})$.
3. If $u \leq \alpha_{\text{RJ}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'}))$, set the next state of the chain to $(k', \mathbf{y}_{k'})$. Otherwise, set it to (k, \mathbf{x}_k) .
4. Go to Step 1.

Looping over the steps described in Algorithm 1 produces Markov chains that are reversible with respect to the target distribution $\pi(\cdot, \cdot \mid \mathbf{D}_n)$. If the chains are in addition irreducible and aperiodic, then they are ergodic (Tierney, 1994), which guarantees, among others, that the law of large numbers holds, implying that ergodic averages are approximations to expectations under $\pi(\cdot, \cdot \mid \mathbf{D}_n)$.

1.2. Problems with RJ

Algorithm 1 takes as inputs: an initial state (k, \mathbf{x}_k) , a total number of iterations, a PMF $g(k, \cdot)$ for each Model k , and functions $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ for each pair (k, k') with Model k' being reachable from Model k . Implementing RJ is well known for being a difficult task considering the large number of functions that need to be specified and the often lack of intuition about how one should achieve their specification, the latter being especially true for functions involved in model switches. Significant amount of work has been carried out to address the specification of the functions $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ involved in model switches in a principled and informed way; see, e.g., Green (2003) and Brooks, Giudici and Roberts (2003). The approaches of these authors are arguably the most popular. Their objective is the following: given $\mathbf{x}_k \sim \pi(\cdot \mid k, \mathbf{D}_n)$, we want to identify $q_{k \rightarrow k'}$, $q_{k' \rightarrow k}$ and $\mathcal{D}_{k \rightarrow k'}$ such that applying the transformation $\mathcal{D}_{k \rightarrow k'}$ to $(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) \sim \pi(\cdot \mid k, \mathbf{D}_n) \otimes q_{k \rightarrow k'}$ leads to $(\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k}) \sim \pi(\cdot \mid k', \mathbf{D}_n) \otimes q_{k' \rightarrow k}$ (at least approximatively). They essentially look for a way to sample from the conditional distributions $\pi(\cdot \mid k', \mathbf{D}_n)$, in this constrained framework. This in turn aims at increasing the acceptance probability α_{RJ} defined in (1) towards

$$\alpha_{\text{MH}}(k, k') := 1 \wedge \frac{\pi(k' \mid \mathbf{D}_n) g(k', k)}{\pi(k \mid \mathbf{D}_n) g(k, k')}, \quad (2)$$

which corresponds to the acceptance probability in a Metropolis–Hastings (MH, Metropolis et al. (1953) and Hastings (1970)) algorithm targeting the PMF $\pi(\cdot \mid \mathbf{D}_n)$. The approach of Green (2003), for instance, proceeds as if the conditional distributions $\pi(\cdot \mid k', \mathbf{D}_n)$ were normal.

Notwithstanding the merit of this objective, it is to be noticed that, even when it is achieved, “half” of the work for maximizing α_{RJ} is done because, as shown in Figure 1, poor models may be proposed more often than they should be if g is not well designed, leading to smaller acceptance rates. Note that when considering candidate proposal distributions $g(k, \cdot)$ having all the same support $\mathbf{N}(k)$, choosing the one that maximizes the acceptance probability α_{RJ}

represents a first step towards proving that it is the best (or at least one of the best) given that they all allow to reach the same models; this choice is indeed expected to improve mixing.




FIG 1. (a) posterior probabilities (represented by the size of the dots) of models forming the neighbourhood of the model with covariates $\{1, 2, 3, 6\}$, in a linear-regression analysis of the prostate-cancer data of Stamey et al. (1989); the models with an additional covariate can be found in the upper part, whereas those with one less variable are in the lower part; (b) probabilities of proposing the models forming the neighbourhood (represented by the size of the dots) when the current model is that with covariates $\{1, 2, 3, 6\}$ and g is the uniform over the neighbourhood; (c) probabilities of proposing the models forming the neighbourhood (represented by the size of the dots) when the current model is that with covariates $\{1, 2, 3, 6\}$ and g is instead an informed version (as defined in Section 1.3)

1.3. Objectives

The specification of g has been overlooked; this PMF is indeed typically set to a uniform distribution as in Figure 1 (b). The first objective of this paper is to propose methodology to incorporate information about neighbouring models in its design to make fully-informed RJ available. To achieve this we rely on a generic technique recently introduced by Zanella (2020) that is used to construct informed samplers for discrete state-spaces. Let us assume for a moment that we have access to the unnormalized posterior model probabilities and that we use a MH algorithm to sample from $\pi(\cdot | \mathbf{D}_n)$; the technique consists in constructing what the author calls *locally-balanced* proposal distributions:

$$g(k, k') \propto h\left(\frac{\pi(k' | \mathbf{D}_n)}{\pi(k | \mathbf{D}_n)}\right) \mathbb{1}(k' \in \mathbf{N}(k)), \tag{3}$$

where h is a continuous positive function such that $h(0) = 0$ and $h(x)/h(1/x) = x$ for all positive x (the square root satisfies this condition for instance), and $\mathbb{1}$ is the indicator function. Such a function h leads to an acceptance probability in the MH sampler given by $\alpha_{\text{MH}}(k, k') = 1 \wedge c(k)/c(k')$, where $c(k)$ is the normalizing constant of $g(k, \cdot)$. A motivation for using this technique is that, in the limit, when the ambient space becomes larger and larger (but the neighbourhoods have a fixed size), there is no need for an accept-reject step anymore; the proposal distributions leave the distribution $\pi(\cdot | \mathbf{D}_n)$ invariant.

Zanella (2020) indeed proves that $c(k)/c(k') \rightarrow 1$ for all possible pairs (k, k') under some conditions. These properties suggest that locally-balanced samplers are efficient, at least in high dimensions. The author in fact empirically shows that they perform better than alternative solutions to sample from PMFs in some practical examples, with a highly marked difference in the examples with high-dimensional spaces.

The first obstacle to achieving our first objective is that we typically do not have direct access to the model probabilities, because they involve integrals over the parameter spaces. Drawing inspiration from the approach of Green (2003) that can be viewed as using asymptotic approximations to $\pi(\cdot | k', \mathbf{D}_n)$ to design RJ, where the limit is with regard to the sample size, we propose to use approximations to the unnormalized version of the model probabilities $\pi(k | \mathbf{D}_n)$ whose accuracy increases with n . We prove that, in some scenarios, RJ using both approximations behave asymptotically as RJ that set $g(k, k') \propto h\left(\frac{\pi(k'|\mathbf{D}_n)}{\pi(k|\mathbf{D}_n)}\right) \mathbb{1}(k' \in \mathbf{N}(k))$ and that sample parameters from the correct conditional distributions $\pi(\cdot | k', \mathbf{D}_n)$. These ideal RJ have acceptance probabilities equal to α_{MH} . All this suggests that the resulting RJ with asymptotically locally-balanced proposal distributions are efficient, at least when n is large enough. We show that it is the case even in a moderate-size real-data example. In this example, the PMF $\pi(\cdot | \mathbf{D}_n)$ concentrates on several models. We also provide evidences that the proposed RJ are efficient when n is large and the PMF $\pi(\cdot | \mathbf{D}_n)$ is highly concentrated so that the mass is non-negligible only for a handful of models.

The approximations on which the proposed methodology is based are more accurate when all parameters of all models take values on the real line. We thus recommend to apply a transformation to parameters for which this is not the case. Such a practice is often beneficial for MCMC methods. It is in fact required when using Hamiltonian Monte Carlo (HMC, see, e.g., Neal (2011)), which is employed to perform parameter updates within RJ in the real-data application of Section 4.

The second objective of this paper is to make clear how each function required for implementation should be specified, at least in some cases, allowing a fully-automated procedure. These cases are those where each of the log conditional densities $\log \pi(\cdot | k, \mathbf{D}_n)$ has a well-defined mode and second derivatives that exist and that are continuous. The procedure can be executed even if the model space is large or infinite, as long as the posterior model PMF concentrates on a reasonable number of models, which is expected in practice.

1.4. Organization of the paper

In Section 2, we address the specification of the inputs required for the implementation of the proposed samplers; in particular, we discuss the design of the PMFs $g(k, \cdot)$, so that both objectives of the paper are achieved. We next present in Section 3 a theoretical result about the asymptotic behaviour of the proposed RJ, and an analysis of the limiting RJ. In Section 4, the methodology is eval-

uated in a real-data variable-selection example. The paper finishes in Section 5 with retrospective comments and possible directions for future research.

The approximations on which the design of the functions $g(k, \cdot)$, $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ rely can be inaccurate when the sample size is not sufficiently large. There exist methods that allow to build upon the approximations used to design $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ to improve the parameter-proposal mechanism. The methods of Karagiannis and Andrieu (2013) and Andrieu et al. (2018) are examples of such methods. An overview of these methods is provided in Section 2, while the details are presented in Appendix A. In Section 2, it is explained that these methods are also useful for improving the model-proposal mechanism. In fact, as the precision parameters of these methods increase without bounds, the resulting samplers converge towards ideal ones that have access to the unnormalized model probabilities to design g and that sample parameters from the correct conditional distributions $\pi(\cdot | k', \mathbf{D}_n)$, for fixed n (see Appendix A for the details). The methods of Karagiannis and Andrieu (2013) and Andrieu et al. (2018) are sophisticated and come at a computational cost (especially that of Karagiannis and Andrieu (2013)), which may not be worth it when n is large enough for the problem at hand. This is the case in our variable-selection example, highlighting that the simple RJ proposed in this paper that use simple approximations are efficient in certain practical situations. The methods of Karagiannis and Andrieu (2013) and Andrieu et al. (2018) are nevertheless useful in the variable-selection example as they allow to show which of the approximations used in the design of $g(k, \cdot)$, $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ explain the gap between the proposed RJ and the ideal ones. Their computational cost is offset by a significant enough improvement in, for instance, the change-point problem presented in Green (1995) (as shown in Karagiannis and Andrieu (2013) and Andrieu et al. (2018)).

Some details of the variable-selection example are presented in Appendix B. All proofs of theoretical results are provided in Appendix C.

2. Input design and proposed RJ

We start in Section 2.1 with the proposed design of g . We next turn in Section 2.2 to the proposed design of the functions $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$, which will be seen to represent a simplified version of that in Green (2003). We present in Section 2.3 the resulting RJ. In Section 2.4, we present an overview of the methods of Karagiannis and Andrieu (2013) and Andrieu et al. (2018) to improve upon the approximations used in the design of the proposal mechanism when these approximations are not accurate enough.

2.1. Specification of the model proposal distributions

The design of g starts with the definition of neighbourhoods around all models which specify the support of $g(k, \cdot)$ for all k . As mentioned in Section 1.1, it is typically possible to define these neighbourhoods in a natural way. It will thus

be considered that the neighbourhoods are given and that we can build on these to address the specification of g .

In practice, $g(k, \cdot)$ is commonly set to the uniform distribution over $\mathbf{N}(k)$: $g(k, k') = 1/|\mathbf{N}(k)|$ for all $k' \in \mathbf{N}(k)$, where $|\mathbf{N}(k)|$ is the cardinality of $\mathbf{N}(k)$. In the simple situation where $\mathbf{u}_{k \rightarrow k'} = \mathbf{y}_{k'} \sim \pi(\cdot | k', \mathbf{D}_n)$ and $|\mathbf{N}(k)| = |\mathbf{N}(k')|$ for all possible pairs (k, k') (the latter is true in variable selection with neighbourhoods defined as in Section 1.1), the RJ acceptance probability reduces to

$$\alpha_{\text{RJ}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'})) = 1 \wedge \frac{\pi(k' | \mathbf{D}_n)}{\pi(k | \mathbf{D}_n)}.$$

It is thus easily seen that a chain may get stuck for several iterations when there are a lot of poor models with $\pi(k' | \mathbf{D}_n) \ll \pi(k | \mathbf{D}_n)$ in the neighbourhood; this is especially true in high dimensions because the number of poor models may be very high. Our goal is to include local neighbourhood information in the PMFs $g(k, \cdot)$ to skew the latter towards high probability models.

We propose here to follow the strategy of Zanella (2020), implying that ideally we would set $g(k, k') \propto h(\pi(k' | \mathbf{D}_n)/\pi(k | \mathbf{D}_n))$ for all $k' \in \mathbf{N}(k)$ (recall (3)), where for all k ,

$$\pi(k | \mathbf{D}_n) = \int \pi(k, \mathbf{x}_k | \mathbf{D}_n) d\mathbf{x}_k.$$

These integrals are typically intractable. We propose to approximate them and for this we assume that each log conditional density $\log \pi(\cdot | k, \mathbf{D}_n)$ has a well-defined mode and that its second derivatives exist and are continuous. In fact, the integrals that we propose to approximate are

$$\int \pi_{\text{un.}}(\mathbf{x}_k | k, \mathbf{D}_n) d\mathbf{x}_k,$$

where $\pi_{\text{un.}}(\cdot | k, \mathbf{D}_n) := \mathcal{L}(\cdot | k, \mathbf{D}_n) \pi(\cdot | k)$ is the unnormalized posterior density under Model k given by the product of the likelihood function and prior parameter density under Model k , and we use that

$$\pi(k | \mathbf{D}_n) \propto \pi(k) \int \pi_{\text{un.}}(\mathbf{x}_k | k, \mathbf{D}_n) d\mathbf{x}_k,$$

$\pi(k)$ being the prior probability assigned to Model k . The approximation that we propose to apply is the Laplace approximation, meaning that we write $\pi_{\text{un.}}(\cdot | k, \mathbf{D}_n) = \exp(\log \pi_{\text{un.}}(\cdot | k, \mathbf{D}_n))$ and $\log \pi_{\text{un.}}(\mathbf{x}_k | k, \mathbf{D}_n)$ is developed using a Taylor series expansion (see, e.g., Davison (1986)). This yields:

$$\hat{\pi}(k | \mathbf{D}_n) \propto \pi(k) (2\pi)^{d_k/2} \pi_{\text{un.}}(\hat{\mathbf{x}}_k | k, \mathbf{D}_n) |\hat{\mathcal{I}}_k|^{-1/2}, \quad \text{for all } k, \quad (4)$$

where $\hat{\mathbf{x}}_k$ is the maximizer of $\pi_{\text{un.}}(\cdot | k, \mathbf{D}_n)$ (and $\pi(\cdot | k, \mathbf{D}_n)$), $\hat{\mathcal{I}}_k$ is minus the matrix of second derivatives of $\log \pi_{\text{un.}}(\cdot | k, \mathbf{D}_n)$ (and $\log \pi(\cdot | k, \mathbf{D}_n)$). We use $\hat{\mathcal{I}}_k$ to denote the matrix of minus the second derivatives because when the prior density of the parameters is uniform, this matrix corresponds to the observed information matrix. It is evaluated at $\hat{\mathbf{x}}_k$, but we make this implicit to simplify

because it will always be the case in this paper. Note that other approximations may be employed. In our framework, we are interested by approximations that are easy to compute (at least in some cases) and asymptotically exact as $n \rightarrow \infty$. This is the case for $\hat{\pi}(k \mid \mathbf{D}_n)$ in (4); it is more precisely a consistent estimator of $\pi(k \mid \mathbf{D}_n)$ under the assumptions mentioned above.

We thus propose to define the model proposal distributions as follows:

$$g(k, k') \propto h(\hat{\pi}(k' \mid \mathbf{D}_n) / \hat{\pi}(k \mid \mathbf{D}_n)),$$

for all $k' \in \mathbf{N}(k)$. The function h needs to be specified. A natural choice (that does not satisfy the conditions mentioned in Section 1.3) is the identity function: $g(k, k') \propto \hat{\pi}(k' \mid \mathbf{D}_n) / \hat{\pi}(k \mid \mathbf{D}_n) \propto \hat{\pi}(k' \mid \mathbf{D}_n)$. The resulting proposal distributions are called *globally-balanced* in Zanella (2020). To understand why, consider the situation where the size of \mathcal{K} is small and it is feasible to switch from any model to any other one, and thus we can set $\mathbf{N}(k) = \mathcal{K}$ for all k . In this case, $g(k, k')$ is exactly equal to $\hat{\pi}(k' \mid \mathbf{D}_n)$, and for large enough n , $\hat{\pi}(k' \mid \mathbf{D}_n) \approx \pi(k' \mid \mathbf{D}_n)$, corresponding to independent Monte Carlo sampling when $\mathbf{u}_{k \rightarrow k'} = \mathbf{y}_{k'} \sim \pi(\cdot \mid k', \mathbf{D}_n)$. Using such *global* proposal distributions $g(k, \cdot)$ with $\mathbf{N}(k) = \mathcal{K}$ thus asymptotically leaves the target distribution invariant without an accept/reject step, hence the name *globally-balanced*. In contrast, when the ambient space becomes larger and larger and the neighbourhoods have a fixed size, i.e. when the locally-balanced proposal distributions asymptotically leave the target distribution invariant, setting h to the identity function instead asymptotically leaves a distribution with $\mathbf{X}_K \mid K \sim \pi(\cdot \mid K, \mathbf{D}_n)$ and $K \sim \pi(\cdot \mid \mathbf{D}_n)^2$ invariant. We recommend to set h to the identity only when it is to be feasible to set $\mathbf{N}(k) = \mathcal{K}$ for all k . Otherwise, it seems better to use locally-balanced proposal distributions; this is the case for instance in Zanella (2020) and our moderate-size (both in n and dimension) variable-selection example.

Two choices of locally-balanced functions h are considered in Zanella (2020): $h(x) = \sqrt{x}$ and $h(x) = x/(1+x)$. The choice $h(x) = x/(1+x)$ yields what the author calls *Barker proposal distributions* because of the connection with Barker (1965)'s acceptance-probability choice:

$$\frac{\hat{\pi}(k' \mid \mathbf{D}_n) / \hat{\pi}(k \mid \mathbf{D}_n)}{1 + \hat{\pi}(k' \mid \mathbf{D}_n) / \hat{\pi}(k \mid \mathbf{D}_n)} = \frac{\hat{\pi}(k' \mid \mathbf{D}_n)}{\hat{\pi}(k' \mid \mathbf{D}_n) + \hat{\pi}(k \mid \mathbf{D}_n)}.$$

The empirical results of Zanella (2020) suggest that this latter choice is superior because h is bounded, which stabilizes the normalizing constants $c(k)$ and $c(k')$ and thus the acceptance probabilities. Livingstone and Zanella (2019) reached the same conclusion using locally-balanced proposal distributions for continuous random variables. In our numerical analyses, both choices lead to similar performances. In practice, a user will choose one; we thus recommend setting $h(x) = x/(1+x)$.

2.2. Specification of the parameter proposal distributions

In this section, we discuss the specification of the functions $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ that are used when $k' \neq k$, i.e. during model-switch attempts. As mentioned,

in Section 1.3, HMC is used during parameter-update attempts. The tuning of HMC free parameters is discussed in Section 2.3.

Green (2003) proposed to set $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ as if the parameters of both Model k and Model k' were normally distributed. More precisely, when $d_{k'} > d_k$ (the other cases are similar), the author proposed to set

$$\mathbf{y}_{k'} = \boldsymbol{\mu}_{k'} + \mathbf{L}_{k'} \begin{pmatrix} \mathbf{L}_k^{-1}(\mathbf{x}_k - \boldsymbol{\mu}_k) \\ \mathbf{u}_{k \rightarrow k'} \end{pmatrix}, \quad (5)$$

where $\boldsymbol{\mu}_{k'}$ and $\boldsymbol{\mu}_k$ are vectors and $\mathbf{L}_{k'}$ and \mathbf{L}_k are matrices. If $\pi(\cdot | k, \mathbf{D}_n)$ and $\pi(\cdot | k', \mathbf{D}_n)$ are normal distributions with means $\boldsymbol{\mu}_k$ and $\boldsymbol{\mu}_{k'}$ and covariance matrices $\mathbf{L}_k \mathbf{L}_k^T$ and $\mathbf{L}_{k'} \mathbf{L}_{k'}^T$, then setting $q_{k \rightarrow k'}$ to a $(d_{k'} - d_k)$ -dimensional standard normal yields $\mathbf{y}_{k'} \sim \pi(\cdot | k', \mathbf{D}_n)$. The function yielding $\mathbf{y}_{k'}$ in (5) defines $\mathcal{D}_{k \rightarrow k'}$. Note that in this case $\mathbf{u}_{k' \rightarrow k}$ is non-existent. Note also that in the following, we will write $\varphi(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ to denote both the distribution and PDF of a normal with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

The rationale behind the approach of Green (2003) is to propose $\mathbf{y}_{k'}$ around a suitable location with a suitable covariance structure using a normal distribution. This approach is asymptotically valid in some cases, by virtue of Bernstein-von Mises theorems (see, e.g., Theorem 10.1 in Van der Vaart (2000) and Kleijn and Van der Vaart (2012)), meaning that for regular models, $\mathbf{X}_k | k, \mathbf{D}_n$ is asymptotically normally distributed as $n \rightarrow \infty$, for all k , when \mathcal{K} and d_k are fixed and do not depend on n . In fact, $\pi(\cdot | k, \mathbf{D}_n)$ concentrates around $\hat{\mathbf{x}}_k$, with $\hat{\mathbf{x}}_k \rightarrow \mathbf{x}_k^*$, at a rate of $1/\sqrt{n}$ with a shape that resembles a normal for large enough n (see Section 2.1 for the definition of $\hat{\mathbf{x}}_k$). When Model k is well specified, \mathbf{x}_k^* is the true parameter value and the covariance matrix is \mathcal{I}_k^{-1} , the latter denoting the inverse information matrix based on the whole data set, evaluated at \mathbf{x}_k^* . We make the dependence on \mathbf{x}_k^* implicit to simplify; the inverse information matrix will always be evaluated at \mathbf{x}_k^* in the following. The information matrix \mathcal{I}_k is equivalent to n times the information of one data point in the independent and identically distributed (IID) setting. In the misspecified case, \mathbf{x}_k^* is the *best possible value* under Model k and the covariance matrix is slightly different than \mathcal{I}_k^{-1} (see Kleijn and Van der Vaart (2012) for the details).

The approach to specify $q_{k \rightarrow k'}$ and $\mathcal{D}_{k \rightarrow k'}$ proposed here is a simpler version of that of Green (2003), but one that is based on the same idea: set

$$\mathbf{y}_{k'} = \mathbf{u}_{k \rightarrow k'},$$

with $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$, for any reachable Model k' , regardless of k (see Section 2.1 for the definition of $\hat{\mathcal{I}}_{k'}$). Consequently, $\mathbf{u}_{k' \rightarrow k} = \mathbf{x}_k$ and $\mathcal{D}_{k \rightarrow k'}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) = (\mathbf{u}_{k' \rightarrow k}, \mathbf{x}_k) = (\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k})$, implying that $q_{k' \rightarrow k} = \varphi(\cdot; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1})$ and $|J_{\mathcal{D}_{k \rightarrow k'}}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'})| = 1$. Thus, if Model k' is a well-specified and regular model, $\mathbf{y}_{k'}$ is asymptotically generated from the correct conditional distribution, i.e. $q_{k \rightarrow k'}$ and $\pi(\cdot | k', \mathbf{D}_n)$ are asymptotically equal, because $\pi(\cdot | k', \mathbf{D}_n)$ is asymptotically equal to $\varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$ (Bernstein-von Mises theorem) and $\hat{\mathcal{I}}_{k'}/n - \mathcal{I}_{k'}/n \rightarrow 0$ in probability, where we say that a matrix converges in probability whenever all entries converge in probability.

It may be the case that $q_{k \rightarrow k'}$ and $\pi(\cdot | k', \mathbf{D}_n)$ are *not* asymptotically equal because Model k' may be misspecified, implying that the limiting covariance of $\pi(\cdot | k', \mathbf{D}_n)$ is not $\hat{\mathcal{I}}_{k'}^{-1}$. The difference between $\hat{\mathcal{I}}_k^{-1}$ and the correct limiting covariance depends on how close Model k' is to the *true model*, in a Kullback-Leibler divergence sense. In general, the true model is unknown; therefore we cannot obtain an analytical expression of the correct limiting covariance. The recommended methodology thus represents in some cases an approximated version of the asymptotically exact one, but it is a methodology that can be easily applied and, fortunately, the resulting algorithms are expected to propose rarely models that are far from the true model because of the design of g , as long as a subset of \mathcal{K} are suitable approximations to the true model. In fact when the latter is true, the resulting algorithms are expected to behave similarly to the asymptotically exact ones. This discussion is made more precise in Section 3. In our numerical example, the resulting algorithms are sufficiently close to the asymptotically exact ones and n is sufficiently large so as to yield efficient samplers. Note that the covariance matrices can be estimated otherwise to obtain a generic asymptotically exact methodology; for instance, they can be estimated using pilot MCMC runs. This has for disadvantage of being computationally expensive and challenging.

We finish this section by noting that the approach of Green (2003) is also an approximated version of the asymptotically exact one when $\boldsymbol{\mu}_k = \hat{\mathbf{x}}_k$ and \mathbf{L}_k is such that $\mathbf{L}_k \mathbf{L}_k^T = \hat{\mathcal{I}}_k^{-1}$ for all k . We also note that a RJ user employing the approach of Green (2003) or that described above, already computes what is required to design $g(k, \cdot)$ as in Section 2.1. This user can therefore implement the proposed methodology without much additional effort.

2.3. Resulting RJ

We now present in Algorithm 2 a special case of Algorithm 1 with inputs specified as in Sections 2.1 and 2.2.

Algorithm 2 Proposed RJ

1. Sample $k' \sim g(k, \cdot)$ with $g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n))$ for all $k' \in \mathbf{N}(k)$.
- 2.(a) If $k' = k$, attempt a parameter update using a MCMC kernel of invariant distribution $\pi(\cdot | k, \mathbf{D}_n)$ while keeping the value of the model indicator k fixed.
- 2.(b) If $k' \neq k$, sample $\mathbf{y}_{k'} \sim \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$ and $u \sim \mathcal{U}[0, 1]$. If

$$u \leq \alpha_{\text{RJ}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'})) = 1 \wedge \frac{\pi(k', \mathbf{y}_{k'} | \mathbf{D}_n) g(k', k) \varphi(\mathbf{x}_k; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1})}{\pi(k, \mathbf{x}_k | \mathbf{D}_n) g(k, k') \varphi(\mathbf{y}_{k'}; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})},$$

set the next state of the chain to $(k', \mathbf{y}_{k'})$. Otherwise, set it to (k, \mathbf{x}_k) .

3. Go to Step 1.
-

Recall that we recommend to set h to the identity function when $\mathbf{N}(k) = \mathcal{K}$ for all k and to a locally-balanced function, namely $h(x) = x/(1+x)$, otherwise.

We recommend to apply HMC to update the parameters in Step 2.(a). HMC generate paths which evolve according to Hamiltonian dynamics, thus exploiting

the local structure of the target through gradient information; the endpoints of these paths are used as proposals within a MH scheme. This proposal mechanism allows the chain to take large steps while not ending up in areas with negligible densities. HMC thus has the ability to decorrelate, which is an interesting feature that is especially important in RJ given that the samplers may not spend a lot of consecutive iterations updating the parameters.

There exist several variants of HMC, the most popular being the *No-U-Turn sampler* (Hoffman and Gelman, 2014). The latter can be run automatically for parameter estimation using the probabilistic programming language *Stan* (see, e.g., Carpenter et al. (2017)). To our knowledge, it is not possible to directly use the Stan implementation within RJ to update the parameters, and given that the No-U-Turn sampler is difficult to implement, we have chosen to employ the “vanilla” version of HMC (see, e.g., Neal (2011)). The implementation of this version for a given model requires the specification of three free parameters: a step size, a trajectory length and a mass matrix. To tune these parameters for a given model, we recommend to first run Stan with the option `static HMC` (meaning that the vanilla HMC is employed), and to extract the step size and marginal empirical standard deviations of all parameters. This step size results from a tuning procedure and can thus directly be used. The mass matrix is set to a diagonal matrix with diagonal elements being these standard deviations. For the trajectory length, a grid search is applied and the best (found) value is used. The merit of each trajectory length is evaluated through the minimum of marginal effective sample sizes. Note that in this trans-dimensional framework, the momentum variables need in theory to be considered with care. Theoretically, we may consider that a momentum refreshment is performed every odd iteration, and that the algorithm proceeds as in Algorithm 2 every even iteration. Also, we need (in theory) to add or withdraw momentum variables when switching models. In practice, we do not have to proceed in this way. Given that momentum variables are only required when updating the parameters and that the expectations that one wants to approximate typically do not depend on these variables, we may generate them only when it is known that a parameter update is proposed (i.e. $k' = k$).

Some authors (for instance, Green (2003)) mentioned that using informed RJ samplers like Algorithm 2 may be problematic when it is required to gather information for each model before running the algorithms, because this is infeasible for large (or infinite) model spaces. The information that is required for an iteration of Algorithm 2 is in fact $\hat{\mathbf{x}}_{k'}$ and $\hat{\mathcal{I}}_{k'}$ for all $k' \in \mathbf{N}(k)$, and possibly free parameters for an HMC step in Step 2.(a). The required information can thus be gathered on the fly as the chains reach new models; it is unnecessary to gather all of it beforehand. The maximizers $\hat{\mathbf{x}}_{k'}$, matrices $\hat{\mathcal{I}}_{k'}$ and free parameters should be stored and reused next time they are needed. We only need to make sure that these are independent of the chain paths to have a valid procedure. For maximizers, any general-purpose optimization tool can be employed, but a starting point, if needed, should consequently not depend on the chain path to make sure there is no interference; the same recommendation follows for the computation of $\hat{\mathcal{I}}_k$ and tuning of HMC free parameters.

This strategy makes the implementation of informed RJ samplers possible, even if the model space is large or infinite, provided that the posterior model PMF concentrates on a reasonable number of models (in the sense that the number of different models visited during algorithm runs is on average reasonable). When the PMF concentrates on few models, this implementation strategy is expected to be highly effective as the information required for model switches and parameter updates will in practice be gathered for these few models and their neighbours only (if the algorithm is well initialized).

2.4. Methods to use when the approximations are not accurate

For finite n , the shapes of the posterior parameter densities under Model k and Model k' involved during a model-switch attempt may be quite different from bell curves. When this is the case, using normal approximations to the distributions, i.e. $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$ and $q_{k' \rightarrow k} = \varphi(\cdot; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1})$, may lead to high rejection rates. The method of Karagiannis and Andrieu (2013) allows to build upon these approximations by generating a sequence $(\mathbf{x}_k^{(0)}, \mathbf{y}_{k'}^{(0)}), \dots, (\mathbf{x}_k^{(T-1)}, \mathbf{y}_{k'}^{(T-1)})$ using inhomogeneous Markov kernels $K_{k \rightarrow k'}^{(t)}$, with $\mathbf{y}_{k'}^{(T-1)}$ being the proposal for the parameters of Model k' . More precisely, the starting point $(\mathbf{x}_k^{(0)}, \mathbf{y}_{k'}^{(0)})$ is such that $\mathbf{x}_k^{(0)} = \mathbf{x}_k$ and $\mathbf{y}_{k'}^{(0)} \sim \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$, and the following $(\mathbf{x}_k^{(t)}, \mathbf{y}_{k'}^{(t)})$ are sampled from $K_{k \rightarrow k'}^{(t)}((\mathbf{x}_k^{(t-1)}, \mathbf{y}_{k'}^{(t-1)}), \cdot)$, $K_{k \rightarrow k'}^{(t)}$ being, for each t , reversible with respect to an annealing intermediate distribution given by

$$\begin{aligned} \rho_{k \rightarrow k'}^{(t)}(\mathbf{x}_k^{(t)}, \mathbf{y}_{k'}^{(t)}) &\propto \left[\pi(k, \mathbf{x}_k^{(t)} \mid \mathbf{D}_n) \varphi(\mathbf{y}_{k'}^{(t)}; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1}) \right]^{1-\gamma_t} \\ &\times \left[\pi(k', \mathbf{y}_{k'}^{(t)} \mid \mathbf{D}_n) \varphi(\mathbf{x}_k^{(t)}; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1}) \right]^{\gamma_t}. \end{aligned} \tag{6}$$

The free parameter T is a positive integer, $\gamma_0 = 0, \gamma_T = 1$ and $\gamma_t = t/T$ for each $t \in \{1, \dots, T-1\}$.

We notice that when switching from Model k to Model k' , we start with distributions $\rho_{k \rightarrow k'}^{(t)}$ close to $\pi(k, \cdot \mid \mathbf{D}_n) \otimes \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$ to finish, after a transition phase, with distributions close to $\varphi(\cdot; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1}) \otimes \pi(k', \cdot \mid \mathbf{D}_n)$. Under some regularity conditions, choosing T large enough allows a smooth transition and makes that the last steps of the sequence, with $t \geq t^*$, are similar to steps of a time-homogeneous Markov chain with $\varphi(\cdot; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1}) \otimes \pi(k', \cdot \mid \mathbf{D}_n)$ as a stationary distribution (because $t/T \geq t^*/T \approx 1$), ensuring that $(\mathbf{x}_k^{(T-1)}, \mathbf{y}_{k'}^{(T-1)})$ is approximately distributed as $\varphi(\cdot; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1}) \otimes \pi(k', \cdot \mid \mathbf{D}_n)$. A proof can be found in Gagnon and Doucet (2021). Also, the acceptance probability in the resulting RJ, given by

$$\alpha_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1)})) := 1 \wedge \frac{g(k', k)}{g(k, k')} r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1)})) \tag{7}$$

with

$$r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1)})) := \prod_{t=0}^{T-1} \frac{\rho_{k \rightarrow k'}^{(t+1)}(\mathbf{x}_k^{(t)}, \mathbf{y}_{k'}^{(t)})}{\rho_{k \rightarrow k'}^{(t)}(\mathbf{x}_k^{(t)}, \mathbf{y}_{k'}^{(t)})}, \tag{8}$$

is such that $r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1)}))$ is a consistent estimator of $\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n)$ as $T \rightarrow \infty$. This shows that the resulting RJ behave asymptotically like the ones in which the parameter proposals are sampled from the correct conditional distributions $\pi(\cdot | k', \mathbf{D}_n)$, for fixed n .

Given that $r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1)}))$ can be seen as an estimator of $\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n)$, Andrieu et al. (2018) proposed to exploit this through a scheme allowing to generate in parallel N estimates to average them to reduce the variance. With T fixed, increasing N thus gets the resulting RJ closer to the ones in which the parameter proposals are sampled from the correct conditional distributions $\pi(\cdot | k', \mathbf{D}_n)$. In Appendix A, it is shown that the methods of Karagiannis and Andrieu (2013) and Andrieu et al. (2018) with their estimators of ratios of posterior probabilities can be used to enhance the approximations $\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)$ in $g(k, \cdot)$.

3. Asymptotic result and analysis

In Section 2.2, we described situations where using an informed but approximate design for $q_{k \rightarrow k'}$ is expected to yield RJ that behave like asymptotically exact ones, the latter being asymptotically equivalent to ideal RJ that sample parameters $\mathbf{y}_{k'}$ during model switches from the correct conditional distributions $\pi(\cdot | k', \mathbf{D}_n)$ and that use the unnormalized version of the model probabilities $\pi(k | \mathbf{D}_n)$ to construct g . In Section 3.1, we make this discussion precise and present a weak convergence result. We next provide in Section 3.2 an analysis of the limiting RJ.

3.1. Asymptotic result

The weak convergence result presented in this section is about weak convergence of Markov chains, but the convergence does not happen for almost all data sets \mathbf{D}_n . It rather occurs with a probability that becomes closer to 1 as $n \rightarrow \infty$, where the randomness comes from \mathbf{D}_n . A first statement of this kind in MCMC recently appeared in Schmon et al. (2021). It allows to exploit Bernstein-von Mises theorems which are about the convergence of posterior parameter distributions towards normal distributions in total variation (TV), with a probability that converges to 1 (thus not for almost all data sets \mathbf{D}_n). In this section, the randomness for all convergence in probability comes from \mathbf{D}_n ; the source of randomness is thus omitted to simplify the statements.

The first condition for the weak convergence result to hold is about an explicit independence between n and the state-space.

Assumption 1. *The model space \mathcal{K} and the parameter dimensions, i.e. d_k for all $k \in \mathcal{K}$, do not change with n .*

The second condition is that $\hat{\pi}(k | \mathbf{D}_n)$ is an asymptotically exact estimator of $\pi(k | \mathbf{D}_n)$, the latter admitting a limit $\bar{\pi}(k)$, for all k . In some cases, the posterior model mass concentrates. For instance when a model, say Model k^* ,

is well specified and regular, $\bar{\pi}(k^*) = 1$ (see, e.g., Johnson and Rossell (2012) in linear regression).

Assumption 2. For each $k \in \mathcal{K}$, the pair of random variables $(\hat{\pi}(k | \mathbf{D}_n), \pi(k | \mathbf{D}_n))$ is such that $|\hat{\pi}(k | \mathbf{D}_n) - \pi(k | \mathbf{D}_n)|$ and $|\pi(k | \mathbf{D}_n) - \bar{\pi}(k)|$ converge in probability towards 0 as $n \rightarrow \infty$.

The last assumption is about the asymptotic equivalence in TV between $q_{k \rightarrow k'}$ that is used to generate $\mathbf{u}_{k \rightarrow k'} = \mathbf{y}_{k'}$ and $\pi(\cdot | k', \mathbf{D}_n)$, for all k' . This happens when a Bernstein-von Mises theorem holds for each Model k' and $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\Sigma}_{k'}/n)$ with $\hat{\mathbf{x}}_{k'}$ and $\hat{\Sigma}_{k'}$ being consistent estimator of $\mathbf{x}_{k'}^*$ and $\Sigma_{k'}$, respectively, where $\Sigma_{k'}$ is the limiting covariance matrix (if we divide by n) of $\pi(\cdot | k', \mathbf{D}_n)$. When, for instance, Model k' is well defined and regular and data points are IID, $\hat{\Sigma}_{k'} = (\hat{I}_{k'}/n)^{-1}$ and $\Sigma_{k'}$ is the inverse information matrix of one data point.

Assumption 3. For each $k' \in \mathcal{K}$, there exist $\hat{\mathbf{x}}_{k'}$, $\mathbf{x}_{k'}^*$ and $\Sigma_{k'}$ such that $\hat{\mathbf{x}}_{k'} \rightarrow \mathbf{x}_{k'}^*$ and $TV(\pi(\cdot | k', \mathbf{D}_n), \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \Sigma_{k'}/n)) \rightarrow 0$, both in probability (Bernstein-von Mises theorem). Also, $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\Sigma}_{k'}/n)$ with $\hat{\Sigma}_{k'} \rightarrow \Sigma_{k'}$ in probability, implying that

$$TV(\varphi(\cdot; \hat{\mathbf{x}}_{k'}, \Sigma_{k'}/n), q_{k \rightarrow k'}) \rightarrow 0 \quad \text{and} \quad TV(\pi(\cdot | k', \mathbf{D}_n), q_{k \rightarrow k'}) \rightarrow 0,$$

in probability, for all k' (thus regardless of k).

As mentioned in Section 2.2, the proposed methodology represents in some cases an approximation to the asymptotically exact one, because for some probable transitions from Model k to Model k' , $TV(\varphi(\cdot; \hat{\mathbf{x}}_{k'}, \Sigma_{k'}/n), q_{k \rightarrow k'})$ does not converge to 0 if Model k' is misspecified. In fact, the proposed methodology is asymptotically exact only if all models are well specified, or at least, if some are misspecified, then $\bar{\pi}$ must assign a probability of 0 to them. Indeed, in the latter case, considering that the chain is currently visiting Model k with $\bar{\pi}(k) > 0$,

$$g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n)/\hat{\pi}(k | \mathbf{D}_n)) \rightarrow 0,$$

for all misspecified models with $k' \neq k$, so that the proposal distributions $q_{k \rightarrow k'}$ with $TV(\varphi(\cdot; \hat{\mathbf{x}}_{k'}, \Sigma_{k'}/n), q_{k \rightarrow k'}) \not\rightarrow 0$ will asymptotically never be used.

In practice, it is expected that all models are misspecified, but that some well approximate the true model. We now briefly explain why the proposed methodology is a good approximation to the asymptotically exact one in this case and next present the weak convergence result. Denote by $\mathcal{K}^* \subset \mathcal{K}$ the subset formed of the good models, and consider that for all $k \in \mathcal{K}^*$ and for large enough n , $TV(\varphi(\cdot; \hat{\mathbf{x}}_{k'}, \Sigma_{k'}/n), q_{k \rightarrow k'})$ is small (because $\hat{\Sigma}_{k'} = (\hat{I}_{k'}/n)^{-1}$ and $\Sigma_{k'}$ are similar), implying that $TV(\pi(\cdot | k', \mathbf{D}_n), q_{k \rightarrow k'})$ is small. Consider also to simplify that all models are either good or poor approximations to the true model, and that for all the poor models, i.e. with $k' \notin \mathcal{K}^*$, $g(k, k') \rightarrow 0$ for $k \in \mathcal{K}^*$ (because $\hat{\pi}(k' | \mathbf{D}_n)/\hat{\pi}(k | \mathbf{D}_n) \rightarrow 0$). Therefore, for large enough n , if the chain is currently visiting Model $k \in \mathcal{K}^*$, then with probability close to 1, a

model in \mathcal{K}^* is proposed, say Model k' , with

$$g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)) \approx h(\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n)),$$

and next, parameters $\mathbf{y}_{k'} = \mathbf{u}_{k \rightarrow k'} \sim q_{k \rightarrow k'}$ are proposed, with $\text{TV}(\pi(\cdot | k', \mathbf{D}_n), q_{k \rightarrow k'})$ small. In this situation, the proposed RJ are well approximating the asymptotically exact ones. We keep this situation in mind for the rest of Section 3.

We now introduce notation that are required to present the weak convergence result. Use $\{(K, \mathbf{X}_K)_n(m) : m \in \mathbb{N}\}$ to denote a Markov chain simulated by a RJ that targets $\pi(\cdot, \cdot | \mathbf{D}_n)$ and sets $g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n))$, $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\Sigma}_{k'} / n)$ and $\mathcal{D}_{k \rightarrow k'}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) = (\mathbf{u}_{k \rightarrow k'}, \mathbf{x}_k) = (\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k})$, for all k and $k' \in \mathbf{N}(k)$. Use $\{(K, \mathbf{X}_K)_{\text{ideal}}(m) : m \in \mathbb{N}\}$ to denote a Markov chain simulated by an ideal RJ that targets $\pi(\cdot, \cdot | \mathbf{D}_n)$ as well, but instead sets $g(k, k') \propto h(\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n))$ and $q_{k \rightarrow k'} = \pi(\cdot | k', \mathbf{D}_n)$ with $\mathcal{D}_{k \rightarrow k'}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) = (\mathbf{u}_{k \rightarrow k'}, \mathbf{x}_k) = (\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k})$, for all k and $k' \in \mathbf{N}(k)$.

Because the target distribution concentrates, we need to apply a transformation to the Markov chains to obtain a non-trivial limit. We apply a transformation reflecting that even for large n , the samplers continue to explore the parameter spaces, but at different scales given that the parameters are continuous variables. In contrast, if the posterior PMF $\pi(\cdot | \mathbf{D}_n)$ concentrates on say, one model, then fewer and fewer models are visited during an algorithm run as more and more of them have negligible mass as $n \rightarrow \infty$. The transformation that is applied aims at reflecting this reality and thus obtaining a limiting situation that represents an approximation to what one encounters in practice. We more precisely standardize the parameter variable \mathbf{X}_K , but leave the model indicator K as is, i.e. we define $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$ and $\{(K, \mathbf{Z}_K)_{\text{ideal}}(m) : m \in \mathbb{N}\}$ such that $(K, \mathbf{Z}_K)_n(m) = (K, \sqrt{n}(\mathbf{X}_K - \hat{\mathbf{x}}_K))_n(m)$ and $(K, \mathbf{Z}_K)_{\text{ideal}}(m) = (K, \sqrt{n}(\mathbf{X}_K - \hat{\mathbf{x}}_K))_{\text{ideal}}(m)$, respectively, for all m .

To be an admissible limit of $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$, $\{(K, \mathbf{Z}_K)_{\text{ideal}}(m) : m \in \mathbb{N}\}$ would need to be independent of n , but it is readily seen to be not the case. For instance, the invariant conditional density of \mathbf{Z}_K given K is $\pi(\cdot / \sqrt{n} + \hat{\mathbf{x}}_k | k, \mathbf{D}_n) / \sqrt{n}$, which is asymptotically equivalent to a normal with mean $\mathbf{0}$ and covariance Σ_k , but not equal to it for all n . We thus introduce a process for which the transformation is independent of n and prove that $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$ and $\{(K, \mathbf{Z}_K)_{\text{ideal}}(m) : m \in \mathbb{N}\}$ weakly converge towards it, showing that they all share a similar behaviour for large enough n .

Denote by $\{(K, \mathbf{X}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ this process. It is an RJ which targets a distribution such that $\mathbf{X}_K | K \sim \varphi(\cdot; \hat{\mathbf{x}}_K, \Sigma_K / n)$ and $K \sim \bar{\pi}$. It sets $g(k, k') \propto h(\bar{\pi}(k') / \bar{\pi}(k))$, $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \Sigma_{k'} / n)$, and $\mathcal{D}_{k \rightarrow k'}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) = (\mathbf{u}_{k \rightarrow k'}, \mathbf{x}_k) = (\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k})$, for all k with $\bar{\pi}(k) > 0$ and $k' \in \mathbf{N}(k)$. Because the chains will be assumed to start in stationarity, there is no need to define the functions for k with $\bar{\pi}(k) = 0$. Denote the standardized version of $\{(K, \mathbf{X}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ by $\{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$, which is such that $(K, \mathbf{Z}_K)_{\text{limit}}(m) = (K, \sqrt{n}(\mathbf{X}_K - \hat{\mathbf{x}}_K))_{\text{limit}}(m)$, for all m . We denote the weak

convergence of, for instance, $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$ to $\{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ by $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\} \implies \{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$.

Theorem 1 (Weak convergence). *Under Assumptions 1 to 3, we have that $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\} \implies \{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ and $\{(K, \mathbf{Z}_K)_{\text{ideal}}(m) : m \in \mathbb{N}\} \implies \{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$, both in probability as $n \rightarrow \infty$, provided that all chains start in stationarity.*

Formally, in Theorem 1, it is considered that, for each sampler, the proposal mechanism for parameter updates and model switches is the same. For instance, in the implementable RJ (that simulate $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$), a normal is used to sample the parameter proposals in parameter-update attempts. This is often not the case in practice; for instance in our numerical example, we use HMC steps. To accommodate for such types of parameter-update mechanisms, an additional, rather technical, assumption is required about the convergence of the associated Markov kernels in some sense as $n \rightarrow \infty$. We proceeded in that way to simplify and for brevity.

3.2. Analysis of the limiting RJ

The stochastic process $\{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ can be used as a proxy to $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$ in the large-sample regime. An analysis of it thus helps understand how $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$, and thus $\{(K, \mathbf{X}_K)_n(m) : m \in \mathbb{N}\}$, behave in this regime. The stationary distribution of $\{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ is such that $\mathbf{Z}_K | K \sim \varphi(\cdot; \mathbf{0}, \Sigma_K)$ and $K \sim \bar{\pi}$.

A theoretical result that can be established in the situation where the posterior model PMF concentrates on one model is the dominance of the limiting RJ over any RJ targeting the same distribution and using the same functions, except $g(k, \cdot)$. This is true because the limiting RJ does not waste iterations trying to propose to switch to models other than the best. This shows the relevance of using informed RJ which automatically incorporate such features of the target.

Let us denote by $P_{\text{limit},1}$ and $P_{\text{limit},2}$ the Markov kernels of the limiting RJ setting $g(k, k') \propto h(\bar{\pi}(k')/\bar{\pi}(k))$ and that of the other RJ with the only difference being in g , respectively. The result is a Peskun-Tierney ordering (Peskun, 1973; Tierney, 1998), meaning an order on the Markov kernels, implying an order on the asymptotic variances. We use $\text{var}(f, P_{\text{limit},i})$ to denote the asymptotic variance of the function f applied to the Markov chain of transition kernel $P_{\text{limit},i}$ at equilibrium, $i = 1, 2$.

Proposition 1. *Assume that there exists k^* with $\bar{\pi}(k^*) = 1$. Then, for any $\mathbf{z}_{k^*} \in \mathbb{R}^{d_{k^*}}$ and measurable set A ,*

$$P_{\text{limit},1}((k^*, \mathbf{z}_{k^*}), A \setminus \{(k^*, \mathbf{z}_{k^*})\}) \geq P_{\text{limit},2}((k^*, \mathbf{z}_{k^*}), A \setminus \{(k^*, \mathbf{z}_{k^*})\}),$$

implying that for any square-integrable function f ,

$$\text{var}(f, P_{\text{limit},1}) \leq \text{var}(f, P_{\text{limit},2}).$$

Note that this result holds for globally-balanced proposal distributions as well. What is required is that $h(x) = 0$ when $x = 0$ and $h(x) > 0$ when $x > 0$, which holds when h is the identity function.

In the limiting RJ, the parameters are sampled from the correct conditional distributions and therefore the transition probabilities for K are the same as a MH sampler targeting $\bar{\pi}$ using $g(k, k') \propto h(\bar{\pi}(k')/\bar{\pi}(k))$ with acceptance probabilities given by

$$1 \wedge \frac{\bar{\pi}(k') g(k', k)}{\bar{\pi}(k) g(k, k')} = 1 \wedge \frac{c(k)}{c(k')}. \quad (9)$$

This implies that any analysis of MH samplers with h set to a locally-balanced function (which is what we recommend in typical cases) applies to the limiting RJ. For a thorough analysis of such samplers, we refer the reader to Zanella (2020). Recall that $c(k)$ and $c(k')$ are the normalizing constants of $g(k, \cdot)$ and $g(k', \cdot)$, respectively. As mentioned in Section 1.3, Zanella (2020) proved that $c(k)/c(k')$ are close to 1 in some situations, suggesting that locally-balanced samplers are efficient.

Note that in practice even when $\pi(\cdot | \mathbf{D}_n)$ concentrates, we have that $\pi(k | \mathbf{D}_n) > 0$ for any finite n and any k , implying that before finding the best model with a significantly larger probability, the chain may move around for a while. For these transitions with acceptance probabilities close to $\alpha_{\text{MH}}(k, k')$ in the large-sample regime, the result of Zanella (2020), stating that $\alpha_{\text{MH}}(k, k') \rightarrow 1$ as the ambient space becomes larger and larger, holds under some conditions; this suggests again an efficiency in exploring the state-space and thus finding the best model.

When h is the globally-balanced function, i.e. the identity function, and $\mathbf{N}(k) = \mathcal{K}$ for all k , then the limiting RJ proposes models and parameters independently of the current state of the chain and these proposals are always accepted. It thus corresponds to independent Monte Carlo sampling, which is often seen as the gold standard in MCMC.

4. Application: Variable selection in wholly-robust linear regression

In this section, we apply the proposed RJ methodology to a variable-selection problem in a wholly-robust linear regression. An overview of wholly-robust linear regression is provided in Section 4.1 and then the computational results are presented and analysed in Section 4.2.

4.1. Wholly-robust linear regression

A new technique emerged to gain robustness against outliers in parametric modelling: replace the traditional distribution assumption (which is a normal assumption in the problems previously studied) by a super-heavy-tailed distribution assumption (Desgagné, 2015; Desgagné and Gagnon, 2019; Gagnon, Desgagné and Bédard, 2020a; Gagnon, Bédard and Desgagné, 2021). The ra-

tionale is that this latter assumption is more adapted to the eventual presence of outliers by giving higher probabilities to extreme values. A proof of effectiveness of the approach resides in the following: the posterior distribution converges towards that based on the non-outliers only (i.e. excluding the outliers) as the outliers move further and further away from the bulk of the data. This theoretical result corresponds to a property in Bayesian statistics called *whole robustness* (Desgagné, 2015), and implies a *conflict resolution* (O'Hagan and Pericchi, 2012). As explained in the papers cited above, the models with super-heavy-tailed distributions have built-in robustness that resolve conflicts due to contradictory sources of information in a sensitive way. It takes full consideration of non-outliers and excludes observations that are undoubtedly outlying; in between these two extremes, it balances and bounds the impact of possible outliers, reflecting that there is a uncertainty about whether or not these observations really are outliers.

These features of wholly-robust models are appealing, but they come at price, mainly a computational one. The super-heavy-tailed distribution used in the papers cited above cannot indeed be represented as a scale mixture of normal distributions, contrarily to the Student distribution which is commonly assumed in Bayesian robust models. A scale mixture representation of the Student allows a straightforward implementation of the Gibbs sampler in, for instance, linear regression, for parameter estimation and variable selection (Verdinelli and Wasserman, 1991). A motivation for assuming a distribution with heavier tails than the Student is that the latter only allows to reach *partial robustness*. In linear regression, partial robustness translates into robust regression-coefficient point estimates, but inflated posterior variances, which contaminates uncertainty assessments and thus model selection (see Hayashi (2020) for a proof of partial robustness in linear regression).

In Gagnon, Desgagné and Bédard (2020a), the convergence of the posterior distribution is proved under the most general linear-regression framework, encompassing analysis of variance and covariance (ANOVA and ANCOVA), and variable selection. In this section, we apply the RJ methodology presented in the previous sections to sample from a joint posterior distribution of wholly-robust linear regression models and their parameters. RJ is thus required for this task. The data analysed are the same prostate-cancer data mentioned in Figure 1.

The models are thus linear regressions, but the errors follow a super-heavy-tailed distribution. See Appendix B for all the details about the models and inputs required for RJ implementation. The super-heavy-tailed distribution used is called *log-Pareto-tailed normal* (LPTN). This distribution was introduced by Desgagné (2015). Its density exactly matches that of the normal on the interval $[-\tau, \tau]$, where $\mathbb{P}(-\tau \leq \varphi(\cdot; 0, 1) \leq \tau) = \rho$. Outside of this area, the tails of this continuous density behave as a log-Pareto: $(1/|x|)(\log |x|)^{-\lambda-1}$, hence its name. The only free parameter of this distribution is ρ : the parameter λ is a function of ρ and τ , the latter being itself a function of ρ . The linear regression with a LPTN is thus expected to behave similarly to the traditional normal one in the absence of outliers. Not only this is the case in the absence of outliers, but the limiting robust-regression posterior distribution (as the distance between

the outliers and the bulk of the data approaches infinity) is also similar to the normal one, the latter being instead based on the non-outliers only.

This resemblance is useful in our informed RJ framework because, first, it suggests that the proposed methodology is an approximation to the asymptotically exact one if some linear regressions are good approximations to the true model. Indeed, in the normal linear-regression framework, Assumptions 1 and 2 of Section 3.1 hold, and a Bernstein-von Mises theorem holds for each model. It is thus expected to be the case in the robust LPTN framework as well, yet it is much more difficult to prove.

Second, the resemblance between the wholly-robust and normal regressions provides us with an approximation to the observed information matrix. We indeed set $\hat{\mathcal{I}}_k$ in the robust LPTN framework to the observed information matrix of the normal regression, but evaluated at the maximizer under the wholly-robust model. The matrix $\hat{\mathcal{I}}_k$ is thus basically the maximum a posteriori estimate of the scale parameter in the robust model multiplied by the observed covariate correlation matrix. The approximation is thus expected to be accurate when there are no severe outliers among the covariate observations. One may instead use a robust version of the correlation matrix (see Gagnon, Bédard and Desgagné (2021) for such a robust alternative), but this is not investigated here for brevity. Whether or not a robust version of the covariate matrix is used is expected to have no impact on our comparison of the algorithms in the next section because they all use $q_{k \rightarrow k'} = \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$ and the matrix $\hat{\mathcal{I}}_{k'}$ only appears in $q_{k \rightarrow k'}$. Indeed, $|\hat{\mathcal{I}}_{k'}|^{-1/2}$ in $\hat{\pi}(k' | \mathbf{D}_n)$ (4) cancels with another term in the same expression.

Finally, the resemblance between the wholly-robust and normal regressions is useful to validate our RJ computer code. Indeed, the robust approach yields an outlier-detection method and a “clean” data set without outliers can thus be identified. Based on this data set, posterior coefficient estimates and model probabilities of the normal models can be computed using their closed-form expressions. These estimates and probabilities are expected to be similar to those obtained under the robust models.

4.2. Results and analysis

The performance of the different algorithms are summarized in Table 1 and Figure 2. The results for the uninformed RJ are based on 1,000 runs of 100,000 iterations, with burn-ins of 10,000. The uninformed RJ is in fact a “non-fully” informed RJ corresponding to Algorithm 2 with the only difference that $g(k, \cdot) = \mathcal{U}\{\mathbf{N}(k)\}$ instead of locally- or globally-balanced proposal distributions such that $g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n))$ for all $k' \in \mathbf{N}(k)$. It may thus be seen as corresponding to the approach of Green (2003). For a fair comparison with this RJ, the number of iterations for which the fully informed RJ, i.e. Algorithm 2, is run is reduced to account for the additional complexity of using informed distributions $g(k, \cdot)$. The number of iterations has to be reduced to 85,000 to reach the same runtime as the uninformed RJ; one run takes about one minute in

TABLE 1

Performance of the uninformed and informed RJ in terms of model-switching acceptance rate, model-visit rate and relative increase in TV with respect to Algorithm 2 with $h(x) = \sqrt{x}$

Algorithms	Acc. rate	Visit rate	Rel. increase in TV
Algorithm 2 w. $h(x) = \sqrt{x}$	66%	55%	—
Algorithm 2 w. $h(x) = x/(1+x)$	67%	53%	2%
Uninformed RJ	30%	27%	28%
Algorithm 2 w. $h(x) = x$	57%	46%	40%

R on a i9 computer.¹ To measure performance, we display the model-switching acceptance rates, model-visit rates and relative increases in TV with respect to the best sampler, which is Algorithm 2 with $h(x) = \sqrt{x}$. The TVs are in between approximated model distributions from RJ outputs and the posterior model PMF².

The model-switching acceptance rates are the acceptance rates, but computed considering only the iterations in which model switches are proposed. The visit rates are the average number of model switches in one run, reported per iteration. For both these measures, we count the number of accepted model switches, and this number is divided by either the number of proposed model switches or total number of iterations. These rates are thus similar but they convey different information: model-switching acceptance rates reflect the quality of the proposal distributions used during model switches, while visit rates measure the propensity to propose a switch to another model (and accept it). They together allow to understand why, for instance, Algorithm 2 with $h(x) = \sqrt{x}$ is (slightly) better than Algorithm 2 with $h(x) = x/(1+x)$ in terms of TV. Indeed, even if it has a (slightly) worse model-switch proposal scheme (model-switching acceptance rates: 66% vs 67%), it proposes more often to switch models (visit rates: 55% vs 53%), which has an impact on the TV.

This robust-variable-selection application shows that even in a moderate-size example, in this case with 8 covariates and 256 models, the locally-balanced proposal distributions lead to RJ that can outperform RJ proposing models uniformly at random, *at fixed computational budget*. In contrast, the globally-balanced proposal distribution does not yield an improvement that is significant enough to compensate for the decrease in number of iterations. We observed that even when the RJ with the function $h(x) = x$ is run for the same number of iterations as the uninformed one, the former only slightly outperforms the latter.

For the informed RJ to be effective, the approximations on which they rely have to be accurate. Recall that their accuracy depends on the sample size (at least in some scenarios regarding the approximations of $\pi(\cdot | k', \mathbf{D}_n)$ by $q_{k \rightarrow k'}$, as explained in Section 3.1). In this example, $n = 97$. To understand

¹We do not claim optimality of the computer code; we coded all algorithms in the same fashion for a fair comparison.

²We used accurate MCMC approximations to the posterior model probabilities. We verified that the TV goes to 0 for all algorithms as the number of iterations increases.

if this is actually large for such a robust-linear-regression problem with 8 covariates, we can compare the model-switching acceptance rates of Algorithm 2 with those of the ideal RJ they approximate. Recall that ideal RJ set $g(k, k') \propto h(\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n))$ and $q_{k \rightarrow k'} = \pi(\cdot | k', \mathbf{D}_n)$. It is possible to evaluate their model-switching acceptance rates because these correspond to the acceptance rates of MH samplers targeting the PMF $\pi(\cdot | \mathbf{D}_n)$ using $g(k, k') \propto h(\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n))$, computed considering only iterations in which $k' \neq k$ is proposed. When $h(x) = x/(1+x)$, the acceptance rate of the MH sampler is 88%, which is sufficiently close to 1 to suggest that the asymptotic regime presented in Zanella (2020) where the size of the discrete space increases and the acceptance probabilities converge to 1 is nearly reached. To understand which approximations explain the gap between Algorithm 2 with $h(x) = x/(1+x)$ and a model-switching acceptance rate of 67% and its ideal counterpart with a rate of 88%, we can look at the model-switching acceptance rate of a RJ using methods to sample parameter proposals from a distribution closer to $\pi(\cdot | k', \mathbf{D}_n)$ and next at that of a RJ additionally using a method making $g(k, k')$ closer to being proportional to $h(\pi(k' | \mathbf{D}_n) / \pi(k | \mathbf{D}_n))$ (recall the overview of the methods presented in Section 2.4). It is 84% for the first RJ (with $T = N = 10$), and it increases to 87% for the second RJ. Using such methods simply make RJ too computationally expensive to be competitive in this example. Recall that their precision parameters can be increased to make the second RJ arbitrarily close to the ideal one with a model-switching acceptance rate of 88%.

These results suggest that n is of moderate size, in the sense that the approximations used in Algorithm 2 are good enough to yield an improvement, but it is expected that the improvement would be even more marked if we had access to more data points. Also, the results suggest that the approximations explaining the most part of the gap between Algorithm 2 with $h(x) = x/(1+x)$ and a model-switching acceptance rate of 67% and its ideal counterpart with a rate of 88% are the approximations of $\pi(\cdot | k', \mathbf{D}_n)$ by $q_{k \rightarrow k'}$. We investigated to see if there are significant differences between the covariance matrices used in $q_{k \rightarrow k'}$ and those of $\pi(\cdot | k', \mathbf{D}_n)$, and it is not the case. This means that the approximations of $\pi(\cdot | k', \mathbf{D}_n)$ by $q_{k \rightarrow k'}$ are not as accurate as we would like because the densities $\pi(\cdot | k', \mathbf{D}_n)$ are significantly different from bell curves. Our analysis also allows to rule out an issue with using the same observed covariance matrix as the normal regression in the robust LPTN framework. In fact, a robust-linear-regression analysis reveals the presence of outliers, but only for models with negligible posterior mass (both in the robust and non-robust frameworks). We finally note that, for models with non-negligible posterior mass, there is no gross violation of the assumptions underlying linear regression.

5. Discussion

In this paper, we proposed simple fully-informed RJ and described the situations in which they are expected to be efficient: when the sample size is sufficiently large, Assumptions 1–2 of Section 3.1 are verified, some models among those considered well approximate the true data generating process and Bernstein-




FIG 2. Number of visited models as a function of the iteration number in a typical run of the uninformed RJ and Algorithm 2 with $h(x) = \sqrt{x}$

von Mises theorems hold for these models. Informed proposal distributions are crucial when the model probabilities and parameter densities vary significantly within neighbourhoods of states of the Markov chains. They are expected to vary significantly in the large-sample regime in which the target concentrates. But we noticed with our numerical example that they may vary greatly even when this large-sample regime is not reached. The proposed RJ show major improvements in this example as the chains spend less iterations at the same state, compared with the RJ that naively proposes models uniformly at random and thus often tries to reach low probability models, leading to higher rejection rates.

Yet, the proposed samplers are reversible which allows them to return to recently visited models often. The next step in the line of research of trans-dimensional samplers for non-nested model selection is to propose sampling schemes which do not suffer from this diffusive behaviour, but instead induce persistent movement in the model-indicator process. A first step in this direction has recently been made by Gagnon and Maire (2020), but their approach cannot be used in all contexts of non-nested model selection. It can nevertheless be applied in any variable-selection problem. It is however expected to be efficient when the posterior model mass is diffused over a large number of models with different number of covariates, which is not the case in our numerical example.

Appendix A: Improving the approximations

In practice, the sample size may not be large enough for the approximations in the proposal mechanisms to be accurate. In Sections A.1 and A.2, we present the details of the methods reviewed in Section 2.4 that allow to improve upon the normal-distribution approximations to $\pi(\cdot | k, \mathbf{D}_n)$. These methods turn out to be useful for improving upon the approximations forming the model proposal distributions $g(k, \cdot)$ as well; this is explained in detail in Section A.3.

A.1. RJ with the method of Karagiannis and Andrieu (2013)

In the definition of the annealing intermediate distributions $\rho_{k \rightarrow k'}^{(t)}$ in (6), it is considered that $\mathcal{D}_{k \rightarrow k'}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'}) = (\mathbf{u}_{k \rightarrow k'}, \mathbf{x}_k) = (\mathbf{y}_{k'}, \mathbf{u}_{k' \rightarrow k})$, implying that

$|J_{\mathcal{D}_{k \rightarrow k'}}(\mathbf{x}_k, \mathbf{u}_{k \rightarrow k'})| = 1$, and $q_{k' \rightarrow k} = \varphi(\cdot; \hat{\mathbf{x}}_k, \hat{\mathcal{I}}_k^{-1})$. For the general definition, see Karagiannis and Andrieu (2013).

We now present in Algorithm 3 an informed RJ incorporating the method of Karagiannis and Andrieu (2013). Recall that the acceptance probability in this RJ is defined in (7). Note that Algorithm 3 corresponds to Algorithm 2 when $T = 1$; no path is sampled in Step 2.(b).

Algorithm 3 RJ incorporating the method of Karagiannis and Andrieu (2013)

1. Sample $k' \sim g(k, \cdot)$ with $g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n))$ for all $k' \in \mathbf{N}(k)$.
 - 2.(a) If $k' = k$, attempt a parameter update using a MCMC kernel of invariant distribution $\pi(\cdot | k, \mathbf{D}_n)$ while keeping the value of the model indicator \mathbf{k} fixed.
 - 2.(b) If $k' \neq k$, sample $\mathbf{y}_{k'}^{(0)} \sim \varphi(\cdot; \hat{\mathbf{x}}_{k'}, \hat{\mathcal{I}}_{k'}^{-1})$ and $u \sim \mathcal{U}[0, 1]$, and set $\mathbf{x}_k^{(0)} := \mathbf{x}_k$. Sample a path $(\mathbf{x}_k^{(1)}, \mathbf{y}_{k'}^{(1)}), \dots, (\mathbf{x}_k^{(T-1)}, \mathbf{y}_{k'}^{(T-1)})$, where $(\mathbf{x}_k^{(t)}, \mathbf{y}_{k'}^{(t)}) \sim K_{k \rightarrow k'}^{(t)}((\mathbf{x}_k^{(t-1)}, \mathbf{y}_{k'}^{(t-1)}), \cdot)$. If $u \leq \alpha_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1)}))$ set the next state of the chain to $(k', \mathbf{y}_{k'}^{(T-1)})$. Otherwise, set it to (k, \mathbf{x}_k) .
 3. Go to Step 1.
-

Using the same proof technique as in Gagnon and Doucet (2021), it can be proved that under regularity conditions a Markov chain simulated by Algorithm 3 converges weakly to that of a RJ which is able to sample from $\pi(\cdot | k, \mathbf{D}_n)$ for all k with acceptance probabilities α_{MH} but using $g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n))$, as $T \rightarrow \infty$, for fixed n . Note that the weak convergence in this case is not in probability because the target is considered non-random (contrarily to the framework under which Theorem 1 is stated).

As T increases, it is expected that the acceptance probabilities increase steadily towards α_{MH} until convergence is reached. A RJ user thus has to find a balance between a RJ close to be ideal and the computational cost associated to this. As explained in Karagiannis and Andrieu (2013), the computational cost scales linearly with T . In some situations, the improvement as a function of T is very marked for $T \leq T_0$, effectively offsetting the computational cost. This is not the case in our numerical example.

Karagiannis and Andrieu (2013) prove that under two conditions Algorithm 3 is valid, in the sense that the target distribution is an invariant distribution. These conditions are the following.

Symmetry condition: For $t = 1, \dots, T - 1$ the pairs of transition kernels $K_{k \rightarrow k'}^{(t)}(\cdot, \cdot)$ and $K_{k' \rightarrow k}^{(T-t)}(\cdot, \cdot)$ satisfy

$$K_{k \rightarrow k'}^{(t)}((\mathbf{x}_k, \mathbf{y}_{k'}), \cdot) = K_{k' \rightarrow k}^{(T-t)}((\mathbf{x}_k, \mathbf{y}_{k'}), \cdot) \quad \text{for any } (\mathbf{x}_k, \mathbf{y}_{k'}). \quad (10)$$

Reversibility condition: For $t = 1, \dots, T - 1$, and for any $(\mathbf{x}_k, \mathbf{y}_{k'})$ and $(\mathbf{x}'_k, \mathbf{y}'_{k'})$,

$$\begin{aligned} & \rho_{k \rightarrow k'}^{(t)}(\mathbf{x}_k, \mathbf{y}_{k'}) K_{k \rightarrow k'}^{(t)}((\mathbf{x}_k, \mathbf{y}_{k'}), (\mathbf{x}'_k, \mathbf{y}'_{k'})) \\ &= \rho_{k \rightarrow k'}^{(t)}(\mathbf{x}'_k, \mathbf{y}'_{k'}) K_{k \rightarrow k'}^{(t)}((\mathbf{x}'_k, \mathbf{y}'_{k'}), (\mathbf{x}_k, \mathbf{y}_{k'})). \end{aligned} \quad (11)$$

As mentioned in Karagiannis and Andrieu (2013), (10) is verified if for all t , $K_{k \rightarrow k'}^{(t)}(\cdot, \cdot)$ and $K_{k' \rightarrow k}^{(T-t)}(\cdot, \cdot)$ are MH kernels sharing the same proposal distributions. (11) is satisfied if $K_{k \rightarrow k'}^{(t)}(\cdot, \cdot)$ is a MH kernel targeting $\rho_{k \rightarrow k'}^{(t)}$. We recommend to use MALA (Metropolis-adjusted Langevin algorithm, Roberts and Tweedie (1996)) proposal distributions whenever this is possible; see Karagiannis and Andrieu (2013) for other examples. If MALA proposal distributions are used, a step size is required for each pair (k, k') defining a transition from Model k to Model k' . For any of these model switches, we recommend to set the step size to $\ell/(d_k + d_{k'})^{1/6}$, according to the findings in Roberts and Rosenthal (1998), and to tune ℓ . To achieve the latter, do a line search with a fixed value of T and choose ℓ that maximizes the model-switching acceptance rate, so that the resulting RJ is the closest to the ideal one.

A.2. RJ with additionally the method of Andrieu et al. (2018)

Denote the N estimates of $\pi(k' | \mathbf{D}_n)/\pi(k | \mathbf{D}_n)$ produced by the scheme of Karagiannis and Andrieu (2013) by $r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1,1)})), \dots, r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1,N)}))$. Denote the average (with simplified notation) by

$$\bar{r}(k, k') := \frac{1}{N} \sum_{j=1}^N r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1,j)})).$$

We now present in Algorithm 4 the RJ additionally incorporating the method of Andrieu et al. (2018).

No additional assumptions to those presented in Section A.1 are required to guarantee that Algorithm 4 is valid. Andrieu et al. (2018) prove that increasing N decreases monotonically the asymptotic variances of the Monte Carlo estimates produced by RJ incorporating their approach.

It is expected that increasing N (as increasing T in the last section) leads to a steady increase in the acceptance probabilities until convergence is reached. As with Algorithm 3, there exists a balance between a RJ close to be ideal and the computational cost associated to this. An advantage of the approach presented in this section is that several operations can be executed in parallel, so that the computational burden is alleviated. The computational cost (over that of using the approach of Karagiannis and Andrieu (2013)) nevertheless depends on N because computational overheads have to be taken into account. In our numerical example, the improvement as a function of N is not significant enough to offset the computational cost.

A.3. Improving the model proposal distribution

We have seen in Section A.2 that $\bar{r}(k, k')$ and the ratios r_{RJ2} forming it are estimators of $\pi(k' | \mathbf{D}_n)/\pi(k | \mathbf{D}_n)$. They can thus be used to enhance the approximations to the posterior probability ratios in $g(k, \cdot)$.

Algorithm 4 RJ additionally incorporating the method of Andrieu et al. (2018)

1. Sample $k' \sim g(k, \cdot)$ with $g(k, k') \propto h(\hat{\pi}(k' | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n))$ for all $k' \in \mathbf{N}(k)$.
- 2.(a) If $k' = k$, attempt a parameter update using a MCMC kernel of invariant distribution $\pi(\cdot | k, \mathbf{D}_n)$ while keeping the value of the model indicator k fixed.
- 2.(b) If $k' \neq k$, generate $u_a, u_c \sim \mathcal{U}(0, 1)$. If $u_c \leq 1/2$ go to Step 2.(b-i), otherwise go to Step 2.(b-ii).
- 2.(b-i) Sample N proposals $\mathbf{y}_{k'}^{(T-1,1)}, \dots, \mathbf{y}_{k'}^{(T-1,N)}$ as in Step 2.(b) of Algorithm 3. Sample j^* from a PMF such that $\mathbb{P}(J^* = j) \propto r_{\text{RJ2}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'}^{(T-1,j)}))$. If

$$u_a \leq \frac{g(k', k)}{g(k, k')} \bar{\tau}(k, k'),$$

set the next state of the chain to $(k', \mathbf{y}_{k'}^{(T-1,j^*)})$. Otherwise, set it to (k, \mathbf{x}_k) .

- 2.(b-ii) Sample one forward path as in Step 2.(b) of Algorithm 3. Denote the endpoint by $\mathbf{y}_{k'}^{(T-1,1)}$. From $\mathbf{y}_{k'}^{(T-1,1)}$, sample $N - 1$ reverse paths again as in Step 2.(b) of Algorithm 3, yielding $N - 1$ proposals for the parameters of Model k . If

$$u_a \leq \frac{g(k', k)}{g(k, k')} \bar{\tau}(k', k)^{-1},$$

set the next state of the chain to $(k', \mathbf{y}_{k'}^{(T-1,1)})$. Otherwise, set it to (k, \mathbf{x}_k) .

3. Go to Step 1.
-

If we want to enhance the PMF $g(k, \cdot)$, we need to improve $\hat{\pi}(l | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)$ for all $l \in \mathbf{N}(k)$ as these are all involved in the construction of the PMF. Also, once k' has been sampled, we need to do the same for $g(k', \cdot)$ given that this PMF comes into play in the computation of the acceptance probabilities (see, e.g., Algorithm 4). We thus need parameter proposals $\mathbf{y}_l^{(T-1,1)}, \dots, \mathbf{y}_l^{(T-1,N)}$ for all Models $l \in \mathbf{N}(k)$, and also for all models belonging to $\mathbf{N}(k')$. The ratios r_{RJ2} are next computed.

There are several ways to combine these ratios with $\hat{\pi}(l | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)$ and $\hat{\pi}(s | \mathbf{D}_n) / \hat{\pi}(k' | \mathbf{D}_n)$ to improve the estimation of $\pi(l | \mathbf{D}_n) / \pi(k | \mathbf{D}_n)$ and $\pi(s | \mathbf{D}_n) / \pi(k' | \mathbf{D}_n)$. We define the improved version of the PMF g as follows to reflect this flexibility:

$$g_{\text{imp.}}(k, k', \mathbf{x}_{k \mapsto \bullet}^{(0:T-1, \bullet)}, \mathbf{y}_{\bullet}^{(0:T-1, \bullet)}) := h \left(\frac{\tilde{\pi}(k' | \mathbf{D}_n)}{\tilde{\pi}(k | \mathbf{D}_n)} \right) / c_{\text{imp.}}(k), \quad (12)$$

where

$$\frac{\tilde{\pi}(k' | \mathbf{D}_n)}{\tilde{\pi}(k | \mathbf{D}_n)} := \varrho \left(\frac{\hat{\pi}(k' | \mathbf{D}_n)}{\hat{\pi}(k | \mathbf{D}_n)}, r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1,1)})), \dots, r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1,N)})) \right),$$

$\mathbf{y}_{\bullet}^{(0:T-1, \bullet)}$ is the vector containing $\mathbf{y}_l^{(0,j)}, \dots, \mathbf{y}_l^{(T-1,j)}$ for all $j \in \{1, \dots, N\}$ and $l \in \mathbf{N}(k)$, and $c_{\text{imp.}}$ is the normalizing constant. ϱ is a function aiming at putting together the information; its choice is discussed below. $\mathbf{x}_{k \mapsto \bullet}^{(0:T-1, \bullet)}$ is the vector

containing $\mathbf{x}_{k \mapsto l}^{(0,j)}, \dots, \mathbf{x}_{k \mapsto l}^{(T-1,j)}$ for all $j \in \{1, \dots, N\}$ and $l \in \mathbf{N}(k)$. Here we added the subscript $k \mapsto l$ and superscript (t, j) to highlight that the sequence generated using Step 2.(b) of Algorithm 3 depends on j (different sequences are generated for different j), and also on l through $\rho_{k \mapsto l}^{(t)}$ (6). Note that $\tilde{\pi}(l | \mathbf{D}_n) / \tilde{\pi}(k | \mathbf{D}_n)$ is an estimator of $\pi(l | \mathbf{D}_n) / \pi(k | \mathbf{D}_n)$ which is in fact a function of $\mathbf{x}_{k \mapsto l}^{(0,j)}, \dots, \mathbf{x}_{k \mapsto l}^{(T-1,j)}$ and $\mathbf{y}_l^{(0,j)}, \dots, \mathbf{y}_l^{(T-1,j)}$ for all $j \in \{1, \dots, N\}$ additionally to k and l ; we used this notation to simplify and make the connection with $\hat{\pi}(l | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)$.

Algorithm 5 includes the idea of improving $\hat{\pi}(l | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)$ using ratios r_{RJ2} in a valid way (as indicated by Proposition 2 below). This algorithm is quite complicated, but once a code for Algorithm 4 has been written, we can use the connections between Algorithm 5 and Algorithm 4 to facilitate the coding of the former. Another drawback of Algorithm 5 is that it requires to perform the computations for $g_{\text{imp.}}(k', \cdot)$ even when $k' = k$. This is because $g_{\text{imp.}}(k, k', \mathbf{x}_{k \mapsto \bullet}^{(0:T-1, \bullet)}, \mathbf{y}_{\bullet}^{(0:T-1, \bullet)})$ is different from $g_{\text{imp.}}(k', k, \mathbf{y}_{k' \mapsto \bullet}^{(0:T-1, j^*)}, \mathbf{z}_{\bullet}^{(0:T-1, j^*)}, \mathbf{z}_{\bullet}^{(0:T-1, \bullet)})$ (see Step 2.(i)), even when $k' = k$. At least, most of computations for the two main steps (Steps 2.(i) and 2.(ii)) can be performed in parallel. Again, computational overheads have to be taken into account. In our numerical example, the improvement is not significant enough to offset the computational cost.

Proposition 2. *Under the two assumptions presented in Section A.1, (10)–(11), Algorithm 5 is valid.*

It is natural to set $\tilde{\pi}(l | \mathbf{D}_n) / \tilde{\pi}(k | \mathbf{D}_n)$ to 1 when $l = k$. This implies that we in fact do not need to sample proposals for Model k in the first parts of Steps 2.(i) and 2.(ii). Also, in the second parts of Steps 2.(i) and 2.(ii), it is not required to sample proposals for $s = k'$ for the same reason. When $k' = k$, we actually perform a “vanilla” parameter-update step (an HMC step in our numerical example). In this case, the MH ratio as to be multiplied by

$$\frac{g_{\text{imp.}}(k', k, \mathbf{y}_{k' \mapsto \bullet}^{(0:T-1, j^*)}, \mathbf{z}_{\bullet}^{(0:T-1, j^*)}, \mathbf{z}_{\bullet}^{(0:T-1, \bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_{k \mapsto \bullet}^{(0:T-1, \bullet)}, \mathbf{y}_{\bullet}^{(0:T-1, \bullet)})}$$

or

$$\frac{g_{\text{imp.}}(k', k, \mathbf{y}_{k' \mapsto \bullet}^{(0:T-1, 1)}, \mathbf{z}_{\bullet}^{(0:T-1, \bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_{k \mapsto \bullet}^{(0:T-1, 1)}, \mathbf{y}_{\bullet}^{(0:T-1, 1)}, \mathbf{y}_{\bullet \mapsto k}^{(0:T-1, \bullet)})},$$

depending if Step 2.(i) or Step 2.(ii) is used.

The function ϱ in (12) specifies the way the information is combined. It may be set for instance to the simple average:

$$\frac{\tilde{\pi}(l | \mathbf{D}_n)}{\tilde{\pi}(k | \mathbf{D}_n)} = \frac{1}{N + 1} \left(\frac{\tilde{\pi}(l | \mathbf{D}_n)}{\tilde{\pi}(k | \mathbf{D}_n)} + \sum_{j=1}^N r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (l, \mathbf{y}_l^{(T-1, j)})) \right). \quad (13)$$

Algorithm 5 RJ additionally improving the model proposal distribution

1. Sample $u_a, u_c \sim \mathcal{U}(0, 1)$. If $u_c \leq 1/2$ go to Step 2.(i), otherwise go to Step 2.(ii).
- 2.(i) Do:
 - (a) For all $l \in \mathbf{N}(k)$, sample proposals as in Step 2.(b-i) of [Algorithm 4](#): from \mathbf{x}_k , sample N proposals $\mathbf{y}_l^{(T-1,1)}, \dots, \mathbf{y}_l^{(T-1,N)}$ as in Step 2.(b) of [Algorithm 3](#).
 - (b) Compute $\tilde{\pi}(l | \mathbf{D}_n) / \tilde{\pi}(k | \mathbf{D}_n)$ for all $l \in \mathbf{N}(k)$ and $g_{\text{imp.}}(k, \cdot)$ as in (12).
 - (c) Sample $k' \sim g_{\text{imp.}}(k, \cdot)$ and j^* from a PMF such that $\mathbb{P}(J^* = j) \propto r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (k', \mathbf{y}_{k'}^{(T-1,j)}))$, and compute $\bar{r}(k, k')$.
 - (d) For all $s \in \mathbf{N}(k') \setminus \{k\}$, sample proposals as in Step 2.(b-ii) of [Algorithm 4](#): from $\mathbf{y}_{k'}^{(T-1,j^*)}$, sample one path as in Step 2.(b) of [Algorithm 3](#). Denote the endpoint by $\mathbf{z}_s^{(0,j^*)}$. From $\mathbf{z}_s^{(0,j^*)}$, sample $N - 1$ reverse paths again as in Step 2.(b) of [Algorithm 3](#), yielding $N - 1$ proposals for the parameters of Model k' that we denote by $\mathbf{z}_{s \rightarrow k'}^{(T-1,j)}, j \in \{1, \dots, N\} \setminus \{j^*\}$.
 - (e) Compute $\tilde{\pi}(s | \mathbf{D}_n) / \tilde{\pi}(k' | \mathbf{D}_n)$ for all $s \in \mathbf{N}(k') \setminus \{k\}$. Compute $g_{\text{imp.}}(k', \cdot)$ using these and $(\tilde{\pi}(k' | \mathbf{D}_n) / \tilde{\pi}(k | \mathbf{D}_n))^{-1}$ computed in (b).
 - (f) If

$$u_a \leq \frac{g_{\text{imp.}}(k', k, \mathbf{y}_{k' \rightarrow \bullet}^{(0:T-1,j^*)}, \mathbf{z}_{\bullet}^{(0:T-1,j^*)}, \mathbf{z}_{\bullet}^{(0:T-1,\bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_{k \rightarrow \bullet}^{(0:T-1,\bullet)}, \mathbf{y}_{\bullet}^{(0:T-1,\bullet)})} \bar{r}(k, k'),$$

set the next state of the chain to $(k', \mathbf{y}_{k'}^{(T-1,j^*)})$. Otherwise, set it to (k, \mathbf{x}_k) .

2.(ii) Do:

- (a) For all $l \in \mathbf{N}(k)$, sample proposals as in Step 2.(b-ii) of [Algorithm 4](#): from \mathbf{x}_k , sample one path as in Step 2.(b) of [Algorithm 3](#). Denote the endpoint by $\mathbf{y}_l^{(T-1,1)}$. From $\mathbf{y}_l^{(T-1,1)}$, sample $N - 1$ reverse paths again as in Step 2.(b) of [Algorithm 3](#), yielding $N - 1$ proposals for the parameters of Model k that we denote by $\mathbf{y}_{l \rightarrow k}^{(0,2)}, \dots, \mathbf{y}_{l \rightarrow k}^{(0,N)}$.
- (b) Compute $\tilde{\pi}(l | \mathbf{D}_n) / \tilde{\pi}(k | \mathbf{D}_n)$ for all $l \in \mathbf{N}(k)$ and $g_{\text{imp.}}(k, \cdot)$ as in (12).
- (c) Sample $k' \sim g_{\text{imp.}}(k, \cdot)$ and compute $\bar{r}(k', k)^{-1}$.
- (d) For all $s \in \mathbf{N}(k') \setminus \{k\}$, sample proposals as in Step 2.(b-i) of [Algorithm 4](#): from $\mathbf{y}_{k'}^{(T-1,1)}$, sample N proposals $\mathbf{z}_s^{(0,1)}, \dots, \mathbf{z}_s^{(0,N)}$ as in Step 2.(b) of [Algorithm 3](#).
- (e) Compute $\tilde{\pi}(s | \mathbf{D}_n) / \tilde{\pi}(k' | \mathbf{D}_n)$ for all $s \in \mathbf{N}(k') \setminus \{k\}$. Compute $g_{\text{imp.}}(k', \cdot)$ using these and $(\tilde{\pi}(k' | \mathbf{D}_n) / \tilde{\pi}(k | \mathbf{D}_n))^{-1}$ computed in (b).
- (f) If

$$u_a \leq \frac{g_{\text{imp.}}(k', k, \mathbf{y}_{k' \rightarrow \bullet}^{(0:T-1,1)}, \mathbf{z}_{\bullet}^{(0:T-1,\bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_{k \rightarrow \bullet}^{(0:T-1,1)}, \mathbf{y}_{\bullet}^{(0:T-1,1)}, \mathbf{y}_{\bullet \rightarrow k}^{(0:T-1,\bullet)})} \bar{r}(k', k)^{-1},$$

set the next state of the chain to $(k', \mathbf{y}_{k'}^{(T-1,1)})$. Otherwise, set it to (k, \mathbf{x}_k) .

3. Go to Step 1.

One may alternatively take the average of $\hat{\pi}(l | \mathbf{D}_n) / \hat{\pi}(k | \mathbf{D}_n)$ and $\bar{r}(k, l)$:

$$\frac{\tilde{\pi}(l | \mathbf{D}_n)}{\tilde{\pi}(k | \mathbf{D}_n)} = \frac{1}{2} \left(\frac{\hat{\pi}(l | \mathbf{D}_n)}{\hat{\pi}(k | \mathbf{D}_n)} + \bar{r}(k, l) \right)$$

$$= \frac{1}{2} \left(\frac{\widehat{\pi}(l \mid \mathbf{D}_n)}{\widehat{\pi}(k \mid \mathbf{D}_n)} + \frac{1}{N} \sum_{j=1}^N r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (l, \mathbf{y}_l^{(T-1,j)})) \right).$$

These reflect a choice of putting more or less weight on $\widehat{\pi}(l \mid \mathbf{D}_n)/\widehat{\pi}(k \mid \mathbf{D}_n)$. We know that if T and N are large enough then $\bar{r}(k, l)$ is close to $\pi(l \mid \mathbf{D}_n)/\pi(k \mid \mathbf{D}_n)$, which may not be the case for $\widehat{\pi}(l \mid \mathbf{D}_n)/\widehat{\pi}(k \mid \mathbf{D}_n)$ when n is not sufficiently large. The latter ratio may thus act as outlying/conflicting information against which these averages above are not robust. A robust approach consists in setting ϱ to be the median of $\widehat{\pi}(l \mid \mathbf{D}_n)/\widehat{\pi}(k \mid \mathbf{D}_n), r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (l, \mathbf{y}_l^{(T-1,1)})), \dots, r_{\text{RJ2}}((k, \mathbf{x}_k^{(0)}), (l, \mathbf{y}_l^{(T-1,N)}))$. We recommend this approach and use it in our numerical example.

Furthermore, as $T, N \rightarrow \infty$, $\tilde{\pi}(l \mid \mathbf{D}_n)/\tilde{\pi}(k \mid \mathbf{D}_n)$ is a consistent estimator of $\pi(l \mid \mathbf{D}_n)/\pi(k \mid \mathbf{D}_n)$ for fixed n when the median or (13) is used (recall the properties of r_{RJ2} and \bar{r} mentioned in Sections A.1 and A.2). Therefore, if the function h is such that $h(x) = x h(1/x)$ for $x > 0$, then the acceptance probabilities in Algorithm 5 converge towards $1 \wedge c(k)/c(k')$, where $c(k)$ and $c(k')$ are the normalizing constants with

$$c(k) = \sum_{l \in \mathbf{N}(k)} h \left(\frac{\pi(l \mid \mathbf{D}_n)}{\pi(k \mid \mathbf{D}_n)} \right).$$

In fact, the same technique as in the proof of Theorem 1 in Gagnon and Doucet (2021) allows to prove that a Markov chain simulated by Algorithm 5 converges weakly for fixed n to that of an ideal RJ which has access to the unnormalized version of the posterior probabilities $\pi(k \mid \mathbf{D}_n)$ and is able to sample from the conditional distributions $\pi(\cdot \mid k, \mathbf{D}_n)$ (and for which the acceptance probabilities are $1 \wedge c(k)/c(k')$), with its good mixing properties as discussed in Section 1.3.

Appendix B: Details of the variable-selection example

We review linear regression and introduce notation in Section B.1. We present in Section B.2 all the details to compute estimates for the normal linear-regression model. Some are useful for the computation of algorithm inputs, such as the observed information matrix, and to verify the validity of our RJ code (as mentioned in Section 4.1). In Section B.3, we turn to the details of the robust linear regressions and the computation of algorithm inputs.

B.1. Linear regression

We first introduce notation. We define $\gamma_1, \dots, \gamma_n \in \mathbb{R}$ to be n data points from the dependent variable, and $\boldsymbol{\gamma}_n := (\gamma_1, \dots, \gamma_n)^T$. We denote the full design matrix containing n observations from all covariates by $\mathbf{C} \in \mathbb{R}^{n \times p}$, where p is a positive integer. For simplicity, we refer to the first column of \mathbf{C} as the *first* covariate even if, as usual, it is a column of 1's. The design matrix associated

with Model k whose columns form a subset of \mathbf{C} is denoted by \mathbf{C}_k , with lines denoted by $\mathbf{c}_{i,k}^T$. We use d_k to denote the number of covariates in Model k ; we therefore slightly abuse notation given that the number of parameters for Model k is $d_k + 1$ (one regression coefficient per covariate plus the scale parameter of the error term).

The linear regression is

$$\gamma_i = \mathbf{c}_{i,K}^T \boldsymbol{\beta}_K + \epsilon_{i,K}, \quad i = 1, \dots, n, \quad K \in \mathcal{K},$$

where K is the model indicator, $\boldsymbol{\beta}_K$ is the random vector containing the regression coefficients of Model K and $\epsilon_{1,K}, \dots, \epsilon_{n,K} \in \mathbb{R}$ are the random errors of Model K . We assume that $\epsilon_{1,K}, \dots, \epsilon_{n,K}$ and $\boldsymbol{\beta}_K$ are $n + 1$ conditionally independent random variables given (K, σ_K) , with $\sigma_K > 0$ being the scale parameter of the errors of Model K . The conditional density of $\epsilon_{i,K}$ is given by

$$\epsilon_{i,K} | K, \sigma_K, \boldsymbol{\beta}_K \stackrel{d}{=} \epsilon_{i,K} | K, \sigma_K \stackrel{d}{\sim} (1/\sigma_K) f(\epsilon_{i,K}/\sigma_K), \quad i = 1, \dots, n.$$

B.2. Normal linear regression

We present in this section a result giving the precise form of the joint posterior density for the normal linear regression.

Proposition 3. *If $f = \varphi(\cdot; 0, 1)$ and $\pi(\boldsymbol{\beta}_k, \sigma_k | k) \propto 1/\sigma_k$, then*

$$\pi(k | \boldsymbol{\gamma}_n) \propto \pi(k) \frac{\Gamma((n - d_k)/2) \pi^{d_k/2}}{\|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^{n-d_k} |\mathbf{C}_k^T \mathbf{C}_k|^{1/2}}, \quad (14)$$

$$\pi(\sigma_k | k, \boldsymbol{\gamma}_n) = \frac{2^{1 - \frac{n-d_k}{2}} \|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^{n-d_k}}{\Gamma((n - d_k)/2) \sigma_k^{n-d_k+1}} \exp \left\{ -\frac{1}{2\sigma_k^2} \|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^2 \right\},$$

and

$$\boldsymbol{\beta}_K | K, \sigma_K, \boldsymbol{\gamma}_n \sim \varphi(\cdot; \hat{\boldsymbol{\beta}}_k, \sigma_K^2 (\mathbf{C}_K^T \mathbf{C}_K)^{-1}),$$

where $\hat{\boldsymbol{\gamma}}_k := \mathbf{C}_k (\mathbf{C}_k^T \mathbf{C}_k)^{-1} \mathbf{C}_k^T \boldsymbol{\gamma}_n$, $\hat{\boldsymbol{\beta}}_k := (\mathbf{C}_K^T \mathbf{C}_K)^{-1} \mathbf{C}_K^T \boldsymbol{\gamma}_n$, and $\|\cdot\|_2$ is the Euclidian norm. Note that the normalization constant of $\pi(k | \boldsymbol{\gamma}_n)$ is the sum over k of the expression on the right-hand side (RHS) in (14).

Note that $\sigma_K^2 | K, \boldsymbol{\gamma}_n$ has an inverse-gamma distribution with shape and scale parameters given by $(n - d_k)/2$ and $\|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^2/2$, respectively. We work on the log-scale for the scale parameters so that all the parameters take values on the real line, and presumably, have densities that are closer to the bell curve. We thus define $\eta_k := \log \sigma_k$. The associated conditional distribution is given by

$$\pi(\eta_k | k, \boldsymbol{\gamma}_n) := \frac{2^{1 - \frac{n-d_k}{2}} \|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^{n-d_k}}{\Gamma((n - d_k)/2) e^{(n-d_k)\eta_k}} \exp \left\{ -\frac{1}{2e^{2\eta_k}} \|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^2 \right\}.$$

The other terms in the joint posterior do not change except that we now consider that $\boldsymbol{\beta}_K | K, \eta_K, \boldsymbol{\gamma}_n \sim \varphi(\cdot; \hat{\boldsymbol{\beta}}_k, \exp(2\eta_K) (\mathbf{C}_K^T \mathbf{C}_K)^{-1})$.

The maximizers of the conditional posterior densities given that $K = k$ are:

$$\hat{\mathbf{x}}_k = (\hat{\boldsymbol{\beta}}_k, \hat{\eta}_k),$$

where

$$\hat{\eta}_k := \log \sqrt{\frac{1}{n} \|\boldsymbol{\gamma}_n - \hat{\boldsymbol{\gamma}}_k\|_2^2}.$$

The observed information matrix evaluated at the maximizers is thus given by

$$\hat{\mathcal{I}}_k = \begin{pmatrix} \mathbf{C}_k^T \mathbf{C}_k / e^{2\hat{\eta}_k} & \mathbf{0} \\ \mathbf{0} & 2n \end{pmatrix}.$$

Then,

$$\hat{\mathcal{I}}_k^{-1} = \begin{pmatrix} e^{2\hat{\eta}_k} (\mathbf{C}_k^T \mathbf{C}_k)^{-1} & \mathbf{0} \\ \mathbf{0} & (2n)^{-1} \end{pmatrix}. \tag{15}$$

Relying on improper priors such as $\pi(\boldsymbol{\beta}_k, \sigma_k | k) \propto 1/\sigma_k$ may lead to inconsistencies in model selection (see, e.g., Casella et al. (2009)). When this problem happens, the phenomenon is referred to as the *Jeffreys-Lindley paradox* (Lindley (1957) and Jeffreys (1967)) in the literature. It can be shown that the Jeffreys-Lindley paradox does not arise under the framework of normal linear regression described above when the prior distribution on K is carefully chosen, and more precisely, set to

$$\pi(k) \propto |\mathbf{C}_k^T \mathbf{C}_k|^{1/2} / n^{d_k/2}.$$

See Gagnon, Bédard and Desgagné (2021) for an analogous proof in a framework of principal component regression.

When the covariates are orthonormal, $|\mathbf{C}_k^T \mathbf{C}_k|^{1/2} = |n\mathbf{I}_{d_k}|^{1/2} = n^{d_k/2}$ (if all covariate observations are standardized using a standard deviation in which the divisor is n). In this case, $\pi(k) \propto |\mathbf{C}_k^T \mathbf{C}_k|^{1/2} / n^{d_k/2} = 1$. The proposed prior on K can thus be seen as a relative adjustment of the volume spanned by the columns of $\mathbf{C}_K^T \mathbf{C}_K$.

B.3. Robust linear regression

In this section, we consider that the density f is a LPTN with parameter $\rho \in (2\Phi(1) - 1, 1) \approx (0.6827, 1)$, i.e.

$$f(x) = \begin{cases} \varphi(x; 0, 1) & \text{if } |x| \leq \tau, \\ \varphi(\tau; 0, 1) \frac{\tau}{|x|} \left(\frac{\log \tau}{\log |x|}\right)^{\lambda+1} & \text{if } |x| > \tau, \end{cases} \tag{16}$$

where $x \in \mathbb{R}$ and $\Phi(\cdot)$ is the cumulative distribution function (CDF) of a standard normal. The terms $\tau > 1$ and $\lambda > 0$ are functions of ρ and satisfy

$$\begin{aligned} \tau &= \{\tau : \mathbb{P}(-\tau \leq Z \leq \tau) = \rho \text{ for } Z \sim \varphi(\cdot; 0, 1)\}, \\ \lambda &= 2(1 - \rho)^{-1} \varphi(\tau; 0, 1) \tau \log(\tau). \end{aligned}$$

Setting ρ to 0.95 has proved to be suitable for practical purposes (see Gagnon, Desgagné and Bédard (2020a)). Accordingly, this is the value that is used in our numerical example.

Using the same prior as the normal regression, the joint posterior density is:

$$\pi(k, \boldsymbol{\beta}_k, \sigma_k \mid \boldsymbol{\gamma}_n) \propto \frac{|\mathbf{C}_k^T \mathbf{C}_k|^{1/2}}{n^{d_k/2}} \frac{1}{\sigma_k} \prod_{i=1}^n \frac{1}{\sigma_k} f\left(\frac{\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k}{\sigma_k}\right).$$

With the change of variable $\eta_k = \log \sigma_k$, we have

$$\pi(k, \boldsymbol{\beta}_k, \eta_k \mid \boldsymbol{\gamma}_n) \propto \frac{|\mathbf{C}_k^T \mathbf{C}_k|^{1/2}}{n^{d_k/2}} \frac{1}{e^{\eta_k n}} \prod_{i=1}^n f\left(\frac{\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k}{e^{\eta_k}}\right).$$

Recall that for the robust models we set $\hat{\mathcal{L}}_k^{-1}$ as in (15) but use the maximizer $\hat{\eta}_k$ of the function above. We need log conditionals and their gradients for optimizers and HMC:

$$\log \pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n) \propto -n\eta_k + \sum_{i=1}^n \log f\left(\frac{\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k}{e^{\eta_k}}\right),$$

where the proportional sign has to be understood in the original scale, and

$$\begin{aligned} \frac{\partial}{\partial \boldsymbol{\beta}_k} \log \pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n) &= \sum_{i=1}^n e^{-2\eta_k} (\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) \mathbf{c}_{i,k} \mathbf{1} \left((\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) / e^{\eta_k} \leq \tau \right) \\ &+ \left[\frac{\operatorname{sgn}(\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) \mathbf{c}_{i,k}}{|\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k|} + (\lambda + 1) \frac{\operatorname{sgn}(\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) \mathbf{c}_{i,k}}{|\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k| \log \left((\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) / e^{\eta_k} \right)} \right] \\ &\times \mathbf{1} \left((\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) / e^{\eta_k} > \tau \right), \end{aligned}$$

and

$$\begin{aligned} \frac{\partial}{\partial \eta_k} \log \pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n) &= -n + \sum_{i=1}^n e^{-2\eta_k} (\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k)^2 \mathbf{1} \left((\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) / e^{\eta_k} \leq \tau \right) \\ &+ \left[1 + (\lambda + 1) \frac{1}{\log \left((\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) / e^{\eta_k} \right)} \right] \mathbf{1} \left((\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k) / e^{\eta_k} > \tau \right), \end{aligned}$$

$\operatorname{sgn}(\cdot)$ being the sign function.

To use the annealing distributions in Algorithms 3, 4 and 5, we work with the log densities; therefore we simply multiply $\log \pi(\cdot \mid k, \boldsymbol{\gamma}_n)$ by $1 - t/T$ and $\log \pi(\cdot \mid k', \boldsymbol{\gamma}_n)$ by t/T to obtain $\log \rho_{k \rightarrow k'}^{(t)}$. To use MALA proposal distributions, we however need to compute the gradient of $\log \rho_{k \rightarrow k'}^{(t)}$. We now compute the

gradient with respect to $\mathbf{x}_k^{(t)}$; the gradient with respect to $\mathbf{y}_{k'}^{(t)}$ is computed similarly. The proportional sign “ \propto ” is with respect to everything that are not the parameters:

$$\begin{aligned} & \pi(\mathbf{x}_k^{(t)} \mid k, \mathbf{D}_n)^{1-t/T} q_{k' \rightarrow k}(\mathbf{x}_k^{(t)})^{t/T} \\ &= [\pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n)]^{1-t/T} \\ & \quad \times \left[\frac{|\mathbf{C}_k^T \mathbf{C}_k|^{1/2}}{(2\pi)^{d_k/2} e^{d_k \hat{\eta}_k}} \exp \left(-\frac{1}{2e^{2\hat{\eta}_k}} (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k)^T (\mathbf{C}_k^T \mathbf{C}_k) (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k) \right) \right]^{t/T} \\ & \quad \times \left[\frac{1}{\sqrt{2\pi(1/(2n))}} \exp \left(-\frac{1}{2(1/(2n))} (\eta_k - \hat{\eta}_k)^2 \right) \right]^{t/T} \\ & \propto [\pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n)]^{1-t/T} \\ & \quad \times \frac{1}{e^{d_k(t/T)\hat{\eta}_k}} \exp \left(-\frac{(t/T)}{2e^{2\hat{\eta}_k}} (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k)^T (\mathbf{C}_k^T \mathbf{C}_k) (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k) \right) \\ & \quad \times \exp \left(-n(t/T)(\eta_k - \hat{\eta}_k)^2 \right), \end{aligned}$$

where we omitted the superscript “(t)” for the variables to simplify. Therefore,

$$\begin{aligned} & \frac{\partial}{\partial \boldsymbol{\beta}_k} \log \pi(\mathbf{x}_k^{(t)} \mid k, \mathbf{D}_n)^{1-t/T} q_{k' \rightarrow k}(\mathbf{x}_k^{(t)})^{t/T} \\ &= (1-t/T) \frac{\partial}{\partial \boldsymbol{\beta}_k} \log \pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n) - (t/T) e^{-2\hat{\eta}_k} (\mathbf{C}_k^T \mathbf{C}_k) (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k), \end{aligned}$$

and

$$\begin{aligned} & \frac{\partial}{\partial \eta_k} \log \pi(\mathbf{x}_k^{(t)} \mid k, \mathbf{D}_n)^{1-t/T} q_{k' \rightarrow k}(\mathbf{x}_k^{(t)})^{t/T} \\ &= (1-t/T) \frac{\partial}{\partial \eta_k} \log \pi(\boldsymbol{\beta}_k, \eta_k \mid k, \boldsymbol{\gamma}_n) - 2n(t/T)(\eta_k - \hat{\eta}_k). \end{aligned}$$

Appendix C: Proofs

We present the proofs of Theorem 1, Proposition 1, Proposition 2 and Proposition 3 in that order.

Proof of Theorem 1. We here only prove that $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\} \implies \{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ in probability as the proof of $\{(K, \mathbf{Z}_K)_{\text{ideal}}(m) : m \in \mathbb{N}\} \implies \{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$ is similar. To prove this result, we use Theorem 2 of Schmon et al. (2021). We thus have to verify the following three conditions.

1. $(K, \mathbf{Z}_K)_n(0) \implies (K, \mathbf{Z}_K)_{\text{limit}}(0)$ in probability as $n \rightarrow \infty$.
2. Use P_n and P_{limit} to denote the transition kernels of $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$ and $\{(K, \mathbf{Z}_K)_{\text{limit}}(m) : m \in \mathbb{N}\}$, respectively. These are such that

$$\sum_k \int |P_n \phi(k, \mathbf{z}_k) - P_{\text{limit}} \phi(k, \mathbf{z}_k)| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{z}_k \rightarrow 0, \quad (17)$$

in probability as $n \rightarrow \infty$ for all $\phi \in \text{BL}$, where $\pi_{K, \mathbf{Z}_K}(\cdot, \cdot | \mathbf{D}_n)$ denotes the stationary distribution/density of $\{(K, \mathbf{Z}_K)_n(m) : m \in \mathbb{N}\}$ and BL denotes the set of bounded Lipschitz functions.

3. The transition kernel P_{limit} is such that $P_{\text{limit}}\phi$ is continuous for any $\phi \in \mathcal{C}_b$ (the set of continuous bounded functions).

We start with Step 1. Given that \mathcal{K} is finite (Assumption 1), it suffices to verify that

$$|\pi(k | \mathbf{D}_n) \mathbb{P}(\mathbf{Z}_{k,n} \in A) - \bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,\text{limit}} \in A)| \rightarrow 0 \quad \text{in probability,}$$

for any k and measurable set A , where $\mathbb{P}(\mathbf{Z}_{k,n} \in A)$ and $\mathbb{P}(\mathbf{Z}_{k,\text{limit}} \in A)$ are computed using the conditional distributions given that $K = k$. Using the triangle inequality, we have that

$$\begin{aligned} & |\pi(k | \mathbf{D}_n) \mathbb{P}(\mathbf{Z}_{k,n} \in A) - \bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,\text{limit}} \in A)| \\ & \leq |\pi(k | \mathbf{D}_n) \mathbb{P}(\mathbf{Z}_{k,n} \in A) - \bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,n} \in A)| \\ & \quad + |\bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,n} \in A) - \bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,\text{limit}} \in A)|. \end{aligned}$$

We now show that both absolute values converge towards 0 in probability which will allow to conclude by Slutsky's theorem and monotonicity of probabilities. We first have that

$$|\pi(k | \mathbf{D}_n) \mathbb{P}(\mathbf{Z}_{k,n} \in A) - \bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,n} \in A)| \leq |\pi(k | \mathbf{D}_n) - \bar{\pi}(k)| \rightarrow 0,$$

in probability by Assumption 2 and the fact that $\mathbb{P}(\mathbf{Z}_{k,n} \in A) \leq 1$. Using now that $\bar{\pi}(k) \leq 1$, we have that

$$\begin{aligned} & |\bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,n} \in A) - \bar{\pi}(k) \mathbb{P}(\mathbf{Z}_{k,\text{limit}} \in A)| \\ & \leq |\mathbb{P}(\mathbf{Z}_{k,n} \in A) - \mathbb{P}(\mathbf{Z}_{k,\text{limit}} \in A)| \\ & = |\mathbb{P}(\mathbf{X}_{k,n} \in A_n) - \mathbb{P}(\mathbf{X}_{k,\text{limit}} \in A_n)| \\ & \leq \int |\pi(\mathbf{x}_k | k, \mathbf{D}_n) - \varphi(\mathbf{x}_k; \hat{\mathbf{x}}_k, \boldsymbol{\Sigma}_k/n)| d\mathbf{x}_k \rightarrow 0, \end{aligned}$$

in probability, by Jensen's inequality and Assumption 3, where A_n is the set A after applying the inverse transformation to retrieve the original random variables. Note that in the last inequality, we also used that $A_n \subset \mathbb{R}^{d_k}$.

We continue with Step 2 and prove (17). We have that

$$\begin{aligned} P_{\text{limit}}((k, \mathbf{z}_k), (k', d\mathbf{y}_{k'})) &= g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(d\mathbf{y}_{k'}; \mathbf{0}, \boldsymbol{\Sigma}_{k'}) \\ & \quad + \delta_{(k, \mathbf{z}_k)}(k', d\mathbf{y}_{k'}) \sum_{l \in \mathbf{N}(k)} g_{\text{limit}}(k, l) (1 - \alpha_{\text{limit}}(k, l)), \end{aligned}$$

where

$$g_{\text{limit}}(k, k') := \frac{h\left(\frac{\bar{\pi}(k')}{\bar{\pi}(k)}\right)}{\sum_{l \in \mathbf{N}(k)} h\left(\frac{\bar{\pi}(l)}{\bar{\pi}(k)}\right)},$$

and

$$\alpha_{\text{limit}}(k, k') := 1 \wedge \frac{\bar{\pi}(k') g_{\text{limit}}(k', k)}{\bar{\pi}(k) g_{\text{limit}}(k, k')}.$$

By definition,

$$\begin{aligned} P_{\text{limit}}\phi(k, \mathbf{z}_k) &= \sum_{k' \in \mathbf{N}(k)} \int \phi(k', \mathbf{y}_{k'}) P_{\text{limit}}((k, \mathbf{z}_k), (k', d\mathbf{y}_{k'})) \\ &= \sum_{k' \in \mathbf{N}(k)} \int \phi(k', \mathbf{y}_{k'}) g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \boldsymbol{\Sigma}_{k'}) d\mathbf{y}_{k'} \\ &\quad + \phi(k, \mathbf{z}_k) \sum_{l \in \mathbf{N}(k)} g_{\text{limit}}(k, l) (1 - \alpha_{\text{limit}}(k, l)). \end{aligned}$$

$P_{\text{limit}}\phi$ is thus continuous for any $\phi \in \mathcal{C}_b$ so Condition 3 above is satisfied.

We also have that

$$\begin{aligned} P_n((k, \mathbf{z}_k), (k', d\mathbf{y}_{k'})) &= g(k, k') \varphi(d\mathbf{y}_{k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) \\ &\quad + \delta_{(k, \mathbf{z}_k)}(k', d\mathbf{y}_{k'}) \\ &\quad \times \sum_{l \in \mathbf{N}(k)} \int (1 - \alpha((k, \mathbf{z}_k), (l, \mathbf{u}_{k \mapsto k'}))) g(k, l) \varphi(\mathbf{u}_{k \mapsto k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'}) d\mathbf{u}_{k \mapsto k'}, \end{aligned}$$

where in this case

$$\alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) = 1 \wedge \frac{\pi_{K, \mathbf{Z}_K}(k', \mathbf{y}_{k'} \mid \mathbf{D}_n) g(k', k) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_k)}{\pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'})}.$$

Therefore,

$$\begin{aligned} P_n\phi(k, \mathbf{z}_k) &= \sum_{k'} \int \phi(k', \mathbf{y}_{k'}) P_n((k, \mathbf{z}_k), (k', d\mathbf{y}_{k'})) \\ &= \sum_{k' \in \mathbf{N}(k)} \int \phi(k', \mathbf{y}_{k'}) g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) d\mathbf{y}_{k'} \\ &\quad + \phi(k, \mathbf{z}_k) \sum_{l \in \mathbf{N}(k)} \int (1 - \alpha((k, \mathbf{z}_k), (l, \mathbf{u}_{k \mapsto k'}))) g(k, l) \varphi(\mathbf{u}_{k \mapsto k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'}) d\mathbf{u}_{k \mapsto k'}. \end{aligned}$$

Consequently,

$$\begin{aligned} &\sum_k \int |P_n\phi(k, \mathbf{z}_k) - P_{\text{limit}}\phi(k, \mathbf{z}_k)| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{z}_k \\ &\leq \sum_k \int \left| \sum_{k' \in \mathbf{N}(k)} \int (\phi(k', \mathbf{y}_{k'}) g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) \right. \\ &\quad \left. - \phi(k', \mathbf{y}_{k'}) g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \boldsymbol{\Sigma}_{k'}) \right) d\mathbf{y}_{k'} \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{z}_k \end{aligned}$$

$$\begin{aligned}
& + \sum_k \int \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \\
& \times \left| \phi(k, \mathbf{z}_k) \sum_{l \in \mathbf{N}(k)} \int (1 - \alpha((k, \mathbf{z}_k), (l, \mathbf{u}_{k \rightarrow k'}))) g(k, l) \varphi(\mathbf{u}_{k \rightarrow k'}; \mathbf{0}, \hat{\Sigma}_{k'}) d\mathbf{u}_{k \rightarrow k'} \right. \\
& \quad \left. - \phi(k, \mathbf{z}_k) \sum_{l \in \mathbf{N}(k)} g_{\text{limit}}(k, l) (1 - \alpha_{\text{limit}}(k, l)) \right| d\mathbf{z}_k, \tag{18}
\end{aligned}$$

using the triangle inequality.

We now show that first term on the RHS in (18) converges towards 0 in probability. The second term converges towards 0 in probability following similar arguments. We will thus be able to conclude by Slutsky's theorem and monotonicity of probabilities.

Firstly,

$$\begin{aligned}
& \sum_k \int \left| \sum_{k' \in \mathbf{N}(k)} \int (\phi(k', \mathbf{y}_{k'}) g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k', n}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'}))) \right. \\
& \quad \left. - \phi(k', \mathbf{y}_{k'}) g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'}) \right| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{z}_k \\
& \leq M \sum_k \sum_{k' \in \mathbf{N}(k)} \int \left| g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k', n}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) \right. \\
& \quad \left. - g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'}) \right| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{y}_{k'} d\mathbf{z}_k,
\end{aligned}$$

because we can assume that $|\phi| \leq M$ with M a positive constant and using Jensen's inequality. Because the sums are on a finite number of terms, it suffices to prove that the integral converges to 0, for any k, k' . Note that we have not properly defined $g_{\text{limit}}(k, k')$ when $\bar{\pi}(k) = 0$. For the proof, we can consider that these $g_{\text{limit}}(k, k')$ with $\bar{\pi}(k) = 0$ have any definition, because we can use that

$$\begin{aligned}
& \int \left| g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) \right. \\
& \quad \left. - g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'}) \right| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{y}_{k'} d\mathbf{z}_k \\
& \leq \int \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{y}_{k'} d\mathbf{z}_k \\
& \quad + \int \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'}) \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) d\mathbf{y}_{k'} d\mathbf{z}_k = 2\pi(k \mid \mathbf{D}_n) \rightarrow 0,
\end{aligned}$$

using the triangle inequality, that $0 \leq g, \alpha, g_{\text{limit}}, \alpha_{\text{limit}} \leq 1$, and that $\pi(k \mid \mathbf{D}_n) \rightarrow \bar{\pi}(k) = 0$.

We can thus restrict our attention to the case where $\bar{\pi}(k) > 0$. Note that for similar reasons, we can restrict our attention to the case where $g_{\text{limit}}(k, k') > 0$ (implying that $\bar{\pi}(k') > 0$). Indeed, consider that $\bar{\pi}(k) > 0$, then all terms involved in $g(k, k')$, i.e. $h(\hat{\pi}(l \mid \mathbf{D}_n) / \hat{\pi}(k \mid \mathbf{D}_n))$ with $l \in \mathbf{N}(k)$, have a well-defined

limit in probability given by $h(\bar{\pi}(l)/\bar{\pi}(k))$ by Slutsky's and continuous mapping theorems and Assumption 2. Given that $g(k, k')$ is a continuous function of a finite number of random variables all converging towards a constant in probability, $g(k, k') \rightarrow g_{\text{limit}}(k, k') = 0$ by Slutsky's theorem, and $g_{\text{limit}}(k, k') = 0$ when $\bar{\pi}(k') = 0$.

Using twice the triangle inequality,

$$\begin{aligned} & \int \left| g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) \right. \\ & \quad \left. - g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'}) \right| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \, d\mathbf{y}_{k'} \, d\mathbf{z}_k \\ & \leq \int g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \\ & \quad \times |\alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) - \alpha_{\text{limit}}(k, k')| \, d\mathbf{y}_{k'} \, d\mathbf{z}_k \\ & \quad + |g(k, k') - g_{\text{limit}}(k, k')| \alpha_{\text{limit}}(k, k') \\ & \quad \times \int \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \, d\mathbf{y}_{k'} \, d\mathbf{z}_k \\ & \quad + g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \\ & \quad \times \int |\varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'})| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \, d\mathbf{y}_{k'} \, d\mathbf{z}_k. \end{aligned} \tag{19}$$

It is easily seen that the second term on the RHS converges to 0 given that $g(k, k') \rightarrow g_{\text{limit}}(k, k')$. The third term can be bounded above as follows:

$$\begin{aligned} & g_{\text{limit}}(k, k') \alpha_{\text{limit}}(k, k') \\ & \quad \times \int |\varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'})| \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \, d\mathbf{y}_{k'} \, d\mathbf{z}_k \\ & \leq \int |\varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \Sigma_{k'})| \, d\mathbf{y}_{k'} \\ & \leq \left[\text{tr}(\Sigma_{k'}^{-1} \hat{\Sigma}_{k'} - \mathbb{I}_{k'}) - \log \det(\hat{\Sigma}_{k'} \Sigma_{k'}^{-1}) \right]^{1/2}, \end{aligned}$$

using that $0 \leq g_{\text{limit}}, \alpha_{\text{limit}}, \pi(k \mid \mathbf{D}_n) \leq 1$ in the first inequality and then Proposition 2.1 of Devroye, Mehrabian and Reddad (2018) in the second one, where $\mathbb{I}_{k'}$ is the identity matrix of size $d_{k'}$, and $\text{tr}(\cdot)$ and $\det(\cdot)$ are the trace and determinant operators, respectively. The upper bound converges to 0 under Assumption 3. For the first term in (19), it is less direct. Using that $1 \wedge x$ is a Lipschitz function with constant 1 for all $x \geq 0$ and the triangle inequality,

$$\begin{aligned} & \int g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \\ & \quad \times |\alpha((k, \mathbf{z}_k), (k', \mathbf{y}_{k'})) - \alpha_{\text{limit}}(k, k')| \, d\mathbf{y}_{k'} \, d\mathbf{z}_k \\ & \leq \int \left| \pi_{K, \mathbf{Z}_K}(k', \mathbf{y}_{k'} \mid \mathbf{D}_n) g(k', k) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) \right. \\ & \quad \left. - g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{K, \mathbf{Z}_K}(k, \mathbf{z}_k \mid \mathbf{D}_n) \right| \frac{\bar{\pi}(k') g_{\text{limit}}(k', k)}{\bar{\pi}(k) g_{\text{limit}}(k, k')} \, d\mathbf{y}_{k'} \, d\mathbf{z}_k \end{aligned}$$

$$\begin{aligned}
&= \int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \pi(k' \mid \mathbf{D}_n) g(k', k) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) \right. \\
&\quad \left. - g(k, k') \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) \pi(k \mid \mathbf{D}_n) \frac{\bar{\pi}(k') g_{\text{limit}}(k', k)}{\bar{\pi}(k) g_{\text{limit}}(k, k')} \right| d\mathbf{y}_{k'} d\mathbf{z}_k \\
&\leq \int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \pi(k' \mid \mathbf{D}_n) g(k', k) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) \right. \\
&\quad \left. - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) \bar{\pi}(k') g_{\text{limit}}(k', k) \right| d\mathbf{y}_{k'} d\mathbf{z}_k \\
&+ \left| 1 - \frac{\pi(k \mid \mathbf{D}_n) g(k, k')}{\bar{\pi}(k) g_{\text{limit}}(k, k')} \right| \bar{\pi}(k') g_{\text{limit}}(k', k) \\
&\quad \times \int \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) d\mathbf{y}_{k'} d\mathbf{z}_k.
\end{aligned}$$

The last term is easily seen to converge towards 0 by Slutsky's theorem.

We now analyse the other term. Using the triangle inequality,

$$\begin{aligned}
&\int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \pi(k' \mid \mathbf{D}_n) g(k', k) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) \right. \\
&\quad \left. - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) \bar{\pi}(k') g_{\text{limit}}(k', k) \right| d\mathbf{y}_{k'} d\mathbf{z}_k \\
&\leq \left| \pi(k' \mid \mathbf{D}_n) g(k', k) - \bar{\pi}(k') g_{\text{limit}}(k', k) \right| \\
&\quad \times \int \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) d\mathbf{y}_{k'} d\mathbf{z}_k \\
&\quad + \bar{\pi}(k') g_{\text{limit}}(k', k) \\
&\quad \times \int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) \right| d\mathbf{y}_{k'} d\mathbf{z}_k.
\end{aligned}$$

The first term is seen to converge towards 0 by Slutsky's theorem.

Using again the triangle inequality, we obtain the following bound on the other term:

$$\begin{aligned}
&\int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) \right| d\mathbf{y}_{k'} d\mathbf{z}_k \\
&\leq \int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) \right| d\mathbf{y}_{k'} d\mathbf{z}_k \\
&\quad + \int \left| \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \pi_{\mathbf{Z}_K}(\mathbf{z}_k \mid k, \mathbf{D}_n) \right| d\mathbf{y}_{k'} d\mathbf{z}_k.
\end{aligned}$$

We prove the convergence to 0 of the first term. The other convergence follows from a similar argument.

Using the triangle inequality,

$$\begin{aligned}
&\int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \varphi(\mathbf{z}_k; \mathbf{0}, \hat{\Sigma}_k) \right| d\mathbf{y}_{k'} d\mathbf{z}_k \\
&= \int \left| \pi_{\mathbf{Z}_K}(\mathbf{y}_{k'} \mid k', \mathbf{D}_n) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\Sigma}_{k'}) \right| d\mathbf{y}_{k'}
\end{aligned}$$

$$\begin{aligned} &\leq \int |\pi_{\mathbf{z}_k}(\mathbf{y}_{k'} | k', \mathbf{D}_n) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \boldsymbol{\Sigma}_{k'})| \, d\mathbf{y}_{k'} \\ &\quad + \int \left| \varphi(\mathbf{y}_{k'}; \mathbf{0}, \boldsymbol{\Sigma}_{k'}) - \varphi(\mathbf{y}_{k'}; \mathbf{0}, \hat{\boldsymbol{\Sigma}}_{k'}) \right| \, d\mathbf{y}_{k'}. \end{aligned}$$

The first term converges to 0 by Assumption 3 after a change of variables, and the second one converges to 0 as well by Proposition 2.1 of Devroye, Mehrabian and Reddad (2018) and Assumption 3 as seen above. This concludes the proof. \square

Proof of Proposition 1. Because $\bar{\pi}(k^*) = 1$, we only have to establish the inequality for any set $A = \{(k, \mathbf{z}_k) \in \{k^*\} \times B\}$. So $A \setminus \{(k^*, \mathbf{z}_{k^*})\}$ implies that a parameter update is proposed and accepted with a parameter proposal, denoted here by \mathbf{y}_{k^*} , in B . When the chain is in stationarity, $P_{\text{limit},1}$ always proposes to update the parameters. Therefore,

$$P_{\text{limit},1}((k^*, \mathbf{z}_{k^*}), A \setminus \{(k^*, \mathbf{z}_{k^*})\}) = \mathbb{P}(\mathbf{y}_{k^*} \in B \text{ is accepted}).$$

In contrast,

$$\begin{aligned} P_{\text{limit},2}((k^*, \mathbf{z}_{k^*}), A \setminus \{(k^*, \mathbf{z}_{k^*})\}) &= g(k^*, k^*) \mathbb{P}(\mathbf{y}_{k^*} \in B \text{ is accepted}) \\ &\leq \mathbb{P}(\mathbf{y}_{k^*} \in B \text{ is accepted}) \\ &= P_{\text{limit},1}((k^*, \mathbf{z}_{k^*}), A \setminus \{(k^*, \mathbf{z}_{k^*})\}). \end{aligned}$$

The fact that $\bar{\pi}(k^*) = 1$ implies the order on the asymptotic variances (Tierney, 1998). \square

Proof of Proposition 2. We prove the result for the case $T = 1$ (without annealing intermediate distributions), to simplify; the general case is proved similarly. When $T = 1$, $r_{\text{RJ}2} = r_{\text{RJ}}$ which is the ratio in (1). We prove that the probability to reach the state $\{k'\} \times \{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}$, from $\{k\} \times \{\mathbf{x}_k \in A_k\}$, is equal to the probability of the reverse move. We denote by P the Markov kernel. We thus prove that

$$\begin{aligned} &\int_{\{\mathbf{x}_k \in A_k\}} \pi(k, \mathbf{x}_k | \mathbf{D}_n) \int_{\{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}} P((k, \mathbf{x}_k), (k', d\mathbf{y}_{k'}^{(j^*)})) \, d\mathbf{x}_k \\ &= \int_{\{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}} \pi(k', \mathbf{y}_{k'}^{(j^*)} | \mathbf{D}_n) \int_{\{\mathbf{x}_k \in A_k\}} P((k', \mathbf{y}_{k'}^{(j^*)}), (k, d\mathbf{x}_k)) \, d\mathbf{y}_{k'}^{(j^*)}. \end{aligned}$$

Note that we abused notation by denoting the integration variables $\mathbf{y}_{k'}^{(j^*)}$ and \mathbf{x}_k because a group of vectors like $\mathbf{y}_{\bullet}^{(0,\bullet)}$ are used in the transitions and they are not of the same dimension as $\mathbf{y}_{k'}$ and \mathbf{x}_k . To simplify, we will use notation like $\mathbf{y}_s^{(j)} := \mathbf{y}_s^{(0,j)}$ to denote the j -th proposal, $j \in \{1, \dots, N\}$.

We now introduce notation to improve readability. We define three joint densities that are used to enhance the approximations when Step 2.(i) is applied

to sample the proposal:

$$\begin{aligned}\bar{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) &:= \prod_{l \in \mathbf{N}(k) \setminus \{k'\}} \prod_{j=1}^N q_{k \mapsto l}(\mathbf{y}_l^{(j)}), \\ \bar{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) &:= \prod_{j=1}^N q_{k \mapsto k'}(\mathbf{y}_{k'}^{(j)}), \\ \bar{\bar{q}}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) &:= \prod_{l \in \mathbf{N}(k') \setminus \{k\}} q_{k' \mapsto l}(\mathbf{z}_l^{(j^*)}) \prod_{j=1(j \neq j^*)}^N q_{l \mapsto k'}(\mathbf{z}_{k'}^{(j)}).\end{aligned}$$

The densities $\bar{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}$ and $\bar{q}_{k \mapsto k'}$ together represent the joint density of the random variables sampled in the first part of Step 2.(i). $\bar{\bar{q}}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}$ represents the joint density of the random variables sampled in the second part of Step 2.(i).

We now define three joint densities that are used to enhance the approximations when Step 2.(ii) is applied to sample the proposal:

$$\begin{aligned}\tilde{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) &:= \prod_{l \in \mathbf{N}(k) \setminus \{k'\}} q_{k \mapsto l}(\mathbf{y}_l^{(j^*)}) \prod_{j=1(j \neq j^*)}^N q_{l \mapsto k}(\mathbf{y}_k^{(j)}), \\ \tilde{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) &:= q_{k \mapsto k'}(\mathbf{y}_{k'}^{(j^*)}) \prod_{j=1(j \neq j^*)}^N q_{k' \mapsto k}(\mathbf{y}_k^{(j)}), \\ \tilde{\tilde{q}}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) &:= \prod_{l \in \mathbf{N}(k') \setminus \{k\}} \prod_{j=1}^N q_{k' \mapsto l}(\mathbf{z}_l^{(j)}).\end{aligned}$$

The densities $\tilde{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}$ and $\tilde{q}_{k \mapsto k'}$ together represent the joint density of the random variables sampled in the first part of Step 2.(ii). $\tilde{\tilde{q}}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}$ represents the joint density of the random variables sampled in the second part of Step 2.(ii).

We have that

$$\begin{aligned}P((k, \mathbf{x}_k), (k', d\mathbf{y}_{k'}^{(j^*)})) &= \frac{1}{2} \bar{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) \bar{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) \\ &\times g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)}) \frac{r_{\text{RJ}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'}^{(j^*)}))}{N \bar{r}(k, k', \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)})} \\ &\times \bar{\bar{q}}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) \\ &\times \left(1 \wedge \frac{g_{\text{imp.}}(k', k, \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)}, \mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)})} \bar{r}(k, k', \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)}) \right) \\ &\times d\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)} d\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)} d\mathbf{y}_{k'}^{(\bullet)} \\ &+ \frac{1}{2} \tilde{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) \tilde{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)}) \frac{1}{N}\end{aligned}$$

$$\begin{aligned}
 & \times \tilde{q}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) \\
 & \times \left(1 \wedge \frac{g_{\text{imp.}}(k', k, \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)}, \mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)})} \bar{r}(k', k, \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)})^{-1} \right) \\
 & \times d\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)} d\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)} d\mathbf{y}_{k'}^{(\bullet)} \\
 & + \delta_{(k, \mathbf{x}_k)}(k', d\mathbf{y}_{k'}^{(j^*)}) \mathbb{P}(\text{rejection} \mid (k', \mathbf{y}_{k'}^{(j^*)})),
 \end{aligned}$$

where $\mathbb{P}(\text{rejection} \mid (k', \mathbf{y}_{k'}^{(j^*)}))$ is the rejection probability given that the current state is $(k', \mathbf{y}_{k'}^{(j^*)})$. Note that we considered that in Step 2.(ii) we set uniformly at random the index of the proposal. This is however in practice not important (which is why in Algorithm 5 we set it to be 1) because of the form of the acceptance ratio. Note also that we use the notation $\bar{r}(k, k', \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)})$ to be clear about which variables is involved.

The probability of reaching the state $\{k'\} \times \{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}$, from $\{k\} \times \{\mathbf{x}_k \in A_k\}$, is thus given by

$$\begin{aligned}
 & \int_{\{\mathbf{x}_k \in A_k\}} \pi(k, \mathbf{x}_k \mid \mathbf{D}_n) \int_{\{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}} \frac{1}{2} \bar{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) \bar{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) \\
 & \times g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)}) \frac{r_{\text{RJ}}((k, \mathbf{x}_k), (k', \mathbf{y}_{k'}^{(j^*)}))}{N \bar{r}(k, k', \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)})} \\
 & \times \bar{q}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) \\
 & \times \left(1 \wedge \frac{g_{\text{imp.}}(k', k, \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)}, \mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)})} \bar{r}(k, k', \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)}) \right) \\
 & \times d\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)} d\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)} d\mathbf{y}_{k'}^{(\bullet)} d\mathbf{x}_k \\
 & + \int_{\{\mathbf{x}_k \in A_k\}} \pi(k, \mathbf{x}_k \mid \mathbf{D}_n) \int_{\{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}} \frac{1}{2} \tilde{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) \tilde{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) \\
 & \times g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)}) \frac{1}{N} \tilde{q}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) \\
 & \times \left(1 \wedge \frac{g_{\text{imp.}}(k', k, \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)}, \mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)})}{g_{\text{imp.}}(k, k', \mathbf{x}_k, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{y}_{k'}^{(\bullet)})} \bar{r}(k', k, \mathbf{x}_k, \mathbf{y}_{k'}^{(\bullet)})^{-1} \right) \\
 & \times d\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)} d\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)} d\mathbf{y}_{k'}^{(\bullet)} d\mathbf{x}_k \\
 & + \int_{\{\mathbf{x}_k \in A_k\}} \pi(k, \mathbf{x}_k \mid \mathbf{D}_n) \\
 & \times \int_{\{\mathbf{y}_{k'}^{(j^*)} \in A_{k'}\}} \delta_{(k, \mathbf{x}_k)}(k', d\mathbf{y}_{k'}^{(j^*)}) \mathbb{P}(\text{rejection} \mid (k', \mathbf{y}_{k'}^{(j^*)})) d\mathbf{x}_k.
 \end{aligned}$$

(20)

We now prove that the first integral can be rewritten as that corresponding to Step 2.(ii) for the reverse move; the second integral corresponds instead to Step 2.(i) for the reverse move, and the last term to the probability of rejecting from $(k', \mathbf{y}_{k'}^{(j^*)})$.

By changing the integration variables $\mathbf{y}_{k'} \leftarrow \mathbf{y}_{k'}^{(j^*)}$, $\mathbf{x}_k^{(j^*)} \leftarrow \mathbf{x}_k$ and $\mathbf{x}_k^{(\bullet)} \leftarrow \mathbf{y}_{k'}^{(\bullet)}$ but with $\mathbf{y}_{k'}$ replaced by $\mathbf{x}_k^{(j^*)}$, the first integral can be rewritten as

$$\begin{aligned} & \int_{\{\mathbf{y}_{k'} \in A_{k'}\} \times \{\mathbf{x}_k^{(j^*)} \in A_k\}} \pi(k', \mathbf{y}_{k'} \mid \mathbf{D}_n) \frac{1}{2} \tilde{q}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) \tilde{q}_{k' \mapsto k}(\mathbf{x}_k^{(\bullet)}) \\ & \times g_{\text{imp.}}(k', k, \mathbf{y}_{k'}, \mathbf{x}_k^{(\bullet)}, \mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) \frac{1}{N} \tilde{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) \\ & \times \left(1 \wedge \frac{g_{\text{imp.}}(k, k', \mathbf{y}_{k'}, \mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}, \mathbf{x}_k^{(\bullet)})}{g_{\text{imp.}}(k', k, \mathbf{y}_{k'}, \mathbf{x}_k^{(\bullet)}, \mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)})} \bar{r}(k, k', \mathbf{y}_{k'}, \mathbf{x}_k^{(\bullet)})^{-1} \right) \\ & \times d\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)} d\mathbf{x}_k^{(\bullet)} d\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)} d\mathbf{y}_{k'}, \end{aligned}$$

given that

$$\begin{aligned} r_{\text{RJ}}((k, \mathbf{x}_k^{(j^*)}), (k', \mathbf{y}_{k'})) &= \frac{\pi(k', \mathbf{y}_{k'} \mid \mathbf{D}_n) q_{k' \mapsto k}(\mathbf{x}_k^{(j^*)})}{\pi(k, \mathbf{x}_k^{(j^*)} \mid \mathbf{D}_n) q_{k \mapsto k'}(\mathbf{y}_{k'})}, \\ \bar{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) &= \tilde{q}_{k \mapsto \mathbf{N}(k) \setminus \{k'\}}(\mathbf{y}_{\mathbf{N}(k) \setminus \{k'\}}^{(\bullet)}) \\ \bar{q}_{k \mapsto k'}(\mathbf{y}_{k'}^{(\bullet)}) \frac{q_{k' \mapsto k}(\mathbf{x}_k^{(j^*)})}{q_{k \mapsto k'}(\mathbf{y}_{k'})} &= \tilde{q}_{k' \mapsto k}(\mathbf{x}_k^{(\bullet)}) \\ \bar{q}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}) &= \tilde{q}_{k' \mapsto \mathbf{N}(k') \setminus \{k\}}(\mathbf{z}_{\mathbf{N}(k') \setminus \{k\}}^{(\bullet)}). \end{aligned}$$

The analysis of the second integral in (20) uses the same arguments. Finally, the third integral in (20) can be rewritten as

$$\int_{\{\mathbf{y}_{k'} \in A_{k'}\}} \pi(k', \mathbf{y}_{k'} \mid \mathbf{D}_n) \int_{\{\mathbf{x}_k \in A_k\}} \delta_{(k', \mathbf{y}_{k'})}(k, d\mathbf{x}_k) \mathbb{P}(\text{rejection} \mid (k, \mathbf{x}_k)) d\mathbf{y}_{k'},$$

which concludes the proof. \square

Proof of Proposition 3. The proof essentially relies on straightforward calculations. We have

$$\begin{aligned} \pi(k, \boldsymbol{\beta}_k, \sigma_k \mid \gamma_n) &\propto \pi(k) \frac{1}{\sigma_k} \prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma_k} \exp\left(-\frac{1}{2\sigma_k^2} (\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k)^2\right) \\ &\propto \pi(k) \frac{1}{\sigma_k^{n+1}} \exp\left(-\frac{1}{2\sigma_k^2} \sum_{i=1}^n (\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k)^2\right). \end{aligned}$$

In Gagnon, Desgagné and Bédard (2020b), it is proved that

$$\sum_{i=1}^n (\gamma_i - \mathbf{c}_{i,k}^T \boldsymbol{\beta}_k)^2 = (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k)^T \mathbf{C}_K^T \mathbf{C}_K (\boldsymbol{\beta}_k - \hat{\boldsymbol{\beta}}_k) + \|\gamma_n - \hat{\gamma}_k\|_2^2,$$

where $\hat{\beta}_k := (\mathbf{C}_k^T \mathbf{C}_k)^{-1} \mathbf{C}_k^T \gamma_n$. Multiplying and dividing by the appropriate terms yield

$$\begin{aligned} \pi(k, \beta_k, \sigma_k \mid \gamma_n) &\propto \pi(k) \frac{\Gamma((n - d_k)/2) \pi^{d_k/2}}{\|\gamma_n - \hat{\gamma}_k\|_2^{n-d_k} |\mathbf{C}_k^T \mathbf{C}_k|^{1/2}} \\ &\times \frac{2^{1-\frac{n-d_k}{2}} \|\gamma_n - \hat{\gamma}_k\|_2^{n-d_k}}{\Gamma((n - d_k)/2) \sigma_k^{n-d_k+1}} \exp\left\{-\frac{1}{2\sigma_k^2} \|\gamma_n - \hat{\gamma}_k\|_2^2\right\} \\ &\times \frac{|\mathbf{C}_k^T \mathbf{C}_k|^{1/2}}{(2\pi)^{d_k/2} \sigma_k^{d_k}} \exp\left(-\frac{1}{2\sigma_k^2} (\beta_k - \hat{\beta}_k)^T \mathbf{C}_K^T \mathbf{C}_K (\beta_k - \hat{\beta}_k)\right), \end{aligned}$$

which concludes the proof. \square

References

- ANDRIEU, C., DOUCET, A., YILDIRIM, S. and CHOPIN, N. (2018). On the utility of Metropolis–Hastings with asymmetric acceptance ratio. [arXiv:1803.09527](#).
- BARKER, A. A. (1965). Monte Carlo calculations of the radial distribution functions for a proton-electron plasma. *Austral. J. Phys.* **18** 119–134.
- BROOKS, S. P., GIUDICI, P. and ROBERTS, G. O. (2003). Efficient construction of reversible jump Markov Chain Monte Carlo proposal distributions. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **65** 3–39. [MR1959092](#)
- CARPENTER, B., GELMAN, A., HOFFMAN, M. D., LEE, D., GOODRICH, B., BETANCOURT, M., BRUBAKER, M., GUO, J., LI, P. and RIDDELL, A. (2017). Stan: A probabilistic programming language. *J. Stat. Softw.* **76**.
- CASELLA, G., GIRÒN, F. J., MARTÍNEZ, M. L. and MORENO, E. (2009). Consistency of Bayesian Procedures for Variable Selection. *Ann. Statist.* **37** 1207–1228. [MR2509072](#)
- DAVISON, A. (1986). Approximate predictive likelihood. *Biometrika* **73** 323–332. [MR0855892](#)
- DESGAGNÉ, A. (2015). Robustness to Outliers in Location–Scale Parameter Model using Log-Regularly Varying Distributions. *Ann. Statist.* **43** 1568–1595. [MR3357871](#)
- DESGAGNÉ, A. and GAGNON, P. (2019). Bayesian robustness to outliers in linear regression and ratio estimation. *Braz. J. Probab. Stat.* **33** 205–221. [arXiv:1612.05307](#). [MR3919021](#)
- DEVROYE, L., MEHRABIAN, A. and REDDAD, T. (2018). The total variation distance between high-dimensional Gaussians. [arXiv:1810.08693](#).
- GAGNON, P., BÉDARD, M. and DESGAGNÉ, A. (2021). An automatic robust Bayesian approach to principal component regression. *J. Appl. Stat.* **48** 84–104. [arXiv:1711.06341](#). [MR4183269](#)
- GAGNON, P., DESGAGNÉ, A. and BÉDARD, M. (2020a). A new Bayesian approach to robustness against outliers in linear regression. *Bayesian Anal.* **15** 389–414. [MR4078719](#)

- GAGNON, P., DESGAGNÉ, A. and BÉDARD, M. (2020b). A new Bayesian approach to robustness against outliers in linear regression – supplementary material. *Bayesian Anal.* **15** 389–414. [MR4078719](#)
- GAGNON, P. and DOUCET, A. (2021). Nonreversible Jump Algorithms for Bayesian Nested Model Selection. *J. Comput. Graph. Statist.* **30** 312–323. [arXiv:1911.01340](#). [MR4270506](#)
- GAGNON, P. and MAIRE, F. (2020). An asymptotic Peskun ordering and its application to lifted samplers. [arXiv:2003.05492](#).
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82** 711–732. [MR1380810](#)
- GREEN, P. J. (2003). Trans-dimensional Markov chain Monte Carlo. In *Highly structured stochastic systems* 179–196. Oxford Univ Press. [MR2082410](#)
- HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** 97–109. [MR3363437](#)
- HAYASHI, Y. (2020). Bayesian analysis on limiting Student- t linear regression model. [arXiv:2008.04522](#).
- HOETING, J. A., MADIGAN, D., RAFTERY, A. E. and VOLINSKY, C. T. (1999). Bayesian model averaging: a tutorial. *Statist. Sci.* 382–401. [MR1765176](#)
- HOFFMAN, M. D. and GELMAN, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15** 1593–1623. [MR3214779](#)
- JEFFREYS, H. (1967). *Theory of Probability*. Oxford Univ. Press, London. [MR0000924](#)
- JOHNSON, V. E. and ROSSELL, D. (2012). Bayesian model selection in high-dimensional settings. *J. Amer. Statist. Assoc.* **107** 649–660. [MR2980074](#)
- KARAGIANNIS, G. and ANDRIEU, C. (2013). Annealed Importance Sampling Reversible Jump MCMC Algorithms. *J. Comp. Graph. Stat.* **22** 623–648. [MR3173734](#)
- KLEIJN, B. J. K. and VAN DER VAART, A. W. (2012). The Bernstein-von-Mises theorem under misspecification. *Electron. J. Stat.* **6** 354–381. [MR2988412](#)
- LINDLEY, D. V. (1957). A Statistical Paradox. *Biometrika* **44** 187–192. [MR0087273](#)
- LIVINGSTONE, S. and ZANELLA, G. (2019). The Barker proposal: combining robustness and efficiency in gradient-based MCMC. [arXiv:1908.11812](#).
- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H. and TELLER, E. (1953). Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21** 1087.
- NEAL, R. M. (2011). MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo* 113–160. CRC Press New York, NY. [MR2858447](#)
- O’HAGAN, A. and PERICCHI, L. (2012). Bayesian heavy-tailed models and conflict resolution: A review. *Braz. J. Probab. Stat.* **26** 372–401. [MR2949085](#)
- PESKUN, P. H. (1973). Optimum Monte-Carlo sampling using Markov chains. *Biometrika* **60** 607–612. [MR0362823](#)
- RICHARDSON, S. and GREEN, P. J. (1997). On Bayesian analysis of mixtures with an unknown number of components. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **59** 731–792. [MR1483213](#)

- ROBERTS, G. O. and ROSENTHAL, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **60** 255–268. [MR1625691](#)
- ROBERTS, G. O. and TWEEDIE, R. L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** 341–363. [MR1440273](#)
- SCHMON, S. M., DELIGIANNIDIS, G., DOUCET, A. and PITT, M. K. (2021). Large-sample asymptotics of the pseudo-marginal method. *Biometrika* **108** 37–51. [MR4226188](#)
- STAMEY, T. A., KABALIN, J. N., MCNEAL, J. E., JOHNSTONE, I. M., FREIHA, F., REDWINE, E. A. and YANG, N. (1989). Prostate specific antigen in the diagnosis and treatment of adenocarcinoma of the prostate. II. Radical prostatectomy treated patients. *J. Urol* **141** 1076–1083.
- TIERNEY, L. (1994). Markov chains for exploring posterior distributions. *Ann. Statist.* 1701–1728. [MR1329166](#)
- TIERNEY, L. (1998). A Note on Metropolis-Hastings Kernels for General State Spaces. *Ann. Appl. Probab.* **8** 1–9. [MR1620401](#)
- VAN DER VAART, A. W. (2000). *Asymptotic statistics* **3**. Cambridge university press. [MR1652247](#)
- VERDINELLI, I. and WASSERMAN, L. (1991). Bayesian analysis of outlier problems using the Gibbs sampler. *Stat. Comput.* **1** 105–117.
- ZANELLA, G. (2020). Informed proposals for local MCMC in discrete spaces. *J. Amer. Statist. Assoc.* **115** 852–865. [MR4107684](#)