

Reversible Jump PDMP Samplers for Variable Selection

Augustin Chevalier, Paul Fearnhead and Matt Sutton
Department of Mathematics and Statistics, Lancaster University *

October 23, 2020

Abstract

A new class of Markov chain Monte Carlo (MCMC) algorithms, based on simulating piecewise deterministic Markov processes (PDMPs), have recently shown great promise: they are non-reversible, can mix better than standard MCMC algorithms, and can use subsampling ideas to speed up computation in big data scenarios. However, current PDMP samplers can only sample from posterior densities that are differentiable almost everywhere, which precludes their use for model choice. Motivated by variable selection problems, we show how to develop reversible jump PDMP samplers that can jointly explore the discrete space of models and the continuous space of parameters. Our framework is general: it takes any existing PDMP sampler, and adds two types of trans-dimensional moves that allow for the addition or removal of a variable from the model. We show how the rates of these trans-dimensional moves can be calculated so that the sampler has the correct invariant distribution. Simulations show that the new samplers can mix better than standard MCMC algorithms. Our empirical results show they are also more efficient than gradient-based samplers that avoid model choice through use of continuous spike-and-slab priors which replace a point mass at zero for each parameter with a density concentrated around zero.

Keywords: Bayesian Statistics; Bouncy Particle Sampler; Model Choice; Monte Carlo; Zig Zag Algorithm

1 Introduction

There currently is much interest in developing MCMC algorithms based on simulating piecewise deterministic Markov processes (PDMPs). These are continuous time Markov

*This research was supported by EPSRC grants EP/R018561 and EP/R034710.

processes that have deterministic dynamics between a set of event times, and the randomness in these processes only comes through the random event times and potentially random transitions at the events (see Davis, 1993, for an introduction to PDMPs).

The idea of simulating PDMPs to sample from a target distribution of interest originated in statistical physics (Peters and de With, 2012; Michel *et al.*, 2014), but has recently been proposed as an alternative to standard MCMC to sample from posterior distributions in Bayesian Statistics, with such algorithms as the Bouncy Particle Sampler (Bouchard-Côté *et al.*, 2018) and the ZigZag algorithm (Bierkens and Roberts, 2017; Bierkens *et al.*, 2019) amongst others (Vanetti *et al.*, 2017; Markovic and Sepehri, 2018; Wu and Robert, 2020; Michel *et al.*, 2020; Bierkens *et al.*, 2020). See Fearnhead *et al.* (2018) for an introduction to this area.

To sample from a density $\pi(\mathbf{x})$ current PDMP samplers introduce a velocity component, \mathbf{v} , of the same dimension as \mathbf{x} , and have deterministic dynamics that correspond to a constant velocity model. At the random events the velocity component changes. Algorithms differ in terms of the event rate and how the velocity changes at each event, but each has a simple recipe for choosing these so that the resulting PDMP has $\pi(\mathbf{x})$ as its invariant distribution. These recipes depend on $\pi(\mathbf{x})$ through the gradient of $\log \pi(\mathbf{x})$, which importantly means that $\pi(\mathbf{x})$ only needs to be known up to proportionality, but also that $\pi(\mathbf{x})$ needs to be differentiable almost everywhere. The advantages of PDMP samplers are that they are non-reversible, and thus can mix more quickly than standard reversible MCMC algorithms (Diaconis *et al.*, 2000), and, when sampling from posterior distributions, they can use a small sample of data points at each iteration whilst still targeting the true posterior distribution (Bierkens *et al.*, 2019).

However, the restriction to sampling from densities that are differentiable means that current PDMP samplers cannot be used in model choice problems. The aim of this paper is to address this limitation, with a particular motivation of PDMP samplers that can be used in variable selection problems that are common in, for example, linear regression and general linear regression. We show how to design efficient PDMP samplers which allow movement between different models.

A simple way to implement PDMP samplers for variable selection problems is to use continuous spike-and-slab priors on the parameters (Ishwaran and Rao, 2005; George and McCulloch, 1993), which, rather than setting some parameters exactly to 0, have priors that place substantial mass close to 0. With such a prior, the resulting posterior density is differentiable, and existing PDMP samplers can be used. However such an approach has three disadvantages. First, under such a prior it can be hard to interpret the results as we do not formally get posterior probabilities on whether certain variables should be included in the model. Second, they introduce an extra tuning parameter to the prior which governs the shape of the spike of the component. Third, as we show in Section 2.2, using PDMP samplers to sample from the resulting posterior can be computationally inefficient: the samplers will need to simulate many events so that the parameters associated to variables that should not be in the model are kept close to 0.

We demonstrate how to adapt existing PDMP samplers to variable selection problems. Specifically they evolve as the PDMP sampler when exploring the posterior associated with a given model, but with two additional events: if any parameter value hits 0 the PDMP jumps to the smaller model where the corresponding variable is removed; whilst with some rate there are events that re-introduce variables into the model. We show in Section 3 how to calculate the rate and transition for these new types of event so that the sampler has the correct invariant distribution. To calculate these we need different techniques than those used for existing PDMP samplers, as we need to account for the behaviour of the process when parameters hit zero. The techniques we use are most similar to those in Bierkens *et al.* (2018), which considers PDMPs with restricted domains. However in that paper the dynamics at the boundary of the domain could be chosen so that the net flow of probability at the boundary is zero; whereas we need to balance the probability flow out of a model which occurs when a parameter hits zero with the flow into the model caused by the events that re-introduce variables.

The approach we present is generic, in that it can take any current PDMP sampler and be used to obtain a version that can be applied to the variable selection problem. We call the new class of samplers reversible jump PDMP samplers, due to the analogy with reversible jump MCMC (Green, 1995). We show how to derive reversible-jump versions of both ZigZag and the Bouncy Particle Sampler in Section 4, before investigating empirically these algorithms on both logistic regression and robust linear regression models.

Proofs of all theorems are relegated to the appendix. Code for implementing the new reversible jump PDMP samplers, and for replicating our examples, is available from <https://github.com/matt-sutton/rjpdmp>.

2 PDMP Samplers and Model Choice

2.1 Variable Selection

We will consider model selection problems that arise from variable selection. The general framework is that we have a vector of parameters, $\theta = (\theta_1, \dots, \theta_p)$, and each model is characterised by setting some subset of the θ_j s to 0. This is a common setting across linear models, generalised linear models and various extensions.

To make ideas concrete, consider a linear model

$$\mathbf{Y} = \sum_{j=1}^p \mathbf{X}_j \theta_j + \epsilon$$

where \mathbf{Y} is a vector of response variables, each \mathbf{X}_j is a vector of covariates, and ϵ is an additive noise vector. When p is large it is common to fit such a model under a sparsity assumption, namely that many of the θ_j s are 0.

In a Bayesian analysis, such a sparsity assumption is encapsulated in our choice of prior on θ . To aid interpretation of the variable selection priors it is common to introduce a

latent variable $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)'$ where $\gamma_j = 1$ if the covariate X_j is included in the model, i.e. if $\theta_j \neq 0$. We let $|\boldsymbol{\gamma}| = \sum \gamma_j$ the number of covariates included in the corresponding model. Indexing $\boldsymbol{\theta}_\gamma$ as the sub-vector of $\boldsymbol{\theta}$ with only the selected variables, any prior can be written in a hierarchical form where we have a prior on $\boldsymbol{\gamma}$, then conditional on $\boldsymbol{\gamma}$ we have a prior on $\boldsymbol{\theta}_\gamma$, and set all remaining entries of $\boldsymbol{\theta}$ to 0.

A special case is where each component of $\boldsymbol{\theta}$ is independent of the others. In which case the prior can be written as

$$\theta_j \sim w_j g_j(\theta_j) + (1 - w_j) \delta_0(\theta_j), \quad j = 1, \dots, p,$$

where $w_j \in (0, 1)$ is the prior probability that $\gamma_j = 1$, δ_0 is a Dirac measure at zero and $g_j(\theta_j)$ is a distribution that models our prior beliefs for θ_j conditional on that variable being included in the model. Bayesian approaches to variable selection that put a probability mass on $\theta_j = 0$ in this way will be referred to as Dirac spike and slab methods. Notable examples of these methods include Mitchell and Beauchamp (1988); Kuo and Mallick (1998); Geweke (1996); Smith and Kohn (1996); Bottolo and Richardson (2010).

While this formulation is natural from a modelling perspective it, sampling from the resulting posterior distribution can be challenging, with, for example, MCMC samplers that use gradient information such as Hamiltonian Monte Carlo (Neal, 2011) not being applicable. To circumvent this issue it is common to use an approximation to this prior which replaces the point mass at 0 with a density that is peaked around 0, such as

$$\theta_j \sim w_j \mathcal{N}(0, \tau_j^2) + (1 - w_j) \mathcal{N}(0, \tau_j^2 c_j^2), \quad j = 1, \dots, p, \quad (1)$$

where c_j is taken small so that $\mathcal{N}(0, c_j^2 \tau_j^2)$ approximates the Dirac spike. We will refer to Bayesian variable selection methods that replace the Dirac in the prior with a continuous approximation as continuous spike and slab methods. This prior was originally proposed in linear regression where it is commonly referred to as the stochastic search variable selection procedure (George and McCulloch, 1993).

2.2 PDMP Samplers

Piecewise Deterministic Markov Processes (PDMPs) are an emerging class of non-reversible continuous-time samplers. To be consistent with commonly used notation for PDMP samplers, we will consider sampling from a distribution with density $\pi(\boldsymbol{x})$ defined on some space \mathcal{X} ; which is a slight change of notation relative to our regression model of the previous section. Current samplers augment the state to include a velocity vector and sample from a distribution on $E = \mathcal{X} \times \mathcal{V}$. In the following, for $\boldsymbol{z} \in E$ we will use the notation $\boldsymbol{z} = (\boldsymbol{x}, \boldsymbol{v})$ with $\boldsymbol{x} \in \mathcal{X}$ a position and $\boldsymbol{v} \in \mathcal{V}$ a velocity.

The PDMP can be defined by (i) deterministic dynamics between a set of random event times; (ii) the state-dependent rate at which events occur, $\lambda(\boldsymbol{z})$; and (iii) a probability distribution for the change in state at each event, with density $q(\boldsymbol{z}'|\boldsymbol{z})$. We will consider

PDMPs whose deterministic trajectories follow a differential equation of the form:

$$\frac{d(\mathbf{x}_t, \mathbf{v}_t)}{dt} = (\mathbf{v}_t, \Phi(\mathbf{x}_t, \mathbf{v}_t))$$

with $\Phi : E \rightarrow \mathcal{V}$ a smooth function. This setting contains the usual PDMP samplers such as ZigZag (Bierkens *et al.*, 2019), the Bouncy Particle Sampler (Bouchard-Côté *et al.*, 2018), or the Coordinate Sampler (Wu and Robert, 2020).

For example the ZigZag algorithm to simulate from $\pi(\mathbf{x})$ for $\mathbf{x} \in \mathbb{R}^p$, would introduce a velocity vector $\mathbf{v} \in \{-1, 1\}^p$, and deterministic dynamics

$$\frac{d(\mathbf{x}_t, \mathbf{v}_t)}{dt} = (\mathbf{v}_t, 0),$$

which are the dynamics of constant velocity model. Events occur at a rate that depends on the gradient of $\pi(\mathbf{x})$ in each component of the velocity, and at an event one component of the velocity is switched. The rate at which the j th component of the velocity, \mathbf{v}_t^j is switched is just

$$\max \left\{ 0, -\mathbf{v}_t^j \frac{\partial}{\partial \mathbf{x}^j} \log \pi(\mathbf{x}_t) \right\}.$$

These rates depend on the target distribution just through the gradient of the log of the target – which importantly means that we need only know the target distribution up to proportionality.

We can apply current PDMPs, such as ZigZag, to the Bayesian variable selection problem if we use the continuous spike-and-slab prior (1). Realisations of such a sampler are shown in the left two plots of Figure 1 as we vary how concentrated the spike distribution is. For the more concentrated case, the sampler becomes inefficient as it involves many switching events when the state variable is close to 0.

Intuitively as we make the variance of the spike component of the prior tend to 0 the prior converges to a prior with a point mass at 0; furthermore we can observe the output of our PDMP sampler “converging” to a process shown in the right-hand plot of Figure 1. Here, rather than the state having periods where it oscillates around 0, it has periods of time where $\theta_j = 0$ and thus has many fewer events.

3 Reversible Jump PDMP Samplers

Let $\mathcal{M} = (\mathcal{M}_1, \mathcal{M}_2, \dots)$ be a set of models indexed by a parameter k . To each model \mathcal{M}_k correspond a state space \mathcal{X}^k of dimension p_k . For the sake of clarity, we will limit ourselves to variable selection, though it is straightforward to apply the results below to more general model selection problems. In our case \mathcal{X}^k has a specific form:

$$\mathcal{X}^k = \prod_i \mathbb{R}^{\gamma_i^k}$$




Figure 1: Sample paths of PDMPs implementing the variable selection priors in 1 dimension. The left and centre plots show the trajectories for a continuous spike-and-slab prior $0.5\mathcal{N}(0, \tau^2) + 0.5\mathcal{N}(0, \tau^2 c^2)$ where $\tau^2 = 16$. As c decreases the spike component in the mixture approaches a Dirac mass. The figure on the right is the limiting process where we set the velocity to zero allowing the variable to stay fixed at zero.

where we abuse notation by using $\mathbb{R}^0 = \{0\}$ and where $(\gamma_i^k)_i$ is a sequence of numbers in $\{0, 1\}$ representing whether a variable is enabled for model \mathcal{M}_k . Let π be the target posterior probability defined on $\mathcal{X} = \cup_k \mathcal{X}^k$. We further assume that the restriction of π to each \mathcal{X}^k has a density, and denote this by $\pi_k(\mathbf{x})$. The first ingredient of our sampler is a collection of PDMPs defined for each model. Each PDMP sampler adds a velocity space \mathcal{V}^k to the space \mathcal{X}^k and samples from the space $E_k = \mathcal{X}^k \times \mathcal{V}^k$. Finally, for each model \mathcal{M}_k , the associated PDMP sampler has an extended infinitesimal generator \mathcal{A}_k (Davis, 1993) with invariant distribution proportional to ν_k with

$$\nu_k(\mathbf{x}, \mathbf{v}) = \pi_k(\mathbf{x})p_k(\mathbf{v}|\mathbf{x}) \quad (\mathbf{x}, \mathbf{v}) \in \mathcal{X}^k \times \mathcal{V}^k,$$

for some set of conditional densities $p_k(\mathbf{v}|\mathbf{x})$.

Remark 1. For samplers such as ZigZag, $p_k(d\mathbf{v}|\mathbf{x})$ is a measure with support on a discrete set. By choosing \mathcal{V}^k to be the support of $p_k(d\mathbf{v}|\mathbf{x})$, $p_k(d\mathbf{v}|\mathbf{x})$ still has a density – and

integrals can be interpreted as sums over the support points. This allows us to treat all samplers within the same framework.

The second ingredient of our reversible jump PDMP sampler is a set of jumps between models. For our case of variable selection, we only allow adding or removing one variable at a time. Hence, let

$$\mathcal{T} = \{(i, j) \mid \sum_l |\gamma_l^i - \gamma_l^j| = 1 \text{ and } |\gamma^i| > |\gamma^j|\}$$

be a set of pairs of transitions between models, with these ordered (i, j) such that model \mathcal{M}_j being obtained from \mathcal{M}_i by removing one of the variables in \mathcal{M}_j . For transition $i \rightarrow j$, we define an active boundary $\Gamma_{i,j} = \mathcal{X}^j \times \mathcal{V}^i$, a subspace of E_i . Each trajectory passing through $\Gamma_{i,j}$ has some probability $p_{i,j}$ of jumping to E_j using a deterministic jump function $g_{i,j}$. We assume that the jump function does not change the position. So, if a trajectory \mathbf{z}_t of our process has a left limit at time t , \mathbf{z}_{t-} in $\Gamma_{i,j}$, then with some probability $p_{i,j}$, $\mathbf{z}_t = g_{i,j}(\mathbf{z}_{t-}) \in E_j$ with $\mathbf{x}_t = \mathbf{x}_{t-}$. For transition $j \rightarrow i$, we introduce $\beta_{i,j}(\mathbf{z})$ a Poisson rate and a jump kernel $Q_{i,j}(\cdot, \mathbf{z})$ such that if the trajectory is in E_j , then with rate $\beta_{i,j}(\mathbf{z}_{t-})$, \mathbf{z}_t is drawn from $Q_{i,j}(\cdot, \mathbf{z}_{t-}) \in \Gamma_{i,j}$. We impose symmetry in the jumps between models so that for any $\mathbf{z}' \sim Q_{i,j}(\cdot, \mathbf{z})$, $g_{i,j}(\mathbf{z}') = \mathbf{z}$.

3.1 Invariant distribution

Let $\nu = \sum_k \nu_k$ be a measure on $\cup_k E_k$. The \mathbf{x} -marginal distribution of ν is π and this section provides conditions on $Q_{i,j}, p_{i,j}$ and $\beta_{i,j}$ that has to be satisfied if ν is to be the invariant distribution of the process. To avoid additional technicalities that complicate the proof, we give results for PDMPs with bounded velocity spaces.

Theorem 1. *Assume \mathcal{V}^k is bounded for each k , and that we can bound the event rate over any compact region; i.e if $K \in E$ is compact then $\max_{\mathbf{z} \in K} \lambda(\mathbf{z}) < \infty$. Suppose a measure μ is a stationary distribution with density on each E_i . Then for each function f in a suitable set \mathcal{F} :*

$$\begin{aligned} 0 &= \sum_i \int_{E_i} \mathcal{A}_i f d\mu + \sum_{(i,j) \in \mathcal{T}} p_{i,j} \int_{\Gamma_{i,j}} [f(g_{i,j}(\mathbf{z})) - f(\mathbf{z})] \mu(\mathbf{z}) |\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle| dz \\ &+ \sum_{(i,j) \in \mathcal{T}} \int_{E_j} \beta_{i,j}(\mathbf{z}) \left(\int_{g_{i,j}^{-1}(\{\mathbf{z}\})} [f(\mathbf{z}') - f(\mathbf{z})] Q_{i,j}(\mathbf{z}'|\mathbf{z}) dz' \right) \mu(\mathbf{z}) dz, \end{aligned} \tag{2}$$

where $\mathbf{n}_{i,j}$ is the normal to $\Gamma_{i,j}$.

The set \mathcal{F} in this theorem is the domain of the generator of our PDMP, and is described precisely in Davis (1993).

To get a directly usable conditions on $\beta_{i,j}$, $p_{i,j}$ and $Q_{i,j}$, some additional notation must be introduced. When jumping from $\mathbf{z} = (\mathbf{x}, \mathbf{v}) \in E_j$ to $\mathbf{z}' = (\mathbf{x}, \mathbf{v}') \in E_i$, the dimension of the velocity vector needs to be increased by one, but the position is unchanged. Hence $g_{i,j}^{-1}(\mathbf{z})$ is a one-dimensional manifold, which can be related to a subset, U say, of \mathbb{R} by introducing a function

$$G_{i,j} : U \times E_j \rightarrow E_i$$

such that for any $\alpha \in U$, we have $G_{i,j}(\alpha, \mathbf{z}) \in g_{i,j}^{-1}(\{\mathbf{z}\})$ and for a fixed $\mathbf{z} \in E_j$, $\alpha \mapsto G_{i,j}(\alpha, \mathbf{z})$ is a one to one mapping from U to $g_{i,j}^{-1}(\mathbf{z})$ (similar ideas are seen in reversible jump MCMC; see Green, 1995). For a given $\mathbf{z} \in E_j$, it is natural to rewrite the jump kernel $Q_{i,j}$ in terms of $\alpha \in U$, and henceforth we abuse notation and write $Q_{i,j}(\cdot, \mathbf{z})$ as a density on U . That is we can simulate the transition from E_j to E_i , by simulating $\alpha \sim Q_{i,j}(\cdot, \mathbf{z})$ and then setting $\mathbf{z}' = G_{i,j}(\alpha, \mathbf{z})$.

Finally, let

$$\nu_{i,j}(\mathbf{z}) = \nu(\mathbf{z}) |\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle|$$

be an unnormalised density on $\Gamma_{i,j}$ and let $\bar{\nu}_{i,j}$ be the pushforward measure on E_j of $\nu_{i,j}$ by $g_{i,j}$:

$$\bar{\nu}_{i,j}(B) = \int_{g_{i,j}^{-1}(B)} \nu_{i,j}(\mathbf{z}) d\mathbf{z} \quad \text{for any } B \subset E_j \text{ measurable.}$$

Informally this is the measure under $\nu_{i,j}$ associated with values of $\mathbf{x} \in \Gamma_{i,j}$ that would be mapped by $g_{i,j}$ to the set B .

We are now in a position to give simple conditions for (2) of Theorem 1 to hold. To do this it is helpful to consider separately the cases where the space of velocities is continuous and the case where it is discrete.

Theorem 2. *Assume the space of velocities is continuous. A sufficient condition for (2) of Theorem 1 to hold for ν is that, for every $(i, j) \in \mathcal{T}$, the following conditions holds*

$$\beta_{i,j}(\mathbf{z}) = p_{i,j} \frac{\bar{\nu}_{i,j}(\mathbf{z})}{\nu(\mathbf{z})} \quad \text{for } \mathbf{z} \in E_j \tag{3}$$

$$Q_{i,j}(\alpha|\mathbf{z}) = \frac{\nu_{i,j}(G_{i,j}(\alpha, \mathbf{z})) |J_{G_{i,j}}(\alpha, \mathbf{z})|}{\bar{\nu}_{i,j}(\mathbf{z})} \quad \text{for } \mathbf{z} \in E_j \text{ and } \alpha \in U, \tag{4}$$

where $J_{G_{i,j}}$ denotes the Jacobian associated with the transformation $G_{i,j}$.

Intuitively, this result can be understood as a detailed balance condition: we balance the probability flow for each jump $\mathbf{z} \rightarrow \mathbf{z}'$ from E_i to E_j , with that of the reverse jump $\mathbf{z}' \rightarrow \mathbf{z}$.

For discrete velocity spaces the result is slightly simpler, as the Jacobian term is not required.

Theorem 3. Assume the velocity space is discrete. A sufficient condition for (2) of Theorem 1 to hold for ν is that, for every $(i, j) \in \mathcal{T}$, the following conditions holds

$$\beta_{i,j}(\mathbf{z}) = p_{i,j} \frac{\bar{\nu}_{i,j}(\mathbf{z})}{\nu(\mathbf{z})} \quad \text{for } \mathbf{z} \in E_j \quad (5)$$

$$Q_{i,j}(\alpha|\mathbf{z}) = \frac{\nu_{i,j}(G_{i,j}(\alpha, \mathbf{z}))}{\bar{\nu}_{i,j}(\mathbf{z})} \quad \text{for } \mathbf{z} \in E_j \text{ and } \alpha \in U. \quad (6)$$

4 Example Samplers

4.1 ZigZag Sampler

We first derive the jump rates and transitions for the ZigZag sampler described in Section 2.2. For this sample the velocity of each active component is either 1 or -1, which simplifies the computation.

We choose $g_{i,j}$ to be the projection that sets to 0 the disabled variable. Let $(i, j) \in \mathcal{T}$ be a transition. For any $\mathbf{v} \in \mathcal{V}^i$, we have $|\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle| = 1$. Therefore on E_i ,

$$\nu_{i,j}(\mathbf{z}) = \nu(\mathbf{z}) = \pi_i(\mathbf{x})2^{-|\gamma_i|}.$$

For $\mathbf{z} \in E_j$, since a velocity component in $\{-1, 1\}$ is projected to 0,

$$\bar{\nu}_{i,j}(\mathbf{z}) = 2\pi_i(\mathbf{x})2^{-|\gamma_i|} = \pi_i(\mathbf{x})2^{-|\gamma_j|}.$$

For the jump to E_i we change the velocity for the component of the model that is added, whilst the other components of the velocity is unchanged. Denoting the new velocity of the added component by α , we have from Theorem 3 that

$$\beta_{i,j} = p_{i,j} \frac{\pi_i(\mathbf{x})}{\pi_j(\mathbf{x})}$$

$$Q_{i,j}(\alpha|\mathbf{z}) = 1/2 \quad \text{for } \alpha \in \{-1, 1\}.$$

For our variable selection problem, the ratio of the posterior density that appears in $\beta_{i,j}$ will simplify to the ratio of the priors as the likelihood terms are common and cancel. If we have independent priors on the parameters for each variable this term will be a constant, which simplifies the simulation of the events at which we add new variables into our model. This comment applies also to the rates for the Bouncy Particle Sampler which we derive next.

4.2 Bouncy Particle Sampler

We consider two versions of the Bouncy Particle Sampler. The first version has velocities on the unit sphere, so

$$\mathcal{V}^i = \{\mathbf{v} \in \mathbb{R}^{|\gamma^k|} \text{ such that } \|\mathbf{v}\| = 1\}.$$

Like the ZigZag sampler, the deterministic dynamics are given by a constant velocity model. The event rate for sampling from a density $\pi_k(\mathbf{x})$ is

$$\lambda(\mathbf{z}) = \max\{0, -\mathbf{v} \cdot \nabla_{\mathbf{x}} \log \pi_k(\mathbf{x})\},$$

with the velocity reflecting in the normal to $\log \pi_k(\mathbf{x})$ at an event. The Bouncy particle also often has refresh events, at which a completely new velocity is sampled.

Extending the Bouncy Particle Sampler to the variable selection problem requires a more careful analysis than for the ZigZag sampler due to the geometry of the velocity space. For the jump to E_i , we will let α denote the new velocity for the component added to the model. To ensure the new velocity lies on the unit sphere, we re-scale the other components of the velocity: so if the old velocity is \mathbf{v} then the new velocity is $\alpha \mathbf{n}_{i,j} + \sqrt{1 - \alpha^2} \mathbf{v}$. The following proposition states how to choose the jump rate, $\beta_{i,j}$ and the density for α .

Proposition 1. *For the Bouncy Particle Sampler with velocities on the unit sphere, (3) and (4) are satisfied if, for all $(i, j) \in \mathcal{T}$ with $|\gamma^j| > 0$:*

$$\beta_{i,j} = p_{i,j} \frac{\pi_i(\mathbf{x})}{\pi_j(\mathbf{x})} \frac{2A_{\text{sphere}}(|\gamma^j|)}{A_{\text{sphere}}(|\gamma^i|)} \frac{1}{|\gamma^j|} \quad (7)$$

$$Q_{i,j}(\alpha|\mathbf{z}) = \frac{|\alpha||\gamma^j|\sqrt{1 - \alpha^2}^{|\gamma^j|-2}}{2} \text{ for } \mathbf{z} \in E_j \text{ and } \alpha \in (-1, 1) \quad (8)$$

with $A_{\text{sphere}}(|\gamma^i|) = \frac{\Gamma(\frac{|\gamma^i|}{2})}{2\pi^{\frac{|\gamma^i|}{2}}}$ the area of the unit sphere of $\mathbb{R}^{|\gamma^i|}$; and if for $|\gamma^j| = 0$, where the Bouncy Particle Sampler and ZigZag are equivalent, we use the ZigZag rates.

The second version of the Bouncy Particle Sampler has velocities in $\mathbb{R}^{|\gamma^k|}$, with their density being standard Gaussian and independent of \mathbf{x} . The dynamics are as previously. This sampler does not satisfy our requirement that the velocity space is bounded, though the argument behind Theorem 1 can be extended to this sampler, and our results would apply to this sampler if we replaced the Gaussian density by a truncated Gaussian (with any arbitrarily large truncation value).

For the jump to E_i , we again let α denote the new velocity for the component added to the model, whilst for this version the other components of the velocity remain unchanged.

Proposition 2. *For the Bouncy Particle Sampler with Gaussian velocities, (3) and (4) are satisfied if, for all $(i, j) \in \mathcal{T}$:*

$$\beta_{i,j} = p_{i,j} \frac{\pi_i(\mathbf{x})}{\pi_j(\mathbf{x})} \frac{2}{\sqrt{2\pi}} \quad (9)$$

$$Q_{i,j}(\alpha|\mathbf{z}) = 2|\alpha|e^{-\frac{1}{2}\alpha^2} \text{ for } \mathbf{z} \in E_j \text{ and } \alpha \in \mathbb{R}. \quad (10)$$

5 Simulation Study

In this section we demonstrate the potential advantage of our new samplers compared to alternative approaches for Bayesian variable selection. To compare between different samplers we consider the Monte Carlo estimates of the posterior probabilities of inclusion, the posterior means for the regression coefficients, and the posterior means conditioned on the model. Similar comparisons methods are used in Zanella and Roberts (2019) for comparing efficiency of MCMC methods on Bayesian variable selection problems. For a given sampler the statistical efficiency is measured by the mean squared error of the sampler denoted by $\sigma_{\text{sampler}}^2$ and is calculated using R runs of the sampler as

$$\sigma_{\text{sampler}}^2 = \frac{1}{R} \sum_{i=1}^R (\hat{q}_r - q)^2 \quad (11)$$

where \hat{q}_r for $r = 1, \dots, R$ are the estimates of a quantity of interest from the R runs, and q is either the exact posterior quantity of interest, if available, otherwise it is the estimate from an independent long run of an MCMC method. In multiple dimensions the statistical efficiency is measured as the median $\sigma_{\text{sampler}}^2$ over all dimensions. To compare the performance of different samplers we also consider a measure of efficiency relative to a reference sampler. If we denote the reference sampler by ref, then we define Relative Statistical Efficiency (RSE) and Relative Efficiency (RE) by

$$\text{RSE} = \frac{\sigma_{\text{ref}}^2 n_{\text{ref}}}{\sigma_{\text{sampler}}^2 n_{\text{sampler}}}, \quad \text{RE} = \frac{\sigma_{\text{ref}}^2 t_{\text{ref}}}{\sigma_{\text{sampler}}^2 t_{\text{sampler}}},$$

where n_{sampler} and n_{ref} are the number of iterations of the algorithms and t_{sampler} and t_{ref} are the computation times of the algorithms. The RSE measures the relative efficiency of the algorithms per iteration whereas the RE measures the efficiency per second. For interpretation an RSE or RE value of 2 implies that the sampler is 2 times *more efficient* than the reference method. The sensitivity of the methods to the choice of reversible jump parameters $p_{i,j}$ and regular PDMP tuning parameters is explored empirically on a simple model in the Supplementary Material. Based on these results we fixed $p_{i,j} = 0.6$ for all i and j in all reversible jump PDMP samplers and fixed the refreshment rate at 0.1 for the reversible jump BPS methods for the following results.

5.1 Logistic regression

First we compare PDMP based samplers with a collapsed Gibbs sampler and a reversible jump sampler on a logistic regression problem. The Gibbs sampler is based on the Poly-Gamma sampler of Polson *et al.* (2013): details of this sampler are given in the Supplementary Material. We implemented reversible jump MCMC using the NIMBLE software package (de Valpine *et al.*, 2017, 2020) using independent Normal proposal for selected

variables.

The logistic regression model has a p -dimensional regression parameter $\boldsymbol{\theta} \in \mathbb{R}^p$, and a binary response $y_i \in \{0, 1\}$ which is distributed as

$$P(y_i = 1 | \mathbf{x}_i, \boldsymbol{\theta}) = \frac{\exp(\sum_{j=1}^p x_{i,j} \theta_j)}{1 + \exp(\sum_{j=1}^p x_{i,j} \theta_j)}$$

where \mathbf{x}_i is the p -vector of covariates for observation i . In our simulation study, each vector \mathbf{x}_i is simulated from a multivariate Normal with mean zero and $p \times p$ covariance matrix Σ . We use a prior that is independent for each θ_j and where $\theta_j \sim \frac{10}{p} \mathcal{N}(0, 10) + (1 - \frac{10}{p}) \delta_0$, which corresponds to a prior that favors models with 10 selected variables. Data was generated using this model and the following choices for $\boldsymbol{\theta}$ and covariance matrix Σ :

1. A pair of correlated covariates, one of which is in the model: $\boldsymbol{\theta} = (1, 0, \dots, 0)^T$ with $\Sigma_{2,1} = \Sigma_{1,2} = 0.9$, with $\Sigma_{i,i} = 1$ and $\Sigma_{i,j} = 0$ if $i \neq j$.
2. Structured correlation between all covariates with six active covariates: $\boldsymbol{\theta} = (3, 3, -2, 3, 3, -2, 0, \dots, 0)^T$ with $\Sigma_{i,j} = \exp(-|i - j|)$.
3. No correlation between covariates and six active covariates: $\boldsymbol{\theta} = (3, 3, -2, 3, 3, -2, 0, \dots, 0)^T$, with $\Sigma_{i,i} = 1$ and $\Sigma_{i,j} = 0$ if $i \neq j$.

These simulation scenarios are analogous to others previously considered in the literature for linear regression. Scenarios 1 and 3 are similar to considered by Wang *et al.* (2011) and Zanella and Roberts (2019) while Scenario 2 is similar to one considered by Yang *et al.* (2016). We present results for both ZigZag and the Bouncy Particle Sampler with Gaussian distributed velocities (almost identical results are obtained for the Bouncy Particle Sampler where the velocities are on the unit sphere). See Section D of the Appendix for more details.

The PDMP methods are competitive with Gibbs sampling in low sample sizes and offer substantial efficiency gains for larger sample sizes. Smaller gains can also be seen when the dimension is increased with fixed sample size. Both BPS and ZigZag methods offer similar relative efficiencies across the experiments. The greatest efficiency gains for reversible jump PDMP methods was seen in Scenario 1 with Scenarios 2 and 3 offering lower gains in performance. This may be due to smaller models being more likely in Scenario 1 and thus less computational effort required for the PDMP methods in gradient calculations and simulation of event times. For all experiments the Gibbs sampler offers a computational advantage over the reversible jump methods in relative efficiency for marginal posterior probabilities of inclusion. This increased performance is unsurprising as the model update steps in the Gibbs sampler have marginalised over the parameter values $\boldsymbol{\theta}$ yielding efficient moves through the model space. In more general settings where it is not possible to integrate over $\boldsymbol{\theta}$ the mixing will be substantially poorer.

The use of subsampling methods for Bayesian variable selection problems is a recently emerging area (Song *et al.*, 2020; Buchholz *et al.*, 2019) though it remains under-studied.

Table 1: Scenario 1 (pair of correlated variables): Relative efficiencies (RE) for methods, against a Reversible Jump algorithm, for the marginal posterior means (Mean) and marginal posterior probabilities of inclusion (PI).

	ZigZag		BPS		Gibbs	
n, p	PI	Mean	PI	Mean	PI	Mean
100, 100	1.75	3.11	1.92	2.28	2.12	4.25
200, 100	5.60	8.75	5.49	11.62	2.76	2.83
400, 100	14.70	25.08	12.47	29.07	3.66	4.14
800, 100	20.18	30.45	17.59	34.38	3.71	3.79
100, 200	1.30	1.38	1.69	1.81	2.54	2.26
200, 200	14.92	20.12	14.85	26.02	2.92	2.64
400, 200	23.30	34.21	21.44	39.83	2.86	2.89
800, 200	43.98	59.56	36.38	62.89	3.14	3.05
100, 400	1.17	1.76	2.56	4.07	1.62	2.15
200, 400	12.70	6.72	14.64	17.22	2.42	2.17
400, 400	43.73	62.07	34.13	61.47	2.41	2.43
800, 400	102.36	152.02	78.26	133.46	2.68	2.78

Table 2: Scenario 2 (General correlation): Relative efficiency (RE) for methods, against a Reversible Jump algorithm, for the marginal posterior means (Mean) and marginal posterior probabilities of inclusion (PI).

	ZigZag		BPS		Gibbs	
n, p	PI	Mean	PI	Mean	PI	Mean
100, 100	0.57	0.83	1.15	1.86	1.05	0.76
200, 100	0.88	0.92	1.71	2.10	1.96	1.27
400, 100	1.39	1.50	2.11	2.32	1.69	1.00
800, 100	4.77	7.30	5.63	9.71	1.99	1.78
100, 200	0.76	1.21	2.03	3.03	1.53	1.37
200, 200	2.28	2.80	4.87	6.84	1.92	1.31
400, 200	4.10	4.51	6.77	8.77	1.30	0.87
800, 200	10.36	13.23	11.12	19.21	1.41	1.12
100, 400	1.09	1.27	2.93	4.01	1.12	0.88
200, 400	3.57	3.37	6.57	7.21	1.24	0.66
400, 400	8.05	8.65	13.45	16.24	1.15	0.83
800, 400	26.93	34.41	28.71	51.59	1.32	1.20

Table 3: Scenario 3 (No correlation): Relative efficiency (RE) for methods, against a Reversible Jump algorithm, for the marginal posterior means (Mean) and marginal posterior probabilities of inclusion (PI).

	ZigZag		BPS		Gibbs	
n, p	PI	Mean	PI	Mean	PI	Mean
100, 100	0.64	0.87	1.25	1.91	1.19	1.04
200, 100	1.40	1.54	2.52	3.12	2.20	1.69
400, 100	2.74	2.79	3.07	3.67	2.22	1.54
800, 100	5.20	8.59	5.64	10.74	1.72	1.50
100, 200	0.76	0.94	1.56	1.77	1.24	1.13
200, 200	2.04	1.63	4.00	3.18	2.07	1.23
400, 200	7.60	12.68	11.90	24.23	1.63	1.38
800, 200	14.26	21.65	16.64	28.09	1.56	1.32
100, 400	1.03	1.44	2.57	4.17	1.17	0.97
200, 400	2.31	1.59	4.18	2.98	1.37	0.84
400, 400	9.38	10.42	14.66	18.60	1.40	1.07
800, 400	30.73	49.94	35.35	71.83	1.42	1.26

One of the attractions of PDMP samplers is that they can be implemented in a way where they only access a small subset of data at each iteration, whilst still targeting the true posterior. We now investigate how these ideas work in the variable selection problem, by comparing the efficiency of three implementations of ZigZag with that of the Gibbs sampler, and see how this depends on the number of observations. These are ZigZag using the full data, ZigZag using subsampling with a global bound, and ZigZag with subsampling control variates (see Bierkens *et al.*, 2019, for details of both subsampling approaches).

Standard application of control variates requires calculation of the gradient at a reference point using the full likelihood. Due to the trans-dimensional nature of variable selection problems, a full gradient calculation is not well defined. For this reason we choose to make use of control variates defined for a fixed model \mathcal{M} where the gradient is well defined. These control variates are only used when the sampler is in this model. Extensions where multiple control variates are used for models with high probability would be straightforward.

Results are shown in Figure 2. Despite our simplistic implementation, these results indicate that Zig-Zag with control variates is becoming increasingly efficient relative to Zig-Zag using the full dataset as the number of samples increases. Furthermore we see evidence of super-efficiency – whilst the computational cost per ESS of the Gibbs sampler is expected to be linear in the number of observations, the relative efficiency plots suggest that this is sub-linear for ZigZag with control variates.




Figure 2: Log-log plots of efficiency, relative to the Gibbs sampler, of different samplers as we vary the number of observations. Plotted are the relative efficiencies for the posterior mean conditional on model \mathcal{M}^* where \mathcal{M}^* corresponds to the true data generated model. The dataset was generated with a 15-dimensional regression parameter $\theta = (1, 1, 0, 0, \dots, 0)$. The methods run are the Zig-Zag applied to the full dataset (zz, black), Zig-Zag with subsampling using global bounds (ss, blue), Zig-Zag with control variates (cv, magenta) and Gibbs sampling (Gibbs, green). All methods were initialised at the location of the control variate. Methods were given the same computational budget, for details see Section D of the Appendix.

5.2 Robust regression

As mentioned in the introduction, a common approach to Bayesian variable selection is to use continuous spike-and-slab priors for each parameter rather than try to sample from the joint posterior of model and parameters. Such an approach is attractive as it enables standard gradient-based samplers, such as Hamiltonian Monte Carlo, to be used. We now compare such an approach, implemented with the popular STAN software (Carpenter *et al.*, 2017), to our PDMP samplers. Our aim is to both investigate the computational efficiencies of the two approaches and to show the differences in posterior that we obtain from these different types of prior. Our comparison is based on a robust linear regression model.

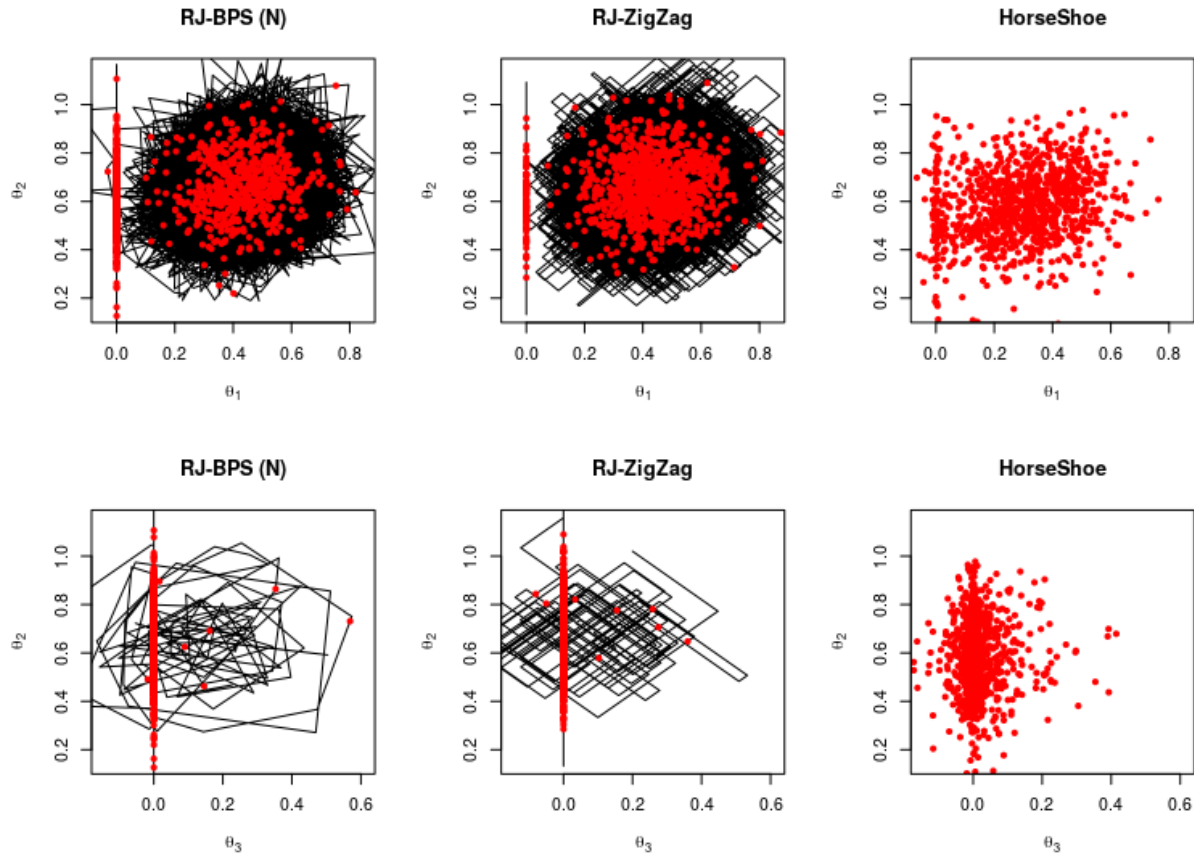


Figure 3: Dynamics of the samplers on a robust regression example with spike and slab or horseshoe prior. The top row shows the posterior for θ_1 and θ_2 , bottom row shows the estimates for θ_2 and θ_3 . The spike and slab distributions are sampled using the reversible jump PDMP samplers with reversible jump parameter 0.6 and refreshment for the BPS methods set to 0.5. All methods are shown with 10^3 samples (red) and the PDMP dynamics are shown in black. Sampling with the Horseshoe prior was implemented in Stan using NUTS. Both Stan and PDMP methods were run for the same computing time. To aid visualisation only the first 30% of the PDMP trajectories are shown.

In particular, we model the errors in our linear regression model as a mixture of Normals with different variances. Thus

$$\mathbf{Y} = \sum_{i=1}^p \mathbf{X}_j \theta_j + \epsilon, \quad \epsilon \sim \frac{1}{2} \mathcal{N}(0, 1) + \frac{1}{2} \mathcal{N}(0, 10^2).$$

The continuous variable selection prior we consider is the regularised horseshoe (Piiroinen

and Vehtari, 2017a,b)

$$\theta_j \sim N(0, \tau^2 \tilde{\lambda}_j), \quad \tilde{\lambda}_j = \frac{c^2 \lambda_j}{c^2 + \tau^2 \lambda_j}, \quad \lambda_j \sim C^+(0, 1)$$

for $j = 1, \dots, p$ where $C^+(0, 1)$ denotes the half-Cauchy distribution for the standard deviation λ_j . The regularised horseshoe is a variation of the horseshoe prior (Carvalho *et al.*, 2010) that offers a continuous approximation of a Dirac spike and slab where the slab is a Normal distribution with finite variance c^2 . The hyper-parameter τ controls the global shrinkage of the variables towards zero. In Carvalho *et al.* (2010) it was shown that for standard linear regression the optimal choice for a fixed value of τ is $\tau_0 = \sigma \frac{p_0}{(p-p_0)\sqrt{n}}$ where p_0 is the number of nonzero variables in the sparse model and σ is the noise variance. In line with this, we compare the regularised horseshoe prior with fixed hyper-parameter $\tau_0 = \frac{p_0}{(p-p_0)\sqrt{n}}$ against the spike-and-slab prior

$$\theta_j \sim \frac{p_0}{p} N(0, c^2) + \left(1 - \frac{p_0}{p}\right) \delta_0$$

for $j = 1, \dots, p$. MCMC for the model using a horseshoe prior was performed by Stan's implementation of NUTS (Hoffman and Gelman, 2014).

We first compare the variable selection dynamics for a simple model with $p = 4$ variables, $n = 120$ observations and regression parameter $\boldsymbol{\theta} = (0.5, 0.5, 0, 0)^T$. The covariate values and residuals were generated as independent draws from a standard Normal and the prior expected model size is set to $p_0 = 1$. Example output for the PDMP samplers and the Stan implementation is shown in Figure 3. The posteriors show the horseshoe prior replicating the effect of the spike-and-slab through shrinking the coefficients towards zero, but it is not able to give exact zeros.

We now compare the reversible jump PDMP methods in terms of their sampling efficiency for a higher dimensional problem. The dataset is generated for $p = 200$ variables and $n = 100$ observations with regression parameter $\boldsymbol{\theta} = (2, 2, 2, 2, 0, 0, 0, \dots, 0)^T$. The covariates for each observation were drawn from an AR(1) process with lag-1 correlation of 0.5. The residuals were generated from a standard Cauchy distribution.

We ran Stan with the default settings for a burnin of 1000 iterations and then for 16, 32, 64, ..., 2048 iterations. We ran the reversible jump PDMP samplers for the same wall clock time for both burnin and subsequent iterations. The experiment was repeated 35 times and the resulting boxplots of the posterior mean of θ_1 are shown in Figure 4.




Figure 4: Sampling efficiency for reversible jump PDMP vs Stan for the robust regression example. The top figure is boxplots of the posterior mean of θ_1 for increasing computational budget, with outliers from the sampler removed for visualisation purposes. These removed outliers correspond to times that the sampler has become stuck in a local mode where $\theta_1 = 0$. The subplot shows the full results including outliers from the sampler. Both the Stan and reversible jump PDMP methods converge to similar solutions despite having different priors. The bottom figure shows the number of times that the samplers did not find the global mode.

For the same computational budget, the reversible jump PDMP methods are able to attain better performance. However, it is also apparent from this simulation that the BPS reversible jump PDMP methods are more sensitive to local modes. It is unclear if this is an artifact of the sampler or due to the choice of refreshment rate (0.6). This sensitivity is seen to diminish as the sampler is run for longer.

The predictive ability of the methods is compared in Figure 5. Here an additional $n = 100$ observations were drawn from the same model and these were used as a hold-out dataset to validate the posterior predictive ability. The predictive ability is defined in terms of the Monte Carlo estimate of the mean square prediction performance $\frac{1}{100} \sum_{i=1}^{100} (y_i - \mathbf{x}_i^T \boldsymbol{\theta})^2$ where $\boldsymbol{\theta}$ is replaced by the samples generated by either Stan or samples given by a discretisation

of time for the reversible jump PDMP samplers. The reversible jump PDMP samplers all give the same predictive performance for large iteration numbers while Stan, which uses a horseshoe prior, performs slightly worse.




Figure 5: Predictive ability of reversible jump PDMP vs Stan for the robust regression example. The predictive ability is measured by Monte Carlo estimates of the mean square predictive performance.

6 Discussion

We have shown how PDMP samplers can be extended so that they can sample from the joint posterior over model and parameters in variable selection problems. There are a number of open challenges that stem from this work. First, as with any MCMC algorithm, the reversible jump PDMP samplers have tuning parameters. Our simulation results were based on choosing these after empirically evaluating the performance of the samplers on one simple problem. Whilst the samplers mixed well, it is likely that better mixing could be achieved if more informed choices of tuning parameters were made, and theory for guiding such choices is needed.

Second, the form of our reversible jump PDMP samplers is based on particular features of the variable selection problem. In other model choice settings, different trans-dimensional moves may be needed. The theory we developed should be able to be adapted to give rules for choosing rates of such moves. Also our trans-dimensional moves are reversible, that is they balance probability flow from model i to model j by the flow of probability from model j to model i – it would be interesting to see if non-reversible trans-dimensional moves could

be constructed.

Third, it is likely that the reversible jump PDMP samplers will still struggle in situations where the posterior is multi-modal with well separated modes. For such cases it would be interesting to try and incorporate ideas such as tempering (Marinari and Parisi, 1992) to allow for better mixing across modes. Also given the close links between PDMP samplers and Hamiltonian Monte Carlo, it would be interesting to see if the ideas we have presented for trans-dimensional jumps could be adapted to be used with Hamiltonian Monte Carlo algorithms.

References

- Bierkens, J. and Roberts, G. (2017). A piecewise deterministic scaling limit of lifted Metropolis–Hastings in the Curie–Weiss model. *The Annals of Applied Probability* **27**(2), 846–882.
- Bierkens, J., Bouchard-Côté, A., Doucet, A., Duncan, A. B., Fearnhead, P., Lienart, T., Roberts, G. and Vollmer, S. J. (2018). Piecewise deterministic Markov processes for scalable Monte Carlo on restricted domains. *Statistics & Probability Letters* **136**, 148–154.
- Bierkens, J., Fearnhead, P., Roberts, G. *et al.* (2019). The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *The Annals of Statistics* **47**(3), 1288–1320.
- Bierkens, J., Grazi, S., Kamatani, K. and Roberts, G. (2020). The boomerang sampler. *arXiv:2006.13777* .
- Bottolo, L. and Richardson, S. (2010). Evolutionary stochastic search for bayesian model exploration. *Bayesian Analysis* **5**, 583–618.
- Bouchard-Côté, A., Vollmer, S. J. and Doucet, A. (2018). The bouncy particle sampler: A nonreversible rejection-free Markov chain Monte Carlo method. *Journal of the American Statistical Association* **113**(522), 855–867.
- Buchholz, A., Ahfock, D. and Richardson, S. (2019). Distributed computation for marginal likelihood based model choice. *arxiv.1910.04672* .
- 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. *Journal of Statistical Software* **76**(1).
- Carvalho, C. M., Polson, N. G. and Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika* **97**(2), 465–480.

- Chipman, H. A., George, E. I. and McCulloch, R. E. (2001). The practical implementation of Bayesian model selection (with discussion). *Model Selection (P. Lahiri ed.)* , 65–134.
- Davis, M. (1993). *Markov Models and Optimization*. Chapman Hall.
- de Valpine, P., Turek, D., Paciorek, C., Anderson-Bergman, C., Temple Lang, D. and Bodik, R. (2017). Programming with models: writing statistical algorithms for general model structures with NIMBLE. *Journal of Computational and Graphical Statistics* **26**, 403–413.
- de Valpine, P., Paciorek, C., Turek, D., Michaud, N., Anderson-Bergman, C., Obermeyer, F., Wehrhahn Cortes, C., Rodríguez, A., Temple Lang, D. and Paganin, S. (2020). *NIMBLE User Manual*. R package manual version 0.10.0.
- Diaconis, P., Holmes, S. and Neal, R. M. (2000). Analysis of a nonreversible Markov chain sampler. *Annals of Applied Probability* **10**, 726–752.
- Fearnhead, P., Bierkens, J., Pollock, M. and Roberts, G. O. (2018). Piecewise deterministic Markov processes for continuous-time Monte Carlo. *Statistical Science* **33**(3), 386–412.
- George, E. I. and McCulloch, R. E. (1993). Variable selection via gibbs sampling. *Journal of the American Statistical Association* **88**(423), 881–889.
- Geweke, J. (1996). Variable selection and model comparison in regression. In: *Bayesian Statistics 5*, Oxford University Press, 609–620.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**(4), 711–732.
- Hoffman, M. D. and Gelman, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* **15**(1), 1593–1623.
- Ishwaran, H. and Rao, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics* **33**(2), 730–773.
- Kuo, L. and Mallick, B. K. (1998). Variable selection for regression models. *Sankhya Series B* **60**, 65–81.
- Marinari, E. and Parisi, G. (1992). Simulated tempering: a new Monte Carlo scheme. *EPL (Europhysics Letters)* **19**(6), 451.
- Markovic, J. and Sepehri, A. (2018). Bouncy hybrid sampler as a unifying device. *arXiv:1802.04366* .

- Michel, M., Kapfer, S. C. and Krauth, W. (2014). Generalized event-chain Monte Carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps. *The Journal of Chemical Physics* **140**(5), 054116.
- Michel, M., Durmus, A. and S en ecal, S. (2020). Forward event-chain Monte Carlo: Fast sampling by randomness control in irreversible Markov chains. *Journal of Computational and Graphical Statistics* , To appear.
- Mitchell, T. J. and Beauchamp, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association* **83**(404), 1023–1032.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In: *Handbook of Markov chain Monte Carlo* (eds. A. Brooks, A. Gelman, G. L. Jones and X. Meng), Chapman & Hall, 113–162.
- Peters, E. A. and de With, G. (2012). Rejection-free Monte Carlo sampling for general potentials. *Physical Review E* **85**(2), 026703.
- Piironen, J. and Vehtari, A. (2017a). On the Hyperprior Choice for the Global Shrinkage Parameter in the Horseshoe Prior. volume 54 of *Proceedings of Machine Learning Research*, PMLR, 905–913.
- Piironen, J. and Vehtari, A. (2017b). Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electron. J. Statist.* **11**(2), 5018–5051.
- Polson, N. G., Scott, J. G. and Windle, J. (2013). Bayesian inference for logistic models using P olya–Gamma latent variables. *Journal of the American statistical Association* **108**(504), 1339–1349.
- Smith, M. and Kohn, R. (1996). Nonparametric regression using bayesian variable selection. *Journal of Econometrics* **75**(2), 317 – 343.
- Song, Q., Sun, Y., Ye, M. and Liang, F. (2020). Extended stochastic gradient Markov chain Monte Carlo for large-scale Bayesian variable selection. *Biometrika* Asaa029.
- Vanetti, P., Bouchard-C ot e, A., Deligiannidis, G. and Doucet, A. (2017). Piecewise-Deterministic Markov Chain Monte Carlo. *arXiv:1707.05296* .
- Wang, S., Nan, B., Rosset, S. and Zhu, J. (2011). Random lasso. *Ann. Appl. Stat.* **5**(1), 468–485.
- Wu, C. and Robert, C. P. (2020). Coordinate sampler: a non-reversible Gibbs-like MCMC sampler. *Statistics and Computing* **30**(3), 721–730.
- Yang, Y., Wainwright, M. J. and Jordan, M. I. (2016). On the computational complexity of high-dimensional bayesian variable selection. *Ann. Statist.* **44**(6), 2497–2532.

Zanella, G. and Roberts, G. (2019). Scalable importance tempering and bayesian variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **81**(3), 489–517.

Supplementary Material: Reversible Jump PDMP Samplers for Variable Selection

A Proofs

A.1 Proof of Theorem 1

Recall theorems 34.15 and 34.19 from Davis (1993):

Theorem 4 (34.15). *Suppose that μ is a stationary distribution for (\mathbf{x}_t) . Then there exists a finite measure σ on Γ (the active boundary) such that for all $t > 0$:*

$$\mathbb{E}_\mu \frac{1}{t} \int_0^t f(\mathbf{x}_{s-}) dp_s^* = \int_\Gamma f(\mathbf{x}) \sigma(d\mathbf{x})$$

where p^* is the counting process that counts the number of jumps for the boundary

Theorem 5 (34.19). *Suppose μ is a stationary distribution for (\mathbf{x}_t) and let σ be the corresponding boundary measure defined above. Then:*

$$\int_E \mathfrak{A}f(\mathbf{x}) \mu(d\mathbf{x}) + \int_\Gamma \mathfrak{C}f(\mathbf{z}) \sigma(d\mathbf{z}) = 0$$

where \mathfrak{A} is the extended generator and $\mathfrak{C}f = (Qf - f)$.

Remark 2. *The construction of PDMPs with an active boundary we presented differs slightly from the one found in (Davis, 1993, p. 57). However our process still fits the definition of Davis (1993): one simply needs to separate each \mathcal{X}^k in two at the boundary and consider the two parts as two separated spaces.*

From the expression of the extended generator given in (26.14) of Davis (1993), it follows that

$$\sum_i \int_{E_i} \mathcal{A}_i f d\mu + \sum_{(i,j) \in \mathcal{T}} \int_{E_j} \beta_{i,j}(\mathbf{z}) \left(\int_{g_{i,j}^{-1}(\{\mathbf{z}\})} [f(\mathbf{z}') - f(\mathbf{z})] Q_{i,j}(\mathbf{z}'|\mathbf{z}) d\mathbf{z}' \right) \mu(\mathbf{z}) d\mathbf{z}$$

corresponds to the first term of Theorem 5. It is now left to show that

$$\sum_{(i,j) \in \mathcal{T}} p_{i,j} \int_{\Gamma_{i,j}} [f(g_{i,j}(\mathbf{z})) - f(\mathbf{z})] \mu(\mathbf{z}) |\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle| d\mathbf{z}$$

corresponds to the boundary term of Theorem 5. Following Davis (1993) notations, for a fixed (i, j) in \mathcal{T} , the operator $\mathfrak{C} = (Qf - f)$ where Q is the jump kernel at the boundary has the following expression:

$$\mathfrak{C}f = p_{i,j} [f(g_{i,j}(\mathbf{z})) - f(\mathbf{z})].$$

Finally, we need to find an expression for measure σ . We start from the definition in Theorem 4:

$$\mathbb{E}_\mu \frac{1}{t} \int_0^t f(\mathbf{z}_{s-}) dp_s^* = \int_{\Gamma_{i,j}} f(\mathbf{z}) \sigma(d\mathbf{z}),$$

where p_s^* is the counting process that counts the number of times we hit $\Gamma_{i,j}$. This holds for all t , and we will derive σ by considering the value of the left-hand side as $t \rightarrow 0$. Intuitively the idea is that, in this limit, the impact of the events will be negligible and we can evaluate the left-hand side by considering a process where \mathbf{z}_0 is simulated from μ but then is purely deterministic, with dynamics given by the deterministic part of the PDMP.

As indicator functions of compact sets separate measures on $\Gamma_{i,j}$ we can restrict ourselves to these functions. Let K be a compact of $\Gamma_{i,j}$ and f its characteristic function, $f = 1_K$. We also introduce some additional notation. Let $\tilde{\mathbf{z}}_t$ be a deterministic process that follows the dynamics of the deterministic part of the PDMP, and let \tilde{p}_t^* be the associated process that counts the number of times $\tilde{\mathbf{z}}_t$ hits $\Gamma_{i,j}$. For our PDMP let $\Omega_0(t)$ be the set of paths for which no events occur by time t , and $\Omega_1(t)$ the set of paths with at least one event prior to time t . We can write

$$\int_0^t 1_K(\mathbf{z}_{s-}) dp_s^* = \int_0^t 1_K(\mathbf{z}_{s-}) 1_{\Omega_0(t)} dp_s^* + \int_0^t 1_K(\mathbf{z}_{s-}) 1_{\Omega_1(t)} dp_s^*,$$

that is split the integral into two, with the first being the contribution from paths with no event and the second the contribution of paths with at least one event.

We now introduce two sets. Let $K_t = \{\mathbf{z} \in E | \mathbf{z}_s \in K \text{ with } \mathbf{z}_0 = \mathbf{z} \text{ and } s \in [0, t]\}$, and $\tilde{K}_t = \{\mathbf{z} \in E_i | \tilde{\mathbf{z}}_s \in K \text{ with } \tilde{\mathbf{z}}_0 = \mathbf{z} \text{ and } s \in [0, t]\}$. The first is the set of all possible starting points for our PDMP process for which it is possible to hit K by time t , the latter is the same set but for the purely deterministic process (equivalently, all PDMP paths with no events by time t). As K is compact so is \tilde{K}_t . Furthermore as the velocity space is bounded and K is compact so is K_t . This, together with our assumption on the event rate means that we can upper bound the event rate for both $\mathbf{z} \in K_t$ and $\mathbf{z} \in \tilde{K}_t$, and denote such an upper bound by λ^+ .

Now $\mathbb{E}_\mu \int_0^t 1_K(\mathbf{z}_{s-}) 1_{\Omega_1(t)} dp_s^*$ can be bounded as for the integral to be non-zero we need $\mathbf{z}_0 \in K_t$ and the number of times the PDMP hits $\Gamma_{i,j}$ is at most twice the number of events. Thus an upper bound is $2\lambda^+ t \Pr_\mu(\mathbf{z}_0 \in K_t)$. As the volume of K_t tends to 0 as $t \rightarrow 0$, this upper bound is $o(t)$.

To bound the contribution from the other integral, we notice that we can bound this from above by the the PDMP process without events, and from below by using our bound on the event rate. Thus

$$\mathbb{E}_\mu \int_0^t 1_K(\tilde{\mathbf{z}}_{s-}) d\tilde{p}_s^* \geq \mathbb{E}_\mu \int_0^t 1_K(\mathbf{z}_{s-}) 1_{\Omega_1(t)} dp_s^* \geq \exp\{-t\lambda^+\} \mathbb{E}_\mu \int_0^t 1_K(\tilde{\mathbf{z}}_{s-}) d\tilde{p}_s^*.$$

Thus

$$\mathbb{E}_\mu \int_0^t 1_K(\mathbf{z}_{s-}) 1_{\Omega_1(t)} dp_s^* = \mathbb{E}_\mu \int_0^t 1_K(\tilde{\mathbf{z}}_{s-}) d\tilde{p}_s^* + o(t).$$

Thus, for any compact subset of $\Gamma_{i,j}$, K , we have

$$\int_{\Gamma_{i,j}} 1_K(\mathbf{z})\sigma(d\mathbf{z}) = \lim_{t \rightarrow 0} \mathbb{E}_\mu \frac{1}{t} \int_0^t 1_K(\mathbf{z}_{s^-}) dp_s^* = \lim_{t \rightarrow 0} \mathbb{E}_\mu \int_0^t 1_K(\tilde{\mathbf{z}}_{s^-}) d\tilde{p}_s^*.$$

The final step is to obtain $\lim_{t \rightarrow 0} \mathbb{E}_\mu \int_0^t 1_K(\tilde{\mathbf{z}}_{s^-}) d\tilde{p}_s^*$, which comes from the following result for solutions of ordinary differential equations:

Theorem 6. *Let E be a subset of \mathbb{R}^d , μ a measure with a density with respect to the Lebesgue measure of E and let X be the vector field associated to the order ordinary differential equation*

$$\frac{d(\mathbf{z}_t)}{dt} = X(\mathbf{z}_t).$$

Let H be an hyperplane of E with normal \mathbf{n}_H and $t_(\mathbf{z})$ be the first hitting time of a trajectory with H starting from \mathbf{z} . Assuming there exists a measure σ on H such that for all $f : H \rightarrow \mathbb{R}$:*

$$\lim_{t \rightarrow 0} \frac{1}{t} \int_E f(\mathbf{z}_{t_*(\mathbf{z}_0)}) 1_{t_*(\mathbf{z}_0) \leq t} \mu(\mathbf{z}_0) d\mathbf{z}_0 = \int_H f(\mathbf{z}) \sigma(\mathbf{z}) d\mathbf{z}$$

then σ has a density with respect to the Lebesgue measure of H and

$$\sigma(d\mathbf{z}) = \mu(\mathbf{z}) |\langle \mathbf{z}, \mathbf{n}_H \rangle| d\mathbf{z}$$

For completeness, this result is proven in Appendix A.6. The intuition behind the $|\langle \mathbf{z}, \mathbf{n}_H \rangle|$ factor is that ‘area’ of space that hits \mathbf{z} within an infinitesimal time interval dt is $|\langle \mathbf{z}, \mathbf{n}_H \rangle| dt$.

Applying this theorem to the deterministic part of our PDMP we get

$$\sigma(\mathbf{z}) d\mathbf{z} = \mu(\mathbf{z}) |\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle| d\mathbf{z}$$

where $\mathbf{n}_{i,j}$ is the normal the the boundary $\Gamma_{i,j}$, which concludes the proof.

A.2 Proof of Theorem 2

Proof. The core of the proof is rewriting (2) of Theorem 1 in a more suitable way. First, notice that

$$\int_{\Gamma_{i,j}} f(g_{i,j}(\mathbf{z})) \mu(\mathbf{z}) |\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle| d\mathbf{z} = \int_{E_j} f(\mathbf{z}) \bar{\nu}_{i,j}(\mathbf{z}) d\mathbf{z}$$

by definition of the pushforward measure $\bar{\nu}_{i,j}$. Second, since $Q_{i,j}$ integrates to 1, we get:

$$\int_{E_j} \beta_{i,j}(\mathbf{z}) \int_{g_{i,j}^{-1}(\{\mathbf{z}\})} f(\mathbf{z}') Q_{i,j}(\mathbf{z}'|\mathbf{z}) d\mathbf{z}' \mu(\mathbf{z}) d\mathbf{z} = \int_{E_j} \beta_{i,j}(\mathbf{z}) f(\mathbf{z}) \mu(\mathbf{z}) d\mathbf{z}.$$

Third, using the definitions of $G_{i,j}$

$$\int_{g_{i,j}^{-1}(\{\mathbf{z}\})} f(\mathbf{z}') Q_{i,j}(\mathbf{z}'|\mathbf{z}) d\mathbf{z}' = \int_U f(G_{i,j}(\alpha, \mathbf{z})) Q_{i,j}(\alpha|\mathbf{z}) d\alpha$$

Finally, using a change the change of variable from $\mathbf{z}' = (\mathbf{x}', \mathbf{v}')$ to α, \mathbf{z} defined as $\mathbf{z}' = G_{i,j}(\alpha, \mathbf{z})$, and the definition of $\nu_{i,j}$:

$$\int_{\Gamma_{i,j}} f(\mathbf{z}') \nu(\mathbf{z}') |\langle \mathbf{v}', \mathbf{n}_{i,j} \rangle| d\mathbf{z}' = \int_{E_j} \int_U f(G_{i,j}(\alpha, \mathbf{z})) \nu_{i,j}(G(\alpha, \mathbf{z})) |J_{G_{i,j}}(\alpha, \mathbf{z})| d\alpha d\mathbf{z}.$$

Thus (2) of Theorem 1, evaluated for $\pi = \nu$, can be restated as:

$$\begin{aligned} 0 &= \sum_i \int_{E_i} \mathcal{A}_i f d\nu + \sum_{(i,j) \in \mathcal{T}} \int_{E_j} [p_{i,j} \bar{\nu}_{i,j}(\mathbf{z}) - \beta_{i,j}(\mathbf{z}) \nu(\mathbf{z})] f(\mathbf{z}) d\mathbf{z} \\ &+ \sum_{(i,j) \in \mathcal{T}} \int_{E_j} \int_U [\beta_{i,j}(\mathbf{z}) Q_{i,j}(\alpha|\mathbf{z}) \nu(\mathbf{z}) - p_{i,j} \nu_{i,j}(G_{i,j}(\alpha, \mathbf{z})) |J_{G_{i,j}}(\alpha, \mathbf{z})|] f(G_{i,j}(\alpha, \mathbf{z})) d\alpha d\mathbf{z} \end{aligned} \quad (12)$$

By construction, for any function $f \in \mathcal{F}$:

$$\sum_i \int_{E_i} \mathcal{A}_i f d\nu = 0.$$

Conditions 3 and 4 of the theorem clearly imply that

$$\sum_{(i,j) \in \mathcal{T}} \int_{E_j} [p_{i,j} \bar{\nu}_{i,j}(\mathbf{z}) - \beta_{i,j}(\mathbf{z}) \nu(\mathbf{z})] f(\mathbf{z}) d\mathbf{z} = 0$$

and

$$\sum_{(i,j) \in \mathcal{T}} \int_{E_j} \int_U [\beta_{i,j}(\mathbf{z}) Q_{i,j}(\alpha|\mathbf{z}) \nu(\mathbf{z}) - p_{i,j} \nu_{i,j}(G_{i,j}(\alpha, \mathbf{z})) |J_{G_{i,j}}(\alpha, \mathbf{z})|] f(G_{i,j}(\alpha, \mathbf{z})) d\alpha d\mathbf{z} = 0$$

respectively, which concludes the proof. \square

A.3 Proof of Theorem 3

Proof. The proof is largely the same as the proof of Theorem 2. The only change is how to treat the change of variable $\mathbf{z}' = G_{i,j}(\alpha, \mathbf{z})$. First, notice that $G_{i,j}$ leaves the position invariant since $g_{i,j}$ leaves the position invariant. Then, since the velocity space is discrete and $\alpha \mapsto G_{i,j}(\alpha, \mathbf{z})$ is a one to one mapping for a fixed $\mathbf{z} \in E_j$:

$$\begin{aligned} \int_{\Gamma_{i,j}} f(\mathbf{z}') \nu_{i,j}(\mathbf{z}') |\langle \mathbf{v}', \mathbf{n}_{i,j} \rangle| d\mathbf{z}' &= \int_{\mathcal{X}^j} \sum_{\mathbf{v}' \in \mathcal{V}^i} f(\mathbf{x}', \mathbf{v}') \nu_{i,j}(\mathbf{x}', \mathbf{v}') d\mathbf{x}' \\ &= \int_{\mathcal{X}^j} \sum_{\mathbf{v} \in E_j} \sum_{\alpha \in U} f(G_{i,j}(\alpha, (\mathbf{x}, \mathbf{v}))) \nu_{i,j}(G_{i,j}(\alpha, (\mathbf{x}, \mathbf{v}))) d\mathbf{x}. \end{aligned}$$

which we can write in integral form as

$$\int_{\Gamma_{i,j}} f(\mathbf{z}') \nu_{i,j}(\mathbf{z}') |\langle \mathbf{v}', \mathbf{n}_{i,j} \rangle| d\mathbf{z}' = \int_{E_j} \int_U f(G_{i,j}(\alpha, (\mathbf{x}, \mathbf{v}))) \nu_{i,j}(G_{i,j}(\alpha, (\mathbf{x}, \mathbf{v}))) d\alpha d\mathbf{z}.$$

□

A.4 Proof of Proposition 1

Proof. Here:

$$\mathcal{V}^i = \{\mathbf{v} \in \mathbb{R}^{|\gamma^i|} \text{ such that } \|\mathbf{v}\| = 1\}$$

For $\mathbf{z} \in E_i$:

$$\nu(\mathbf{z}) = \pi_i(\mathbf{x}) \frac{1}{A_{\text{sphere}}(|\gamma^i|)}$$

with $A_{\text{sphere}}(|\gamma^i|) = \frac{\Gamma(\frac{|\gamma^i|}{2})}{2\pi^{\frac{|\gamma^i|}{2}}}$ the area of the unit sphere of $\mathbb{R}^{|\gamma^i|}$.

In this case we can define $G_{i,j} : (-1, 1) \times E_j \rightarrow E_i$ such that $\mathbf{z}' = (\mathbf{x}', \mathbf{v}') = G_{i,j}(\alpha, \mathbf{z})$, where $\mathbf{z} = (\mathbf{x}, \mathbf{v})$, as $\mathbf{x}' = \mathbf{x}$ and

$$\mathbf{v}' = \sqrt{1 - \alpha^2} \mathbf{v} + \alpha \mathbf{n}_{i,j}.$$

Let $B \subset \mathcal{V}^j$

$$\begin{aligned} \int_B \bar{\nu}_{i,j}(\mathbf{x}, \mathbf{v}) d\mathbf{v} &= \pi_i(\mathbf{x}) \int_{g_{i,j}^{-1}(B)} |\langle \mathbf{v}', \mathbf{n}_{i,j} \rangle| \frac{1}{A_{\text{sphere}}(|\gamma^i|)} d\mathbf{v}' \\ &= \pi_i(\mathbf{x}) \frac{1}{A_{\text{sphere}}(|\gamma^i|)} \int_{[-1,1] \times B} |\langle \sqrt{1 - \alpha^2} \mathbf{v} + \alpha \mathbf{n}_{i,j}, \mathbf{n}_{i,j} \rangle| |J_{G_{i,j}}| d\alpha d\mathbf{v} \\ &= 2\pi_i(\mathbf{x}) \frac{1}{A_{\text{sphere}}(|\gamma^i|)} \int_{(0,1] \times B} \alpha |J_{G_{i,j}}| d\alpha d\mathbf{v}, \end{aligned}$$

where the last inequality uses the symmetry of the integral with respect to α , and that $\langle \mathbf{v}, \mathbf{n}_{i,j} \rangle = 0$ by definition of velocities in E_j .

The determinant of $G_{i,j}$ must be carefully computed since the velocities lives in the sphere and not on the full vector space. We get:

$$|J_{G_{i,j}}(\alpha, \mathbf{v})| = \sqrt{1 - \alpha^2}^{|\gamma^j| - 2}.$$

Therefore for $|\gamma^j| > 0$:

$$\bar{\nu}_{i,j}(\mathbf{z}) = \pi_i(\mathbf{x}) \frac{2}{A_{\text{sphere}}(|\gamma^i|)} \cdot \frac{1}{|\gamma^j|}$$

Hence for $|\gamma^j| > 0$:

$$\beta_{i,j} = p_{i,j} \frac{\pi_i(\mathbf{x})}{\pi_j(\mathbf{x})} \frac{2A_{\text{sphere}}(|\gamma^j|)}{A_{\text{sphere}}(|\gamma^i|)} \frac{1}{|\gamma^j|}, \quad (13)$$

and

$$Q_{i,j}(\alpha|\mathbf{z}) = \frac{|\alpha||\gamma^j|\sqrt{1-\alpha^2}^{|\gamma^j|-2}}{2} \text{ for } \mathbf{z} \in E_j \text{ and } \alpha \in (-1, 1) \quad (14)$$

For $|\gamma^j| = 0$, BPS and ZigZag are equivalent, thus we use ZigZag rates. \square

A.5 Proof of Proposition 2

Proof. Here:

$$\mathcal{V}^i = \mathbb{R}^{|\gamma^i|}$$

Therefore, for $\mathbf{z} \in E_i$:

$$\nu(\mathbf{z}) = \pi_i(\mathbf{x}) \frac{1}{(2\pi)^{\frac{|\gamma^i|}{2}}} e^{-\frac{1}{2}\|\mathbf{v}\|^2}$$

For $(i, j) \in \mathcal{T}$ with $|\gamma^j| > 0$, we can define $G_{i,j} : (-1, 1) \times E_j \rightarrow E_i$ such that $\mathbf{z}' = (\mathbf{x}', \mathbf{v}') = G_{i,j}(\alpha, \mathbf{z})$, where $\mathbf{z} = (\mathbf{x}, \mathbf{v})$, as $\mathbf{x}' = \mathbf{x}$ and

$$\mathbf{v}' = \mathbf{v} + \alpha \mathbf{n}_{i,j}.$$

Clearly, we have $|J_{G_{i,j}}(\alpha, \mathbf{z})| = 1$. Let $B \subset \mathcal{V}^j$

$$\begin{aligned} \int_B \bar{\nu}_{i,j}(\mathbf{x}, \mathbf{v}) d\mathbf{v} &= \pi_i(\mathbf{x}) \int_{g_{i,j}^{-1}(B)} |\langle \mathbf{v}', \mathbf{n}_{i,j} \rangle| \frac{1}{(2\pi)^{\frac{|\gamma^i|}{2}}} e^{-\frac{1}{2}\|\mathbf{v}\|^2} d\mathbf{v}' \\ &= \pi_i(\mathbf{x}) \frac{1}{(2\pi)^{\frac{|\gamma^i|}{2}}} \int_{\mathbb{R} \times B} |\langle \mathbf{v} + \alpha \mathbf{n}_{i,j}, \mathbf{n}_{i,j} \rangle| e^{-\frac{1}{2}\|\mathbf{v} + \alpha \mathbf{n}_{i,j}\|^2} |J_{G_{i,j}}| d\alpha d\mathbf{v} \\ &= \pi_i(\mathbf{x}) \frac{1}{(2\pi)^{\frac{|\gamma^i|}{2}}} \int_{\mathbb{R} \times B} |\alpha| e^{-\frac{1}{2}(\alpha^2 + \|\mathbf{v}\|^2)} d\alpha d\mathbf{v} \\ &= \pi_i(\mathbf{x}) 2 \frac{1}{(2\pi)^{\frac{|\gamma^i|}{2}}} \int_B e^{-\frac{1}{2}(\|\mathbf{v}\|^2)} d\mathbf{v}. \end{aligned}$$

Therefore for $|\gamma^j| > 0$:

$$\bar{\nu}_{i,j}(\mathbf{z}) = \pi_i(\mathbf{x}) 2 \frac{1}{(2\pi)^{\frac{|\gamma^i|}{2}}} e^{-\frac{1}{2}(\|\mathbf{v}\|^2)},$$

and

$$\beta_{i,j} = p_{i,j} \frac{\pi_i(\mathbf{x})}{\pi_j(\mathbf{x})} \frac{2}{\sqrt{2\pi}} \quad (15)$$

Furthermore, $|J_{G_{i,j}}| = 1$ thus:

$$Q_{i,j}(\alpha|\mathbf{z}) = 2|\alpha|e^{-\frac{1}{2}\alpha^2} \text{ for } \mathbf{z} \in E_j \text{ and } \alpha \in \mathbb{R} \quad (16)$$

□

A.6 Proof of Theorem 6

Proof. Let $Z(t, \mathbf{z})$ be the flow associated to the ODE. Fix \mathbf{z}_0 in H such that $a = \langle X(\mathbf{z}_0), \mathbf{n}_H \rangle \neq 0$. In order to show that σ has the required form it is enough to prove that

$$\lim_{r \rightarrow 0} \frac{\sigma(H \cap B(\mathbf{z}_0, r))}{\text{Volume}_H(H \cap B(\mathbf{z}_0, r))} = |a|\mu(\mathbf{z}_0)$$

where $B(\mathbf{z}_0, r)$ is the ball of radius r centered on \mathbf{z}_0 , $H \cap B(\mathbf{z}_0, r)$ is the slice of this ball that lies on H , and Volume_H is the Lebesgue measure of the hyperplane H .

For r and $t > 0$, let

$$E_{r,t} = \{\mathbf{z} \in E : 0 \leq t_*(\mathbf{z}) \leq t, Z(t_*(\mathbf{z}), \mathbf{z}) \in B(\mathbf{z}_0, r) \cap H\}$$

We can re-write the definition of σ in terms of $E_{r,t}$, from which we are reduced to show that

$$\lim_{r \rightarrow 0} \lim_{t \rightarrow 0} \frac{\frac{1}{t}\mu(E_{r,t})}{\text{Volume}_H(H \cap B(\mathbf{z}_0, r))} = |a|\mu(\mathbf{z}_0)$$

The idea of the proof is that to calculate $\mu(E_{r,t})$ we need to find the volume of $E_{r,t}$ as $r, t \rightarrow 0$, and we will do this by bounding $E_{r,t}$ from both above and below by cylinders. We can approximate these cylinders by the set $E_{r,t}$ that we would obtain if $X(\mathbf{z}) = X(\mathbf{z}_0)$ (i.e. the derivative was constant). The following lemma enables us to quantify the order of error for such an approximation.

Lemma 1. 1. Let $r_0 > 0$. When $t > 0$ is small enough, $Z(s, \mathbf{z}) \in B(\mathbf{z}_0, r_0)$ for all $\mathbf{z} \in B(\mathbf{z}_0, r_0/2)$ and $s \in [-t, t]$.

2. Let $e(s, \mathbf{z})$ be the function defined by $e(s, \mathbf{z}) = Z(s, \mathbf{z}) - (\mathbf{z} + sX(\mathbf{z}_0))$. Then:

$$\lim_{(s, \mathbf{z}) \rightarrow (0, \mathbf{z}_0), s \neq 0} \frac{1}{s} e(s, \mathbf{z}) = 0.$$

Proof. For part (1): Z is continuous as a function from $\mathbb{R} \times E$ to E . Therefore for t small enough, $Z(s, z) \in B(\mathbf{z}_0, r_0)$ for all $\mathbf{z} \in B(\mathbf{z}_0, r_0/2)$ and $s \in [-t, t]$.

For part (2): We define

$$C_{\mathbf{z}_0, r} = \sup_{\mathbf{z} \in B(\mathbf{z}_0, r)} \|X(Z(s, \mathbf{z})) - X(\mathbf{z}_0)\|.$$

We clearly have $C_{\mathbf{z}_0, r} < \infty$ and $C_{\mathbf{z}_0, r} \xrightarrow{r \rightarrow 0} 0$. We set $r_0 = 2\|\mathbf{z} - \mathbf{z}_0\|$, and using the mean value theorem with the function $e_{\mathbf{z}}(s) = e(s, \mathbf{z})$ whose derivative is $X(Z(s, \mathbf{z})) - X(\mathbf{z}_0)$ with the fact that $e(0, \mathbf{z}) = 0$ implies

$$\|e(s, \mathbf{z})\| \leq sC_{\mathbf{z}_0, 2\|\mathbf{z} - \mathbf{z}_0\|}$$

which concludes the proof. \square

The following two lemmas then give, respectively, the cylinders that lower and upper bound $E_{r,t}$.

Lemma 2. *For $\mathbf{z} \in B(\mathbf{z}_0, r) \cap H$ and $s \in [-t, 0]$, there exists a constant $C > 0$ and function $o(r) \geq 0$ with $\lim_{r \rightarrow 0} o(r)/r = 0$ such that*

$$E_{r,t} \subset (B(\mathbf{z}_0, r + tC) \cap H) + [-at - o(r)t, 0]\mathbf{n}_H$$

The addition on right-hand side of this statement should be interpreted as addition of sets on a product space. That is the set on the right-hand side is the cylinder of points that can be written as $\mathbf{z} + s\mathbf{n}_H$ where $\mathbf{z} \in B(\mathbf{z}_0, r + tC) \cap H$ and $s \in [-at - o(r)t, 0]$.

Proof. We have that

$$E_{r,t} = \{Z(s, \mathbf{z}) : -t \leq s \leq 0, \mathbf{z} \in H \cap B(\mathbf{z}_0, r)\}.$$

Let \mathbf{z} be in $B(\mathbf{z}_0, r) \cap H$ and $s \in [-t, 0]$

$$\begin{aligned} Z(s, \mathbf{z}) &= \mathbf{z} + sX(\mathbf{z}_0) + e(s, \mathbf{z}) \\ &= \mathbf{z} + s(X(\mathbf{z}_0) - a\mathbf{n}_H) + sa\mathbf{n}_H + e(s, \mathbf{z}) \end{aligned}$$

Notice that $X(\mathbf{z}_0) - a\mathbf{n}_H$ is in H hence $\mathbf{z} + s(X(\mathbf{z}_0) - a\mathbf{n}_H)$ is also in H . Lemma 1 allows us to bound $\frac{1}{s}e(s, \mathbf{z})$ on $[0, t] \times B(\mathbf{z}_0, r) \cap H$ which concludes the proof. \square

Lemma 3. *For $\mathbf{z} \in B(\mathbf{z}_0, r) \cap H$ and $s \in [-t, 0]$, there exists a constant $C > 0$ and function $o(r) \geq 0$ with $\lim_{r \rightarrow 0} o(r)/r = 0$ such that*

$$(B(\mathbf{z}_0, r - tC) \cap H) + [-at + o(r)t, 0]\mathbf{n}_H \subset E_{r,t}.$$

Proof. Let $\mathbf{z} \in B(\mathbf{z}_0, r - tC) \cap H - [0, at - o(r)t]n_H$.

$$\langle Z(s, \mathbf{z}), \mathbf{n}_H \rangle = \langle \mathbf{z}, \mathbf{n}_H \rangle + sa + \langle e(s, \mathbf{z}), \mathbf{n}_H \rangle$$

Using Lemma 1, there exists $s \in [0, at - o(r)]$ (intuitively $s \approx \langle \mathbf{z}_0 - \mathbf{z}, \mathbf{n}_H \rangle / a \leq \frac{at}{a} = t$) such that $\langle Z(s, \mathbf{z}), \mathbf{n}_H \rangle = \langle \mathbf{z}_0, \mathbf{n}_H \rangle$, which implies $Z(s, \mathbf{z}) \in H$. Hence $\mathbf{z} \in E_{r,t}$. \square

Using Lemmas 2 and 3, we have constants C_1 and C_2 and a function $o_1(r)$ and $o_2(r)$ with $o_k(r)/r \rightarrow 0$ ($k = 1, 2$) such that

$$\begin{aligned} \frac{\text{Volume}(E_{r,t})}{t} &\geq (a - o_1(r)) \text{Volume}_H(B(\mathbf{z}_0, r - tC_1) \cap H) \\ \frac{\text{Volume}(E_{r,t})}{t} &\leq (a + o_2(r)) \text{Volume}_H(B(\mathbf{z}_0, r + tC_2) \cap H) \end{aligned}$$

where Volume and Volume_H are with respect to the Lebesgue measure of E and H respectively. Since μ has continuous densities, for any $\epsilon > 0$ there exists r_0 such that for every $r < r_0$,

$$\left| \frac{\mu(E_{r,t})}{t} - \frac{\mu(\mathbf{z}_0) \text{Volume}(E_{r,t})}{t} \right| < \epsilon.$$

Hence for every $\epsilon > 0$ there exists $r_0 > 0$ and $T > 0$ such that for every $r < r_0$ and $0 < t < T$,

$$\left| \frac{\frac{1}{t} \mu(E_{r,t})}{\text{Volume}_H(H \cap B(\mathbf{z}_0, r))} - |a| \mu(\mathbf{z}_0) \right| < \epsilon$$

which concludes the proof. \square

B Sensitivity to tuning parameters

The tuning parameters for the reversible jump PDMP methods consist of:

1. the reversible jump parameter $p_{i,j}$ introduced in Section 3;
2. the refreshment distribution for velocities of BPS; and
3. the refreshment rate λ^{ref} for BPS.

To investigate the performance of our methods we consider sampling from a 100-dimensional spike and slab distribution with $\theta_j \sim 0.5\mathcal{N}(0.5, 1) + 0.5\delta_0$, independently for $j = 1, \dots, 100$. Each sampler was run 100 times for a computational budget of 10.




Figure 6: Plots of the Monte Carlo variance for different samplers with varying choices of tuning parameters. The Monte Carlo variance is defined relative to the estimates of the posterior means and probabilities of inclusion.

We compare the methods based on the statistical efficiency of the estimates of marginal means and probability of inclusions. The PDMP methods appear to perform optimally in estimation of the marginal mean when they are able to spend a reasonable amount of time exploring the model space and exploring the parameter space. This can be seen in Figure 6 where optimal statistical efficiency for the marginal means occurs when $p_{i,j}$ is the range 0.2 to 0.7 across all samplers. Figure 6 also shows that higher values of the reversible jump parameter $p_{i,j}$ gives better estimation of the probabilities of inclusions. Having a high refreshment rate appears to impact BPS with Normal velocity distribution less than BPS with velocities distributed uniformly on the hyper-sphere. This is seen as the magenta curve (refreshment = 10) is lower for BPS N than BPS S in Figure 6 across all reversible jump parameters. All BPS methods appear to out-perform Zig-Zag in terms of marginal mean estimation when the refreshment rate is low.

C General implementation details

Inference in Bayesian model selection relies on expectations with respect to a posterior target distribution, $\pi(\boldsymbol{\theta}, \boldsymbol{\gamma})$. The parameters are $\boldsymbol{\theta}$ while $\boldsymbol{\gamma}$ is a vector which indexes the model with elements $\gamma_j = 1$ if the j th variable is included and $\gamma_j = 0$ otherwise. The posterior has the form

$$\pi(\boldsymbol{\theta}, \boldsymbol{\gamma}) \propto L(y^{1:n} | \boldsymbol{\theta}, \boldsymbol{\gamma}) \pi_0(\boldsymbol{\theta} | \boldsymbol{\gamma}) \pi_0(\boldsymbol{\gamma})$$

where $L(y^{1:n} | \boldsymbol{\theta}, \boldsymbol{\gamma})$ defines a likelihood function for observations $y^{1:n}$, $\pi_0(\boldsymbol{\theta} | \boldsymbol{\gamma})$ and $\pi_0(\boldsymbol{\gamma})$ denote prior distribution for $\boldsymbol{\theta}$ and $\boldsymbol{\gamma}$. We abuse notation writing $\boldsymbol{\theta}_\gamma$ to denote the sub-vector of $\boldsymbol{\theta}$ with only the elements where $\gamma_j = 1$. Moreover, we write $\pi(\boldsymbol{\theta}_\gamma)$ for $\pi(\boldsymbol{\theta} | \boldsymbol{\gamma})$ where $\pi(\boldsymbol{\theta} | \boldsymbol{\gamma}) = 0$ whenever $|\theta_j| > 0$ with corresponding $\gamma_j = 0$.

When simulating from the reversible jump PDMP sampler there are two types of events: normal events for the PDMP sampler within a model $\boldsymbol{\gamma}$ and model jump events. The standard PDMP events are taken with respect to $\pi(\boldsymbol{\theta} | \boldsymbol{\gamma})$ so rates to sample are given using the usual Bouncy Particle Sampler or Zig-Zag rates on the conditioned model

$$\begin{aligned} \lambda^{BPS}(s) &= (-\mathbf{v}_\gamma \cdot \nabla_{\boldsymbol{\theta}_\gamma} \log \pi(\boldsymbol{\theta}_\gamma + s\mathbf{v}_\gamma))^+, \\ \lambda_i^{ZZ}(s) &= (-v_i \nabla_{\theta_i} \log \pi(\boldsymbol{\theta}_\gamma + s\mathbf{v}_\gamma))^+, \quad \text{for } i \in \{i : \gamma_i = 1\}. \end{aligned}$$

In practice to simulate these events we first bound the rates by a simple function, which for the examples we consider will be linear in time. We then simulate events from a Poisson process with this linear-in-time rate, which can be done exactly, and use thinning to generate the actual events in the PDMP. Derivations of the linear-in-time bounds on the rates that we use are now given, before we give the rates for jumps between models.

C.1 Rates for logistic and robust regression

Here we give details on simulating rates for the logistic regression example. Taking a simple independent Gaussian prior the reintroduction rate simplifies to

$$\beta_{\boldsymbol{\gamma} \rightarrow \boldsymbol{\gamma}'} = p_{\boldsymbol{\gamma} \rightarrow \boldsymbol{\gamma}'} \frac{1}{\sqrt{2\pi\sigma^2}} \frac{w}{(1-w)}$$

where $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ are defined as above. The standard PDMP rates are used for $\pi(\boldsymbol{\theta} | \boldsymbol{\gamma}) = \pi(\boldsymbol{\theta}_\gamma)$. So to simplify notation we will assume a fixed dimension and write $\boldsymbol{\theta}$ dropping the indexing with $\boldsymbol{\gamma}$. The log posterior for both logistic and robust regression can be written in the form

$$-\log \pi(\boldsymbol{\theta}) = \sum_{i=1}^n g(e_i) + \frac{\boldsymbol{\theta}^T \boldsymbol{\theta}}{2\sigma^2}.$$

For logistic regression $e_i = -\mathbf{x}_i^T \boldsymbol{\theta}$ and $g(e_i) = -\log\left(\frac{\exp(y_i)}{1 + \exp(e_i)}\right)$, and for robust regression $e_i = y_i - \mathbf{x}_i^T \boldsymbol{\theta}$ and $g(e_i) = -\log\left(\exp(-\frac{1}{2}e_i^2) + \frac{1}{10}\exp(-\frac{1}{200}e_i^2)\right)$. We consider bounding the

event rates for the Bouncy Particle Sampler and ZigZag below.

Bouncy Particle Sampler: Let $f(\boldsymbol{\theta} + t\mathbf{v}) = -\nabla_{\boldsymbol{\theta}} \log \pi(\boldsymbol{\theta} + t\mathbf{v})$ the event rate depends on the quantity

$$\langle \mathbf{v}, f(\boldsymbol{\theta} + t\mathbf{v}) \rangle = \langle \mathbf{v}, \sum_{i=1}^n \mathbf{x}_i^T g'(e_i(t)) + \frac{1}{\sigma^2}(\boldsymbol{\theta} + t\mathbf{v}) \rangle.$$

where $g'(e_i(t))$ is the derivative of g evaluated at $e_i(t) = -\mathbf{x}_i^T(\boldsymbol{\theta} + t\mathbf{v})$ for logistic regression and $e_i(t) = y_i - \mathbf{x}_i^T(\boldsymbol{\theta} + t\mathbf{v})$ for robust regression. The in-time derivative of this quantity is

$$\frac{d}{dt} \langle \mathbf{v}, f(\boldsymbol{\theta} + t\mathbf{v}) \rangle = \langle \mathbf{v}, \sum_{i=1}^n \mathbf{x}_i^T g''(e_i(t)) \mathbf{x}_i^T \mathbf{v} + \frac{1}{\sigma^2} \mathbf{v} \rangle = \sum_{i=1}^n g''(e_i(t)) (\mathbf{x}_i^T \mathbf{v})^2 + \frac{1}{\sigma^2} \mathbf{v}^T \mathbf{v}.$$

If the in-time derivative can be bounded by a constant we can simulate using linear rates. For logistic regression $g''(e_i) \leq \frac{1}{4}$ and for robust regression $g''(e_i) < 1$. The Bouncy Particle Sampler rate is bounded by the linear rate

$$\max(0, \langle \mathbf{v}, f(\boldsymbol{\theta} + t\mathbf{v}) \rangle) \leq \max(0, \langle \mathbf{v}, f(\boldsymbol{\theta}) \rangle) + t \left(\frac{1}{\sigma^2} \mathbf{v}^T \mathbf{v} + c \sum_{i=1}^n (\mathbf{x}_i^T \mathbf{v})^2 \right)$$

where c is chosen according to the application. Inversion methods for thinning a Poisson process can be used to simulate the events (Bierkens *et al.*, 2018).

ZigZag: Let $f(\boldsymbol{\theta} + t\mathbf{v}) = -\frac{d}{d\theta_i} \log \pi(\boldsymbol{\theta} + t\mathbf{v})$ the event rate depends on the quantity

$$v_i f(\boldsymbol{\theta} + t\mathbf{v}) = v_i \sum_{i=1}^n x_{ij} g'(e_i(t)) + v_i \frac{1}{\sigma^2} (\theta_i + t v_i)$$

where $g'(e_i(t))$ is defined as in the BPS rate. The in-time derivative of this quantity is

$$\frac{d}{dt} v_i f(\boldsymbol{\theta} + t\mathbf{v}) = v_i \sum_{i=1}^n x_{ij} g''(e_i(t)) \mathbf{x}_i^T \mathbf{v} + v_i^2 \frac{1}{\sigma^2} = \sum_{i=1}^n g''(e_i(t)) x_{ij} v_i \mathbf{x}_i^T \mathbf{v} + v_i^2 \frac{1}{\sigma^2}.$$

Using the same method as before the ZigZag rate is bounded by the linear rate

$$\max(0, v_i f(\boldsymbol{\theta} + t\mathbf{v})) \leq \max(0, v_i f(\boldsymbol{\theta})) + t \left(v_i^2 \frac{1}{\sigma^2} + \sum_{i=1}^n |c x_{ij} v_i \mathbf{x}_i^T \mathbf{v}| \right)$$

where c is chosen according to the application.

C.2 Rates of jumps between models

Model jump events occur when a parameter, θ_i , hits a hyper-plane $\theta_i = 0$ and with probability $p_{\gamma \rightarrow \gamma'}$ we jump to model γ' where $\gamma'_i = 0$. The other type of model jump event occurs when a variable is reintroduced. For each of the deactivated variables ($\gamma_i = 0$), we simulate a time to reintroduce the variable (switching to γ' where $\gamma'_{-i} = \gamma_{-i}$ with $\gamma'_i = 1$). The rate to reintroduce the variable in the Bayesian inference problem is

$$\beta_{\gamma' \rightarrow \gamma} = p_{\gamma' \rightarrow \gamma} \frac{L(y^{1:n} | \boldsymbol{\theta}_{\gamma'}) \pi_0(\boldsymbol{\theta}_{\gamma'}) \pi_0(\boldsymbol{\gamma}')}{L(y^{1:n} | \boldsymbol{\theta}_{\gamma}) \pi_0(\boldsymbol{\theta}_{\gamma}) \pi_0(\boldsymbol{\gamma})},$$

where often computational savings are possible since the reintroduced variable will be zero $\theta_i = 0$ and it is often the case that $L(y^{1:n} | \boldsymbol{\theta}_{\gamma'}) = L(y^{1:n} | \boldsymbol{\theta}_{\gamma})$. In these cases the rate to reintroduce a variable will only depend on the choice of prior.

Both examples we consider had a Gaussian spike and slab prior, of the form

$$\begin{aligned} \boldsymbol{\theta}_{\gamma} &\sim \mathcal{N}(\boldsymbol{\mu}_{\gamma}, \boldsymbol{\Sigma}_{\gamma}) \\ \boldsymbol{\gamma} &\sim w \sum_{j=1}^p \gamma_j (1-w)^{p-\sum_{j=1}^p \gamma_j}, \end{aligned}$$

for a fixed w . The rate to reintroduce the i th variable, jumping from model γ to γ' where $\gamma'_{-i} = \gamma_{-i}$ with $\gamma'_i = 1$ and $\gamma_i = 0$ is given by

$$\beta_{\gamma \rightarrow \gamma'} = p_{\gamma \rightarrow \gamma'} \frac{\pi(\boldsymbol{\theta}_{\gamma})}{\pi(\boldsymbol{\theta}_{\gamma'})} \frac{\pi(\boldsymbol{\gamma})}{\pi(\boldsymbol{\gamma}')} = \frac{\pi(\boldsymbol{\theta}_{\gamma})}{\pi(\boldsymbol{\theta}_{\gamma'})} \frac{w}{(1-w)}$$

Denoting $V_{\gamma} = \boldsymbol{\Sigma}_{\gamma}^{-1}$, the ratio simplifies as,

$$\frac{\pi(\boldsymbol{\theta}_{\gamma})}{\pi(\boldsymbol{\theta}_{\gamma'})} = \frac{|V_{\gamma}|^{1/2} \exp\left(-\frac{1}{2}(\boldsymbol{\theta}_{\gamma} - \boldsymbol{\mu}_{\gamma})^T V_{\gamma} (\boldsymbol{\theta}_{\gamma} - \boldsymbol{\mu}_{\gamma})\right)}{\sqrt{2\pi} |V_{\gamma'}|^{1/2} \exp\left(-\frac{1}{2}(\boldsymbol{\theta}_{\gamma'} - \boldsymbol{\mu}_{\gamma'})^T V_{\gamma'} (\boldsymbol{\theta}_{\gamma'} - \boldsymbol{\mu}_{\gamma'})\right)}.$$

In our examples the prior is independent across components and this ratio simplifies to a constant. As the prior mean is 0 and, if we denote the prior variance for θ_i for any active covariate i as σ^2 , we have

$$\beta_{\gamma \rightarrow \gamma'} = p_{\gamma \rightarrow \gamma'} \frac{1}{\sqrt{2\pi\sigma^2}} \frac{w}{(1-w)}.$$

C.3 Pólya-Gamma Gibbs sampling for logistic regression

The Polya-Gamma Gibbs sampling approach is an auxiliary variable approach for Bayesian Logistic regression. A Polya-Gamma random variable $\omega \sim \mathcal{PG}(b, 0)$, $b > 0$, with probability density $p(\omega)$ has the property (Polson *et al.*, 2013) that for any $\psi \in \mathbb{R}$ and $a \in \mathbb{R}$

$$\frac{\exp(\psi)^a}{(1 + \exp(\psi))^b} = 2^{-b} \exp\left(\left(a - \frac{b}{2}\right)\psi\right) \int_0^{\infty} \exp\left(-\omega \frac{\psi^2}{2}\right) p(\omega) d\omega.$$

Thus the implied conditional distribution for ψ , given auxiliary variable ω , is Gaussian. The advantage of this approach is that when updating the model γ we can integrate over the parameters θ yielding much more efficient moves. The updates for the collapsed Gibbs sampling procedure follow the form:

- (1) sample $\gamma \sim \gamma \mid \omega$;
- (2) sample $\theta \sim \theta \mid \omega, \gamma$;
- (3) sample $\omega \sim \omega \mid \theta, \gamma$.

Simulation step 1.

Let $\tilde{\pi}(\gamma \mid \omega)$ be a density proportional to $\pi(\gamma \mid \omega)$ such that

$$\begin{aligned} \tilde{\pi}(\gamma \mid \omega) &= \pi_0(\gamma) \int_{\theta_\gamma} L(y^{1:n} \mid \theta_\gamma, \omega) \pi_0(\theta_\gamma) d\theta_\gamma \\ &= \frac{\pi_0(\gamma)}{\sqrt{\det(2\pi\sigma^2\mathbf{I}_\gamma)}} \int_{\theta_\gamma} \prod_{i=1}^n \exp\left((y_i - 0.5)(\mathbf{X}_\gamma \theta_\gamma)_i - \frac{\omega_i}{2}((\mathbf{X}_\gamma \theta_\gamma)_i)^2 - \frac{1}{2\sigma^2} \theta_\gamma^T \theta_\gamma\right) d\theta_\gamma \\ &= \pi_0(\gamma) \sqrt{\frac{\det(2\pi\mathbf{V}_\gamma)}{\det(2\pi\sigma^2\mathbf{I}_\gamma)}} \exp\left(\frac{1}{2} \kappa^T \mathbf{X}_\gamma \mathbf{V}_\gamma \mathbf{X}_\gamma^T \kappa\right) \end{aligned}$$

where $(\mathbf{X}_\gamma \theta_\gamma)_i$ denotes the i th element of $\mathbf{X}_\gamma \theta_\gamma$, the matrix $\mathbf{V}_\gamma = (\mathbf{X}_\gamma^T \Omega \mathbf{X}_\gamma + \frac{1}{\sigma^2} \mathbf{I}_\gamma)^{-1}$, the column vector $\kappa = y^{1:n} - 0.5$ and $\Omega = \text{diag}(\omega_1, \dots, \omega_n)$.

The update for γ is taken by updating component-wise from the conditionals $\gamma_j \mid \gamma_{(-j)}, \omega$ where $\gamma_{(-j)} = (\gamma_1, \dots, \gamma_{j-1}, \gamma_{j+1}, \dots, \gamma_p)$. Such a proposal can be implemented using the relationship (Chipman *et al.*, 2001)

$$Pr(\gamma_j = 1 \mid \gamma_{(-j)}, \omega) = \frac{\tilde{\pi}(\gamma_j = 1 \mid \gamma_{(-j)}, \omega)}{\tilde{\pi}(\gamma_j = 0 \mid \gamma_{(-j)}, \omega)} \left(1 + \frac{\tilde{\pi}(\gamma_j = 1 \mid \gamma_{(-j)}, \omega)}{\tilde{\pi}(\gamma_j = 0 \mid \gamma_{(-j)}, \omega)}\right)^{-1}$$

Simulation step 2.

The conditional for θ_γ is

$$\theta_\gamma \mid \omega, \gamma \sim \mathcal{N}(\mathbf{m}_\gamma, \Sigma_\gamma)$$

where $\Sigma_\gamma = (\mathbf{X}_\gamma^T \Omega \mathbf{X}_\gamma + \frac{1}{\sigma^2} \mathbf{I}_\gamma)$ and $\mathbf{m}_\gamma = \Sigma_\gamma \mathbf{X}_\gamma \kappa$.

Simulation step 3.

The conditional for ω is $\omega_i \mid \theta, \gamma \sim \mathcal{PG}(0, (\mathbf{X}_\gamma \theta_\gamma)_i)$ for $i = 1, \dots, n$.

D Computation of relative efficiencies in Section 5.1

In order to compute the relative efficiency we need an estimate of the statistical efficiency (11). We estimate this quantity using a reference estimate q from an independent 6-hour

run of the Gibbs sampling method for each combination of n , p and Scenario in Tables 1-3 and the results of Figure 2. For the results in Tables 1-3 the quantities of interest are the estimation of the posterior marginal inclusion probabilities $\pi(|\theta_j| > 0)$ (PPI) and the marginal posterior means $\mathbb{E}[\theta_j]$ (Mean). These two quantities allowed us to see how efficient the sampler was in terms of exploring both the parameter and model space. For the simulations in the subsampling comparison (Figure 2) the quantity of interest was the posterior mean conditioned on being in model $\gamma = (1, 1, 0, 0, \dots, 0)$. We estimate the mean square error of these terms by running 100 independent runs of each algorithm and comparing to the corresponding long Gibbs run. Methods used in Tables 1-3 were initialised at the maximum a posteriori estimate for a model with active variables chosen from an initial LASSO fit. For the subsampling comparison, methods were initialised at the location of the control variate (the maximum a posteriori estimate with using the true nonzero variables). For each algorithm in the simulations of Tables 1-3 we use a computational budget of 10^6 iterations with a maximum run time of 2 minutes. For the simulations in the subsampling comparison we used a computational budget of 10^6 iterations with a maximum run time of 6 minutes. Algorithms were then compared on the basis of relative computational efficiency using RE or relative efficiency per iteration using RSE. An iteration for the Gibbs sampler is considered to be a full update of all parameters (i.e. one run of all steps in Section C.3) whereas an iteration of the PDMP methods is considered to be one simulated event time.