
Appendix: Doubly Robust Bayesian Inference for Non-Stationary Streaming Data with β -Divergences

Contents

1	Student-t Experiments	2
2	Proof of Theorem 1	4
3	Variational Bayes Approximation for β-divergence based General Bayesian Inference with the Bayesian Linear Model	6
3.1	Q_1	7
3.2	Q_2	9
3.3	ELBO	12
3.4	Differentiation	12
3.4.1	Derivative with respect to L	13
3.4.2	Derivative with respect to $\hat{\mu}_n$	15
3.4.3	Derivative with respect to \hat{a}_n	16
3.4.4	Derivative with respect to \hat{b}_n	17
4	Timing and performance comparisons: Markov Chain Monte Carlo vs Structured Variational Bayes	18
5	Initialization for β_p	19
6	Recursive Optimization for β_{rlm}	19
7	Proof of Theorem 2	21
8	Complexity Analysis of Inference	22
9	Additional Details on Experiments	23
9.1	Well-log data	23
9.2	Air Pollution data	23
9.3	Optimizing β	24

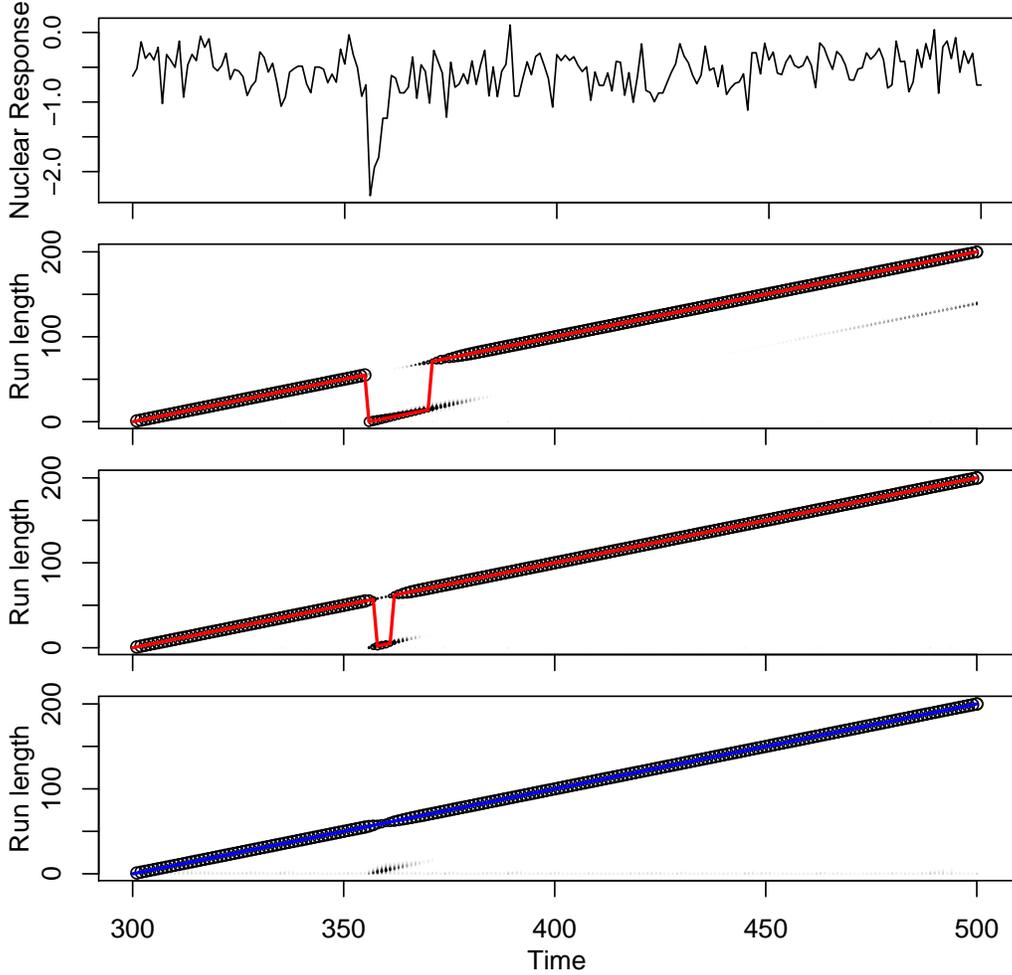


Figure 1: **Top:** Plot of the *well-log* data between $t = 300$ and $t = 500$ with one obvious outlying period. **Middle Top:** KLD run length posterior under the Gaussian error model with the MAP of the run length posterior at each time point overlaid in red. **Middle Bottom:** KLD run length posterior under the Student's t with 5 degrees of freedom error model with the MAP of the run length posterior at each time point overlaid in red. **Bottom:** β -D run length posterior under the Gaussian error model with $\beta_p = 0.25$ and $\beta_{rim} = 0.5$ with the MAP of the run length posterior at each time point overlaid in blue.

1 Student-t Experiments

In Section 2.3 (Quantifying Robustness) of the paper we argue that substituting the Gaussian error model in the BLR setting for a Student's t error model – a traditional solution for robust parameter inference – will be insufficient to ensure that standard Bayesian run-length posteriors are robust. Here, the type of robustness we refer to is defined in Theorem 1. To demonstrate this, we implement a version of BOCPD using both the Gaussian error model and the Student's t error model on two subsets of the *well-log* data. The Student's t distribution is no longer an exponential family and thus cannot be implemented in analytical form or via our structural variational approximation. Hence, we used *stan* [Carpenter et al., 2016] for MCMC sampling from the parameter posterior under the Student's t error model. For comparability, hyperparameters were fixed for both the Gaussian and Student's t error models at $\mu_0 = 0$, $\Sigma_0 = \sqrt{5}$, $a_0 = 0.5$, $b_0 = 2$, $h(r_{t+1}) = 0.01 \forall r_{t+1}$, where $N = 1000$ values were sampled from the parameter posterior, $M = 25$ run lengths were stored and the degrees of freedom of the Student's t error model were set to be $\nu = 5$. Figures 1 and 2 plot the KLD run-length posteriors of the Gaussian and Student's- t error models as well as the β -D run-length

posteriors of the Gaussian error models for the two subsets of the *well-log* data. In both examples, the KLD run-length posteriors favor declaring a CP under both the Gaussian and Student's t error model at the first sign of an outlier. In the second example, the outlier is severe enough to permanently disrupt the run-length inference for both KLD-based methods, while the β -D-based method remains robust. Theorem 1 outlines situations where this desirable behaviour of β -D-based inference can be guaranteed to happen when it would not happen under the KLD with *any* error model.

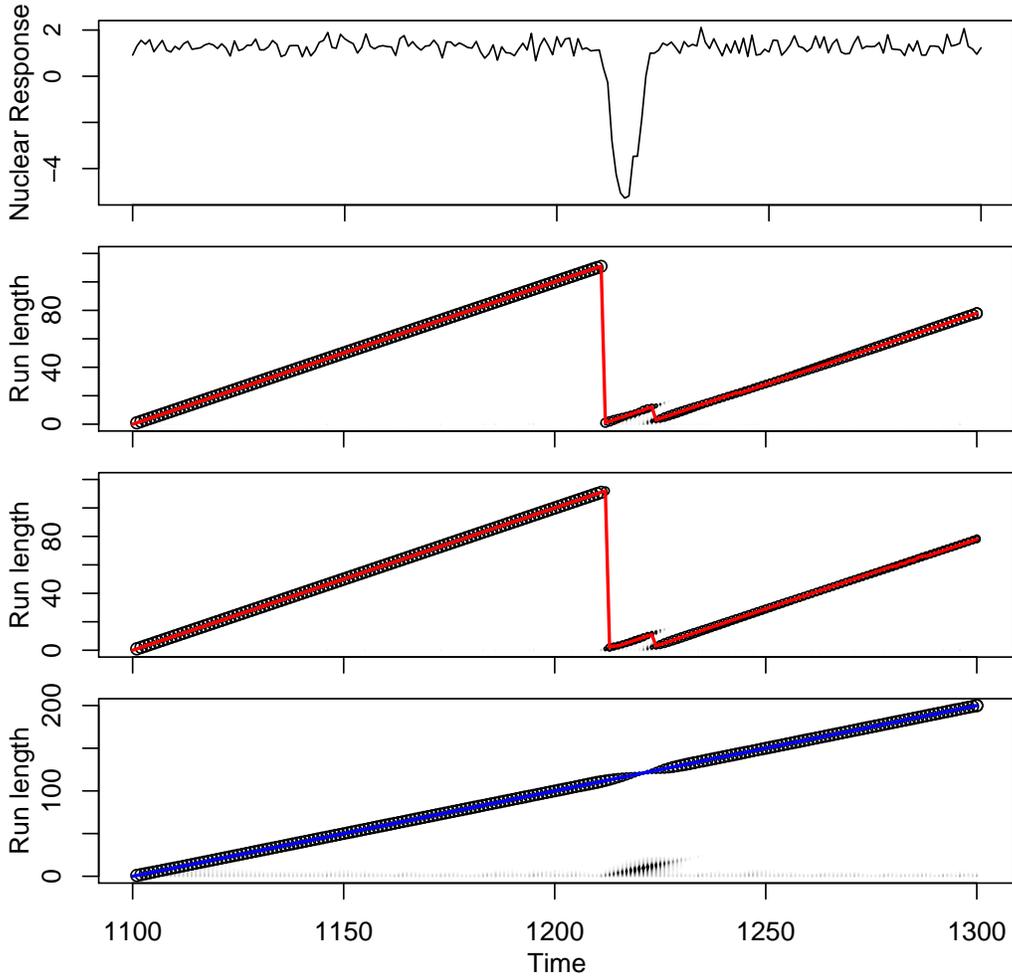


Figure 2: **Top:** Plot of the *well-log* data between $t = 1100$ and $t = 1300$ with one obvious outlying period. **Middle:** KLD run length posterior under the Gaussian error model with the MAP of the run length posterior at each time point overlaid in red. **Second:** KLD run length posterior under the Student's t with 5 degrees of freedom error model with the MAP of the run length posterior at each time point overlaid in red. **Bottom:** β -D run length posterior under the Gaussian error model with $\beta_p = 0.5$ and $\beta_{rim} = 1$ with the MAP of the run length posterior at each time point overlaid in blue.

2 Proof of Theorem 1

Proof. This proof looks at the run length posterior parameterised by β_{rlm} , however to ease notation we refer to $\beta_{\text{rlm}} = \beta$ throughout. Condition on the event $r_t = r$ then after one time step either $r_{t+1} = r + 1$ or $r_{t+1} = 0$. The odds of these two possibilities are as in Thm. 1. Now substituting the definitions of $f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)$ and $f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)$ leaves

$$\frac{f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)}{f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)} = \frac{\exp\left(\frac{1}{\beta}p(\mathbf{y}_{t+1}|\mathbf{y}_{1:t})^\beta - \frac{1}{1+\beta}\int p(\mathbf{z}|\mathbf{y}_{1:t})^{1+\beta}d\mathbf{z}\right)}{\exp\left(\frac{1}{\beta}p(\mathbf{y}_{t+1}|\mathbf{y}_0)^\beta - \frac{1}{1+\beta}\int p(\mathbf{z}|\mathbf{y}_0)^{1+\beta}d\mathbf{z}\right)} \quad (1)$$

$$= \exp\left(\frac{1}{\beta}\left(p(\mathbf{y}_{t+1}|\mathbf{y}_{1:t})^\beta - p(\mathbf{y}_{t+1}|\mathbf{y}_0)^\beta\right) - \frac{1}{1+\beta}\int p(\mathbf{z}|\mathbf{y}_{1:t})^{1+\beta} - p(\mathbf{z}|\mathbf{y}_0)^{1+\beta}d\mathbf{z}\right). \quad (2)$$

This proof first seeks a lower bound for this ratio. A lower bound on $\frac{1}{\beta}p(\mathbf{y}_{t+1}|\mathbf{y}_{1:t})^\beta$ is 0, while the maximal value of $\frac{1}{\beta}p(\mathbf{y}_{t+1}|\mathbf{y}_0)^\beta$ will occur at the prior mode. For the multivariate t -distribution prior predictive with NIG hyperparameters $a_0, b_0, \mu_0, \Sigma_0$ of dimensions p the prior mode has density

$$p(\mu_0|\nu_0, \mu_0, \mathbf{V}_0, p) = \frac{\Gamma((\nu_0 + p)/2)}{\Gamma(\nu_0/2)\nu_0^{p/2}\pi^{p/2}|\mathbf{V}_0|^{1/2}} \left[1 + \frac{1}{\nu_0}(\mu_0 - \mu_0)\Sigma_0^{-1}(\mu_0 - \mu_0)\right]^{-(\nu_0+p)/2} \quad (3)$$

$$= \frac{\Gamma((\nu_0 + p)/2)}{\Gamma(\nu_0/2)\nu_0^{p/2}\pi^{p/2}|\mathbf{V}_0|^{1/2}} \quad (4)$$

$$= \frac{\Gamma(a_0 + p/2)}{\Gamma(a_0)(2b_0\pi)^{p/2}|\mathbf{I} + \mathbf{X}\Sigma_0\mathbf{X}^T|^{1/2}}. \quad (5)$$

As a result the only term in the lower bound of $f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)/f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)$ that does not solely depend on the prior parameters is $\frac{1}{1+\beta}\int p(\mathbf{z}|\mathbf{y}_{1:t})^{1+\beta}d\mathbf{z}$. This term appears in the negative and thus to lower bound $f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)/f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)$, an upper bound for $\frac{1}{1+\beta}\int p(\mathbf{z}|\mathbf{y}_{1:t})^{1+\beta}d\mathbf{z}$ must be found. The multivariate t -distribution can be integrated as

$$\frac{1}{1+\beta}\int \text{MVSt}_\nu(\mathbf{z}|\boldsymbol{\mu}, \mathbf{V})^{1+\beta}d\mathbf{z} = \frac{\Gamma((\nu + p)/2)^{\beta+1}\Gamma((\beta\nu + \beta p + \nu)/2)}{\Gamma(\nu/2)^{\beta+1}\Gamma((\beta\nu + \beta p + \nu + p)/2)} \frac{1}{(1+\beta)(\nu\pi)^{(\beta p)/2}|\mathbf{V}|^{\beta/2}} \quad (6)$$

$$= \frac{\Gamma((\nu + p)/2)^\beta\Gamma((\nu + p)/2)\Gamma((\beta\nu + \beta p + \nu)/2)}{\Gamma(\nu/2)^\beta\Gamma(\nu/2)\Gamma((\beta\nu + \beta p + \nu + p)/2)} \frac{1}{(1+\beta)(\pi\nu)^{(\beta p)/2}|\mathbf{V}|^{\beta/2}} \quad (7)$$

$$\leq \frac{\Gamma((\nu + p)/2)^\beta}{\Gamma(\nu/2)^\beta} \frac{1}{(1+\beta)(\pi\nu)^{(\beta p)/2}|\mathbf{V}|^{\beta/2}}. \quad (8)$$

The inequality is derived from the fact that $\frac{\Gamma(x+\frac{\nu}{2})}{\Gamma(x)}$ is increasing in x and as $\beta \geq 0$ and $\nu \geq 0$ then $(\beta\nu + \beta p + \nu)/2 \geq \nu/2$ which implies $\frac{\Gamma((\nu+p)/2)\Gamma((\beta\nu+\beta p+\nu)/2)}{\Gamma(\nu/2)\Gamma((\beta\nu+\beta p+\nu+p)/2)} \leq 1$.

Now employing the well-known result using Stirling's formula to bound the gamma function

$$(2\pi)^{1/2}x^{x-1/2}\exp(-x) \leq \Gamma(x) \leq (2\pi)^{1/2}x^{x-1/2}\exp(1/(12x) - x) \quad (9)$$

we can therefore rewrite the ratio of gamma functions leaving

$$\frac{1}{1+\beta} \int \text{MVSt} - t_\nu(\mathbf{z}|\boldsymbol{\mu}, \mathbf{V})^{1+\beta} d\mathbf{z} \leq \frac{\Gamma((\nu+p)/2)^\beta}{\Gamma(\nu/2)^\beta} \frac{1}{(1+\beta)(\pi\nu)^{(\beta p)/2} |\mathbf{V}|^{\beta/2}} \quad (10)$$

$$\leq \frac{(\sqrt{2\pi}((\nu+p)/2)^{(\nu+p-1)/2}) \exp(-(\nu+p)/2 + 1/6(\nu+p))^\beta}{(\sqrt{2\pi}(\nu/2)^{(\nu-1)/2} \exp(-\nu/2))^\beta (1+\beta)(\pi\nu)^{(\beta p)/2} |\mathbf{V}|^{\beta/2}} \quad (11)$$

$$= \left(1 + \frac{p}{\nu}\right)^{\beta(\nu+p-1)/2} \exp(\beta(1/(6(\nu+p)) - p/2)) \frac{1}{(1+\beta)(\pi)^{(\beta p)/2} |\mathbf{V}|^{\beta/2}}. \quad (12)$$

Clearly $\exp(\beta(1/(6(\nu+p)) - p/2))$ is decreasing in ν for all p and to demonstrate when $\left(1 + \frac{p}{\nu}\right)^{\beta(\nu+p-1)/2}$ is decreasing in ν we examine its derivative

$$w = \left(1 + \frac{p}{\nu}\right)^{\beta(\nu+p-1)/2} \quad (13)$$

$$= \exp\left(\left(\beta(\nu+p-1)/2\right) \log\left(\left(1 + \frac{p}{\nu}\right)\right)\right) \quad (14)$$

$$\frac{dw}{d\nu} = \frac{\beta}{2} \left(\log\left(1 + \frac{p}{\nu}\right) - (\nu+p-1) \frac{\frac{p}{\nu^2}}{1 + \frac{p}{\nu}}\right) \left(1 + \frac{p}{\nu}\right)^{\beta(\nu+p-1)/2}. \quad (15)$$

The sign of $\frac{dw}{d\nu}$ is dictated by $\left(\log\left(1 + \frac{p}{\nu}\right) - (\nu+p-1) \frac{\frac{p}{\nu^2}}{1 + \frac{p}{\nu}}\right)$, which can be demonstrated to be positive always if $p = 1$ and negative always if $p > 1$.

Case 1: when $p > 1$, $\frac{1}{1+\beta} \int p(\mathbf{z}|\mathbf{y}_{1:t})^{1+\beta} d\mathbf{z}$ is decreasing in ν and thus we can upper bound it by substituting the smallest value of ν . Here we bound ν above 1 in order to enforce that the mean of the predictive t -distribution exists. Under the KLD posterior it is clear that a_0 rises as more data is seen and while we do not have closed forms associated with the variational approximation to the β -D posterior we expect this to be the case here. As more data is seen the finite sampling uncertainty, represented by ν in the NIG case, should be decreasing. Therefore provided a_0 is set such that $2a_0 > 1$, then this lower bound should never be violated.

Case 2: when $p = 1$, Stirling's formula has failed to provide a decreasing upper bound for $\frac{1}{1+\beta} \int p(\mathbf{z}|\mathbf{y}_{1:t})^{1+\beta} d\mathbf{z}$. However in the univariate case

$$\frac{1}{1+\beta} \int \text{St}_\nu(\mathbf{z}|\boldsymbol{\mu}, \mathbf{V})^{1+\beta} d\mathbf{z} \leq \frac{\Gamma((\nu+1)/2)^\beta}{\Gamma(\nu/2)^\beta} \frac{1}{(1+\beta)(\nu|\mathbf{V}|)^{(\beta)/2} \pi^{(\beta)/2}} \quad (16)$$

$$\leq \frac{1}{(1+\beta) |\mathbf{V}|^{(\beta)/2} \pi^{(\beta)/2}} \quad (17)$$

Where $p = 1$ is substituted into the bound from equation (8) and the inequality comes from that fact that $\frac{\Gamma((x+1)/2)}{\Gamma(x/2)} \leq \sqrt{x}$. This bound conveniently does not depend on the degrees of freedom ν at all.

We can therefore lower bound $f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)/f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)$ as

$$\frac{f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)}{f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)} \geq \begin{cases} \exp\left\{-\frac{1}{\beta} \left(\frac{\Gamma(a_0+1/2)}{\Gamma(a_0)(2b_0\pi)^{1/2} |\mathbf{I} + \mathbf{X}\boldsymbol{\Sigma}_0\mathbf{X}^T|^{1/2}}\right)^\beta - \frac{1}{(1+\beta)|\mathbf{V}|^{(\beta)/2} \pi^{(\beta)/2}} + \frac{\Gamma(a_0+1/2)^{\beta+1} \Gamma(\beta a_0 + \beta/2 + a_0)}{\Gamma(a_0)^{\beta+1} \Gamma(\beta a_0 + \beta/2 + a_0 + 1/2)} \frac{1}{(1+\beta)(2\pi b_0)^{(\beta)/2} |\mathbf{I} + \mathbf{X}\boldsymbol{\Sigma}_0\mathbf{X}^T|^{(\beta)/2}}\right\} & \text{if } p = 1 \\ \exp\left\{-\frac{1}{\beta} \left(\frac{\Gamma(a_0+p/2)}{\Gamma(a_0)(2b_0\pi)^{p/2} |\mathbf{I} + \mathbf{X}\boldsymbol{\Sigma}_0\mathbf{X}^T|^{1/2}}\right)^\beta + \frac{\Gamma(a_0+p/2)^{\beta+1} \Gamma(\beta a_0 + \beta p/2 + a_0)}{\Gamma(a_0)^{\beta+1} \Gamma(\beta a_0 + \beta p/2 + a_0 + p/2)} \frac{1}{(1+\beta)(2\pi b_0)^{(\beta p)/2} |\mathbf{I} + \mathbf{X}\boldsymbol{\Sigma}_0\mathbf{X}^T|^{(\beta p)/2}} - \frac{1}{(1+\beta)(\pi)^{(\beta p)/2} |\mathbf{V}|^{\beta/2}}\right\} & \text{if } p > 1 \end{cases} \quad (18)$$

Now fixing $p, a_0, b_0, \mu_0, \Sigma_0$ and $|V|_{\min}$ which values of β and $H(r_t, r_{t+1})$ would leave

$$\frac{1 - H(r_t, r_{t+1})}{H(r_t, r_{t+1})} \frac{f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)}{f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)} \geq 1? \quad (19)$$

We demonstrate this for $p > 1$ but it is straightforward to see that it extends to when $p = 1$. Rearranging the inequality in equation (18) gives us that (19) holds providing

$$\frac{1}{|\mathbf{V}|^{\beta/2}} \leq \left(\frac{\Gamma(a_0 + p/2)^\beta}{\Gamma(a_0)^\beta (2b_0\pi)^{\beta p/2} |\mathbf{I} + \mathbf{X}\Sigma_0\mathbf{X}^T|^{\beta/2}} \left(\frac{\Gamma(a_0 + p/2)\Gamma(\beta a_0 + \beta p/2 + a_0)}{\Gamma(a_0)\Gamma(\beta a_0 + \beta p/2 + a_0 + p/2)} \frac{1}{(1 + \beta)} - \frac{1}{\beta} \right) \right. \\ \left. + \log \left(\frac{1 - H(r_t, r_{t+1})}{H(r_t, r_{t+1})} \right) \right) \frac{(1 + \beta)(\pi)^{(\beta p)/2}}{\left((1 + \frac{p}{2a_0})^{\alpha(2a_0 + p - 1)/2} \exp(\beta(1/(6(2a_0 + p)) - p/2)) \right)} \quad (20)$$

We define the set defined by inequality (20) as $S(p, \beta, a_0, b_0, \mu_0, \Sigma_0, |\mathbf{V}|_{\min}) = \{(\beta, H(r_t, r_{t+1})) : (\beta, H(r_t, r_{t+1})) \text{ satisfy (20) for } p, \beta, a_0, b_0, \mu_0, \Sigma_0, |\mathbf{V}|_{\min}\}$. As a result we can see that for fixed of $a_0, b_0, \mu_0, \Sigma_0$ and $|\mathbf{V}| \geq |\mathbf{V}|_{\min}$ it is always possible to choose values of β and $H(r_t, r_{t+1})$ such that this holds. To see this consider fixing β , the the upper bound is simply increasing in $\log \left(\frac{1 - H(r_t, r_{t+1})}{H(r_t, r_{t+1})} \right)$ which takes values in \mathbb{R} and thus can be set large enough so that the inequality holds. \square

We note that in practice this result is likely to be stronger than is necessary. The observation that is most likely to generate a change-point will have 0 mass under the predictive associated with the current segment but also appears at the prior mode. While this was necessary to demonstrate this result for all situations this is incredibly unlikely to occur. The requirement for $|\mathbf{V}|_{\min}$ is a result of the beta-divergence loss function depending on $\int p(z|\mathbf{y}_{1:t})^{1+\beta} dz$. In the proof of this result we demonstrate that $f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathcal{F}_t)/f_{m_t}^\beta(\mathbf{y}_{t+1}|\mathbf{y}_0)$ is increasing in $|\mathbf{V}|$ and as a result if it is allowed to get too small the inequality in equation (20) would not hold. This is an undesirable consequence of the beta-divergence score not being completely local, that is to say not solely depending on the predictive probability of the observation, thus the score under the prior can be quite a lot bigger than the score under the continuing run length independent of the observations seen and solely based on the predictive covariances.

3 Variational Bayes Approximation for β -divergence based General Bayesian Inference with the Bayesian Linear Model

For ease of notation, we use $\beta = \beta_p$. We wish to approximate the posterior belief distribution $\pi_{\text{DPD}}^\beta(\boldsymbol{\mu}, \sigma^2|\mathbf{y})$ which for observations $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n)^T$ with $\mathbf{y}_i \in \mathbb{R}^d$, prior $\text{NIG}^0(\boldsymbol{\mu}, \sigma^2|a_0, b_0, \boldsymbol{\mu}_0, \Sigma_0)$, model likelihood f and density power divergence (DPD) loss

$$\ell^\beta(\boldsymbol{\mu}, \sigma^2|\mathbf{y}_i) = \frac{1}{\beta} f(\mathbf{y}_i|\boldsymbol{\mu}, \sigma^2)^\beta - \frac{1}{1 + \beta} \int_{\mathcal{Y}} f(\mathbf{y}_i|\boldsymbol{\mu}, \sigma^2)^{1+\beta} d\mathbf{y} \quad (21)$$

is given by

$$\pi_{\text{DPD}}^\beta(\boldsymbol{\mu}, \sigma^2|\mathbf{y}) = \text{NIG}^0(\boldsymbol{\mu}, \sigma^2|a_0, b_0, \boldsymbol{\mu}_0, \Sigma_0) \cdot \exp \left\{ - \sum_{i=1}^n \ell^\beta(\boldsymbol{\mu}, \sigma^2|\mathbf{y}_i) \right\}. \quad (22)$$

In particular, we want to approximate it with a posterior $\text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2|\hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\Sigma}_n)$ via Variational Bayes. This can be done by minimizing the variational parameters in a Kullback-Leibler sense:

$$(a^*, b^*, \boldsymbol{\mu}^*, \Sigma^*) = \underset{(\hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\Sigma}_n)}{\text{argmin}} \left\{ \text{KL} \left(\pi_{\text{DPD}}^\beta(\boldsymbol{\mu}, \sigma^2|\mathbf{y}) \parallel \text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2|\hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\Sigma}_n) \right) \right\}. \quad (23)$$

It is straightforward to rewrite the objective function for the above minimization as the Evidence Lower Bound (ELBO) induced by the DPD:

$$\text{ELBO}_{\text{DPD}} = \underbrace{-\text{KL} \left(\text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\boldsymbol{\Sigma}}_n) \parallel \text{NIG}^0(\boldsymbol{\mu}, \sigma^2 | a_0, b_0, \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) \right)}_{=Q_1} - \underbrace{\mathbb{E}_{\text{VB}} \left[-\sum_{i=1}^n \ell^\beta(\boldsymbol{\mu}, \sigma^2 | \mathbf{y}_i) \right]}_{=Q_2}. \quad (24)$$

In what follows, closed forms are derived for both Q_1 and Q_2 . Some algebraic tricks will be applied multiple times, and will be referred to by the following symbols:

■ Completion of Squares, i.e. $\mathbf{u}'\mathbf{A}\mathbf{u} - 2\mathbf{v}'\mathbf{u} = (\mathbf{u} - \mathbf{A}^{-1}\mathbf{v})'\mathbf{A}(\mathbf{u} - \mathbf{A}^{-1}\mathbf{v}) - \mathbf{v}'\mathbf{A}^{-1}\mathbf{v}$;

$I(\mathcal{N})$ Integrating out the Normal density;

$I(\mathcal{IG})$ Integrating out the Inverse Gamma density.

Throughout, the dimensionality of $\boldsymbol{\mu}$ is $p \in \mathbb{N}$, $\mathcal{N}(\boldsymbol{\mu} | \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ refers to a normal pdf in $\boldsymbol{\mu}$ with expectation $\boldsymbol{\mu}_0$, variance $\boldsymbol{\Sigma}_0$ and $\mathcal{IG}(\sigma^2 | a, b)$ to an inverse gamma pdf in σ^2 with shape a and scale b .

3.1 Q_1

First, note that by definition,

$$Q_1 = \int_{\boldsymbol{\mu}, \sigma^2} \underbrace{\log \left(\frac{\text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\boldsymbol{\Sigma}}_n)}{\text{NIG}^0(\boldsymbol{\mu}, \sigma^2 | a_0, b_0, \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)} \right)}_{=Q_1^{\log}} \text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\boldsymbol{\Sigma}}_n) d\boldsymbol{\mu} d\sigma^2. \quad (25)$$

Writing out Q_1^{\log} , one obtains a natural sum of three components $C_1, C_2(\sigma^2), C_3(\sigma^2, \boldsymbol{\mu})$:

$$\begin{aligned} Q_1^{\log} &= \log \left(\frac{|\hat{\boldsymbol{\Sigma}}_n|^{-0.5} \frac{\hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n)} (\sigma^2)^{-0.5p-\hat{a}_n-1} \exp \left\{ -\frac{1}{2\sigma^2} [(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n)'\hat{\boldsymbol{\Sigma}}_n^{-1}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n) + 2\hat{b}_n] \right\}}{|\boldsymbol{\Sigma}_0|^{-0.5} \frac{b_0^{a_0}}{\Gamma(a_0)} (\sigma^2)^{-0.5p-a_0-1} \exp \left\{ -\frac{1}{2\sigma^2} [(\boldsymbol{\mu} - \boldsymbol{\mu}_0)'\boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\mu} - \boldsymbol{\mu}_0) + 2b_0] \right\}} \right) \\ &= \underbrace{\log \left(\frac{\hat{b}_n^{\hat{a}_n} \Gamma(a_0)}{b_0^{a_0} \Gamma(\hat{a}_n)} \right)}_{=C_1} + 0.5 \log |\boldsymbol{\Sigma}_0 \hat{\boldsymbol{\Sigma}}_n^{-1}| + \underbrace{(\hat{a}_n - a_0) \log \left(\frac{1}{\sigma^2} \right)}_{=C_2(\sigma^2)} \\ &\quad - \underbrace{\frac{1}{2\sigma^2} [(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n)'\hat{\boldsymbol{\Sigma}}_n^{-1}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n) - (\boldsymbol{\mu} - \boldsymbol{\mu}_0)'\boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\mu} - \boldsymbol{\mu}_0) + 2(\hat{b}_n - b_0)]}_{=C_3(\sigma^2, \boldsymbol{\mu})}. \quad (26) \end{aligned}$$

Next, note that $C_3(\sigma^2, \boldsymbol{\mu})$ further decomposes into

$$\underbrace{\frac{1}{2\sigma^2} \left[\boldsymbol{\mu}'(\hat{\boldsymbol{\Sigma}}_n^{-1} - \boldsymbol{\Sigma}_0^{-1})\boldsymbol{\mu} - 2\boldsymbol{\mu}'(\hat{\boldsymbol{\Sigma}}_n^{-1}\hat{\boldsymbol{\mu}}_n - \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\mu}_0) \right]}_{=C_4(\sigma^2, \boldsymbol{\mu})} + \underbrace{\frac{1}{\sigma^2} \left[\frac{1}{2}\hat{\boldsymbol{\mu}}_n'\hat{\boldsymbol{\Sigma}}_n^{-1}\hat{\boldsymbol{\mu}}_n - \frac{1}{2}\boldsymbol{\mu}_0'\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\mu}_0 + (\hat{b}_n - b_0) \right]}_{=C_5} \underbrace{\quad}_{=C_6(\sigma^2)}.$$

Notice that we have isolated the random variable $\boldsymbol{\mu}$ inside $C_4(\sigma^2, \boldsymbol{\mu})$ and that by definition, $\text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\boldsymbol{\Sigma}}_n) = \mathcal{N}^{\text{VB}}(\boldsymbol{\mu} | \hat{\boldsymbol{\mu}}_n, \sigma^2 \hat{\boldsymbol{\Sigma}}_n) \cdot \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n)$, meaning that

$$\begin{aligned} Q_1 &= C_1 + \int_{\sigma^2} \{C_2(\sigma^2) - C_6(\sigma^2)\} \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2 \\ &\quad - \underbrace{\int_{\sigma^2} \left\{ \int_{\boldsymbol{\mu}} C_4(\sigma^2, \boldsymbol{\mu}) \mathcal{N}^{\text{VB}}(\boldsymbol{\mu} | \hat{\boldsymbol{\mu}}_n, \sigma^2 \hat{\boldsymbol{\Sigma}}_n) d\boldsymbol{\mu} \right\}}_{=C_7(\sigma^2)} \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2. \quad (27) \end{aligned}$$

The inner integral is available in closed form, and naturally decomposes as

$$\begin{aligned}
C_7(\sigma^2) &= \frac{1}{2\sigma^2} \mathbb{E}_{\mathcal{N}^{\text{VB}}} \left[\boldsymbol{\mu}' \left(\hat{\boldsymbol{\Sigma}}_n^{-1} - \boldsymbol{\Sigma}_0^{-1} \right) \boldsymbol{\mu} \right] - \frac{2}{2\sigma^2} \mathbb{E}_{\mathcal{N}^{\text{VB}}} \left[\boldsymbol{\mu}' \right] \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n - \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 \right) \\
&= \frac{1}{2\sigma^2} \mathbb{E}_{\mathcal{N}^{\text{VB}}} \left[\text{tr} \left(\left(\hat{\boldsymbol{\Sigma}}_n^{-1} - \boldsymbol{\Sigma}_0^{-1} \right) \boldsymbol{\mu} \boldsymbol{\mu}' \right) \right] - \frac{1}{\sigma^2} \hat{\boldsymbol{\mu}}_n' \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n - \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 \right) \\
&= \frac{1}{2\sigma^2} \text{tr} \left(\left(\hat{\boldsymbol{\Sigma}}_n^{-1} - \boldsymbol{\Sigma}_0^{-1} \right) \mathbb{E}_{\mathcal{N}^{\text{VB}}} \left[\boldsymbol{\mu} \boldsymbol{\mu}' \right] \right) - \frac{1}{\sigma^2} \hat{\boldsymbol{\mu}}_n' \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n - \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 \right) \\
&= \frac{1}{2\sigma^2} \text{tr} \left(\left(\hat{\boldsymbol{\Sigma}}_n^{-1} - \boldsymbol{\Sigma}_0^{-1} \right) \left[\sigma^2 \hat{\boldsymbol{\Sigma}}_n - \hat{\boldsymbol{\mu}}_n \hat{\boldsymbol{\mu}}_n' \right] \right) - \frac{1}{\sigma^2} \hat{\boldsymbol{\mu}}_n' \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n - \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 \right) \\
&= \underbrace{\frac{1}{2} \text{tr} \left(I - \boldsymbol{\Sigma}_0^{-1} \hat{\boldsymbol{\Sigma}}_n \right)}_{=C_8} - \underbrace{\frac{1}{\sigma^2} \left[\frac{1}{2} \hat{\boldsymbol{\mu}}_n' \left(\hat{\boldsymbol{\Sigma}}_n^{-1} - \boldsymbol{\Sigma}_0^{-1} \right) \hat{\boldsymbol{\mu}}_n - \hat{\boldsymbol{\mu}}_n' \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n - \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 \right) \right]}_{=C_9} \quad (28) \\
&\hspace{15em} = C_{10}(\sigma^2)
\end{aligned}$$

We may now rewrite Q_1 so as to integrate out σ^2 next:

$$Q_1 = C_1 - C_8 + \int_{\sigma^2} \{C_2(\sigma^2) - C_6(\sigma^2) - C_{10}(\sigma^2)\} \mathcal{I}\mathcal{G}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2.$$

Using the additivity of integrals, we consider its three components separately and then add them up together afterwards. For $C_2(\sigma^2)$, (I) apply a change of variable with $z = \frac{\sigma^2}{\hat{b}_n}$ and then use (II) that $\frac{d}{dx} a^{-x} = -a^{-x} \cdot \log(a) = a^{-x} \cdot \log(a^{-1})$ together with Fubini's Theorem (III) to find that

$$\begin{aligned}
C_{11} &= \int_{\sigma^2} C_2(\sigma^2) \mathcal{I}\mathcal{G}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2 \\
&= (\hat{a}_n - a_0) \int_{\sigma^2} \log \left(\frac{1}{\sigma^2} \right) \frac{\hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n)} (\sigma^2)^{-\hat{a}_n-1} \exp \left\{ -\frac{\hat{b}_n}{\sigma^2} \right\} d\sigma^2 \\
&\stackrel{\text{(I)}}{=} (\hat{a}_n - a_0) \int_z \log \left(\frac{1}{z \hat{b}_n} \right) \frac{\hat{b}_n^{\hat{a}_n+1}}{\Gamma(\hat{a}_n)} (z \hat{b}_n)^{-\hat{a}_n-1} \exp \left\{ -\frac{1}{z} \right\} dz \\
&= (\hat{a}_n - a_0) \frac{1}{\Gamma(\hat{a}_n)} \int_z \left(-\log(z) - \log(\hat{b}_n) \right) z^{-\hat{a}_n-1} \exp \left\{ -\frac{1}{z} \right\} dz \\
&\stackrel{\text{(II)}}{=} (\hat{a}_n - a_0) \left[\frac{1}{\Gamma(\hat{a}_n)} \int_z \left(-\log(z) \right) z^{-\hat{a}_n-1} \exp \left\{ -\frac{1}{z} \right\} dz - \log(\hat{b}_n) \right] \\
&\stackrel{\text{(III)}}{=} (\hat{a}_n - a_0) \left[\frac{1}{\Gamma(\hat{a}_n)} \int_z \frac{d}{d\hat{a}_n} \left\{ z^{-\hat{a}_n-1} \exp \left\{ -\frac{1}{z} \right\} \right\} dz - \log(\hat{b}_n) \right] \\
&\stackrel{\text{(III)}}{=} (\hat{a}_n - a_0) \left[\frac{1}{\Gamma(\hat{a}_n)} \frac{d}{d\hat{a}_n} \left\{ \underbrace{\int_z z^{-\hat{a}_n-1} \exp \left\{ -\frac{1}{z} \right\} dz}_{\stackrel{\text{(II)}}{=} \Gamma(\hat{a}_n)} \right\} - \log(\hat{b}_n) \right] \\
&= (\hat{a}_n - a_0) \left(\frac{\Gamma'(\hat{a}_n)}{\Gamma(\hat{a}_n)} - \log(\hat{b}_n) \right) \\
&= (\hat{a}_n - a_0) \left(\Psi(\hat{a}_n) - \log(\hat{b}_n) \right),
\end{aligned}$$

where Ψ is the digamma function. For $C_6(\sigma^2)$, one obtains the closed form as

$$\begin{aligned}
C_{12} &= \int_{\sigma^2} C_6(\sigma^2) \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2 \\
&= C_5 \int_{\sigma^2} \frac{\hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n)} (\sigma^2)^{-\hat{a}_n-1} \exp\left\{-\frac{\hat{b}_n}{\sigma^2}\right\} d\sigma^2 \\
&\stackrel{I(\mathcal{IG})}{=} C_5 \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)}.
\end{aligned} \tag{29}$$

Using the exact same steps for $C_{10}(\sigma^2)$, one finds

$$\begin{aligned}
C_{13} &= \int_{\sigma^2} C_{10}(\sigma^2) \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2 \\
&\stackrel{I(\mathcal{IG})}{=} C_9 \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)},
\end{aligned} \tag{30}$$

finally yielding

$$\begin{aligned}
Q_1 &= C_1 - C_8 + C_{11} - C_{12} - C_{13} \\
&= \log\left(\frac{\hat{b}_n^{\hat{a}_n} \Gamma(a_0)}{b_0^{a_0} \Gamma(\hat{a}_n)}\right) + 0.5 \log\left|\Sigma_0 \hat{\Sigma}_n^{-1}\right| - \frac{1}{2} \text{tr}\left(I - \Sigma_0^{-1} \hat{\Sigma}_n\right) + (\hat{a}_n - a_0) \left(\Psi(\hat{a}_n) - \log(\hat{b}_n)\right) \\
&\quad - \left[\frac{1}{2} \hat{\boldsymbol{\mu}}_n' \hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n - \frac{1}{2} \boldsymbol{\mu}_0' \Sigma_0^{-1} \boldsymbol{\mu}_0 + (\hat{b}_n - b_0)\right] \cdot \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)} \\
&\quad - \left[\frac{1}{2} \hat{\boldsymbol{\mu}}_n' (\hat{\Sigma}_n^{-1} - \Sigma_0^{-1}) \hat{\boldsymbol{\mu}}_n - \hat{\boldsymbol{\mu}}_n' (\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n - \Sigma_0^{-1} \boldsymbol{\mu}_0)\right] \cdot \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)} \\
&= \log\left(\frac{\hat{b}_n^{\hat{a}_n} \Gamma(a_0)}{b_0^{a_0} \Gamma(\hat{a}_n)}\right) + 0.5 \log\left|\Sigma_0 \hat{\Sigma}_n^{-1}\right| - \frac{1}{2} \text{tr}\left(I - \Sigma_0^{-1} \hat{\Sigma}_n\right) + (\hat{a}_n - a_0) \left(\Psi(\hat{a}_n) - \log(\hat{b}_n)\right) \\
&\quad + \frac{1}{2} \left[(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n)' \Sigma_0^{-1} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n) + 2(b_0 - \hat{b}_n)\right] \cdot \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)}
\end{aligned} \tag{31}$$

$$\tag{32}$$

3.2 Q_2

Noting that one can write Q_2 as

$$\begin{aligned}
&= E_{\text{VB}} \left[\sum_{i=1}^n \ell^\beta(\boldsymbol{\mu}, \sigma^2 | \mathbf{y}_i) \right] \\
&= \int_{\boldsymbol{\mu}, \sigma^2} \left\{ \sum_{i=1}^n \left[\frac{1}{\beta} f(\mathbf{y}_i | \boldsymbol{\mu}, \sigma^2)^\beta - \frac{1}{1 + \beta} \int_{\mathcal{Y}} f(\mathbf{y} | \boldsymbol{\mu}, \sigma^2)^{1+\beta} d\mathbf{y} \right] \text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\Sigma}_n) \right\} d\sigma^2 d\boldsymbol{\mu} \\
&= \sum_{i=1}^n \left[\int_{\boldsymbol{\mu}, \sigma^2} \left\{ \frac{1}{\beta} f(\mathbf{y}_i | \boldsymbol{\mu}, \sigma^2)^\beta - \frac{1}{1 + \beta} \int_{\mathcal{Y}} f(\mathbf{y} | \boldsymbol{\mu}, \sigma^2)^{1+\beta} d\mathbf{y} \right\} \text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\Sigma}_n) d\sigma^2 d\boldsymbol{\mu} \right]
\end{aligned}$$

The last equation implies that it is sufficient to concern ourselves with the integral for a single term. To this end, observe that the likelihood for a single observation \mathbf{y}_i with regressor matrix \mathbf{X}_i is given by

$$f(\mathbf{y}_i | \boldsymbol{\mu}, \sigma^2) = \mathcal{N}(\mathbf{y}_i | \mathbf{X}_i' \boldsymbol{\mu}, \sigma^2 I_d), \tag{34}$$

where I_d is the identity matrix of dimension d . Looking at the likelihood terms inside ℓ^β , the β -exponentiated likelihood term can be rewritten as

$$\begin{aligned}
\frac{1}{\beta} f(\mathbf{y}_i | \boldsymbol{\mu}, \sigma^2)^\beta &= \underbrace{\frac{1}{\beta} (2\pi)^{-0.5d\beta} (\sigma^2)^{-0.5d\beta}}_{=D_1(\sigma^2)} \cdot \exp \left\{ -\frac{\beta}{2\sigma^2} [(\mathbf{y}_i - \mathbf{X}_i' \boldsymbol{\mu})' (\mathbf{y}_i - \mathbf{X}_i' \boldsymbol{\mu})] \right\} \\
&= D_1(\sigma^2) \cdot \exp \left\{ -\frac{\beta}{2\sigma^2} \left[\mathbf{y}_i' \mathbf{y}_i + \boldsymbol{\mu}' \underbrace{(\mathbf{X}_i \mathbf{X}_i')}_{=\check{\boldsymbol{\Sigma}}_i^{-1}} \boldsymbol{\mu} - 2(\mathbf{y}_i' \mathbf{X}_i) \boldsymbol{\mu} \right] \right\} \\
&\stackrel{\blacksquare}{=} D_1(\sigma^2) \cdot \exp \left\{ -\frac{1}{2\sigma^2} \left[\beta(\boldsymbol{\mu} - \underbrace{\check{\boldsymbol{\Sigma}}_i(\mathbf{X}_i' \mathbf{y}_i)}_{=\check{\boldsymbol{\mu}}_i})' \check{\boldsymbol{\Sigma}}_i^{-1} (\boldsymbol{\mu} - \check{\boldsymbol{\mu}}_i) + \beta \underbrace{[\mathbf{y}_i' \mathbf{y}_i - (\mathbf{y}_i \mathbf{X}_i') \check{\boldsymbol{\Sigma}}_i (\mathbf{X}_i \mathbf{y}_i)]}_{=D_{2,i}} \right] \right\} \\
&= D_1(\sigma^2) \cdot \exp \left\{ -\frac{1}{2\sigma^2} \left[\underbrace{\beta(\boldsymbol{\mu} - \check{\boldsymbol{\mu}}_i)' \check{\boldsymbol{\Sigma}}_i^{-1} (\boldsymbol{\mu} - \check{\boldsymbol{\mu}}_i)}_{=D_{3,i}(\boldsymbol{\mu})} + D_{2,i} \right] \right\} \\
&= D_1(\sigma^2) \cdot \exp \left\{ -\frac{1}{2\sigma^2} [D_{3,i}(\boldsymbol{\mu}) + D_{2,i}] \right\}, \tag{35}
\end{aligned}$$

while the integral is available in closed form as

$$\frac{1}{1+\beta} \int_{\mathbf{y}} f(\mathbf{y} | \boldsymbol{\mu}, \sigma^2)^{1+\beta} d\mathbf{y} \stackrel{I(\mathcal{N})}{=} (\sigma^2)^{-0.5p\beta} \underbrace{(2\pi)^{-0.5d\beta} (1+\beta)^{-0.5d-1}}_{=D_4} \tag{36}$$

One can see a neat separation between terms involving σ^2 and terms involving $\boldsymbol{\mu}$ again, allowing us to rewrite the integral in equation (33) such as to exploit the conditional structure of the normal inverse-gamma distribution in Eqs. (36), (35). Looking at integrating out σ^2 from (35) first, note that

$$\begin{aligned}
L_1 &= \int_{\sigma^2} \left\{ \frac{1}{1+\beta} \int_{\mathbf{y}} f(\mathbf{y} | \boldsymbol{\mu}, \sigma^2)^{1+\beta} d\mathbf{y} \right\} \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2 \\
&= D_4 \int_{\sigma^2} (\sigma^2)^{-0.5d\beta - \hat{a}_n - 1} \frac{\hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n)} \exp \left\{ -\frac{\hat{b}_n}{\sigma^2} \right\} d\sigma^2 \\
&\stackrel{I(\mathcal{N})}{=} D_4 \cdot \frac{\Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n) \hat{b}_n^{0.5d\beta}}. \tag{37}
\end{aligned}$$

For the β -exponentiated likelihood term, one finds that

$$\begin{aligned}
L_{2,i} &= \int_{\sigma^2, \boldsymbol{\mu}} \frac{1}{\beta} f(\mathbf{y}_i | \boldsymbol{\mu}, \sigma^2)^\beta \text{NIG}^{\text{VB}}(\boldsymbol{\mu}, \sigma^2 | \hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\boldsymbol{\Sigma}}_n) d\sigma^2 d\boldsymbol{\mu} \\
&= \int_{\sigma^2} D_1(\sigma^2) \cdot \exp \left\{ -\frac{1}{2\sigma^2} D_{2,i} \right\} \underbrace{\left[\int_{\boldsymbol{\mu}} \exp \left\{ -\frac{1}{2\sigma^2} D_{3,i}(\boldsymbol{\mu}) \right\} \mathcal{N}^{\text{VB}}(\boldsymbol{\mu} | \hat{\boldsymbol{\mu}}_n, \sigma^2 \hat{\boldsymbol{\Sigma}}_n) d\boldsymbol{\mu} \right]}_{=D_{5,i}(\sigma^2)} \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2,
\end{aligned}$$

where we have again exploited the conditional structure of our assumed posterior. The inner integral equals

$$D_{5,i}(\sigma^2) = (2\pi)^{-0.5p} \left| \sigma^2 \hat{\boldsymbol{\Sigma}}_n \right|^{-0.5} \underbrace{\int_{\boldsymbol{\mu}} \exp \left\{ -\frac{1}{2\sigma^2} \left[\underbrace{D_{3,i}(\boldsymbol{\mu}) + (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n)' \hat{\boldsymbol{\Sigma}}_n^{-1} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n)}_{=D_{6,i}(\boldsymbol{\mu})} \right] \right\}}_{=D_{7,i}(\sigma^2)} d\boldsymbol{\mu} \tag{38}$$

indicating that the closed form for the integral is available if one rewrites it as a normal density. To this end, one can use completion of squares to rewrite

$$\begin{aligned}
D_{6,i}(\boldsymbol{\mu}) &= \beta(\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i)' \tilde{\boldsymbol{\Sigma}}_i^{-1} (\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i) + (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n)' \hat{\boldsymbol{\Sigma}}_n^{-1} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_n) \\
&= \boldsymbol{\mu}' \underbrace{\left[\hat{\boldsymbol{\Sigma}}_n^{-1} + \beta \tilde{\boldsymbol{\Sigma}}_i^{-1} \right]}_{=\tilde{\boldsymbol{\Sigma}}_i^{-1}} \boldsymbol{\mu} - 2 \left[\hat{\boldsymbol{b}}_n' \hat{\boldsymbol{\Sigma}}_n^{-1} + \beta \tilde{\boldsymbol{\mu}}_i' \tilde{\boldsymbol{\Sigma}}_i^{-1} \right] \boldsymbol{\mu} + \left[\hat{\boldsymbol{\mu}}_n' \hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \tilde{\boldsymbol{\mu}}_i' \tilde{\boldsymbol{\Sigma}}_i^{-1} \tilde{\boldsymbol{\mu}}_i \right] \\
&\equiv \left(\boldsymbol{\mu} - \underbrace{\tilde{\boldsymbol{\Sigma}}_i \left[\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{b}}_n + \beta \tilde{\boldsymbol{\Sigma}}_i^{-1} \tilde{\boldsymbol{\mu}}_i \right]}_{=\tilde{\boldsymbol{\mu}}_i} \right)' \tilde{\boldsymbol{\Sigma}}_i^{-1} (\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i) + \\
&\quad \underbrace{\hat{\boldsymbol{\mu}}_n' \hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \tilde{\boldsymbol{\mu}}_i' \tilde{\boldsymbol{\Sigma}}_i^{-1} \tilde{\boldsymbol{\mu}}_i - \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \tilde{\boldsymbol{\Sigma}}_i^{-1} \tilde{\boldsymbol{\mu}}_i \right)' \tilde{\boldsymbol{\Sigma}}_i \left(\hat{\boldsymbol{\Sigma}}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \tilde{\boldsymbol{\Sigma}}_i^{-1} \tilde{\boldsymbol{\mu}}_i \right)}_{=D_{8,i}} \\
&= (\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i)' \tilde{\boldsymbol{\Sigma}}_i^{-1} (\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i) + D_{8,i}, \tag{39}
\end{aligned}$$

which then allows integrating out $\boldsymbol{\mu}$ from $D_{7,i}(\sigma^2)$ using the density of a normal random variable:

$$\begin{aligned}
D_{7,i}(\sigma^2) &= \exp \left\{ -\frac{1}{2\sigma^2} D_{8,i} \right\} \int_{\boldsymbol{\mu}} \exp \left\{ -\frac{1}{2\sigma^2} (\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i)' \tilde{\boldsymbol{\Sigma}}_i^{-1} (\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}_i) \right\} d\boldsymbol{\mu} \\
&\stackrel{I(\mathcal{N})}{=} \exp \left\{ -\frac{1}{2\sigma^2} D_{8,i} \right\} (2\pi)^{0.5p} |\sigma^2 \tilde{\boldsymbol{\Sigma}}_i|^{0.5}, \tag{40}
\end{aligned}$$

so we can finally rewrite the entire integral as

$$D_{5,i}(\sigma^2) = |\hat{\boldsymbol{\Sigma}}_n^{-1} \tilde{\boldsymbol{\Sigma}}_i|^{0.5} \exp \left\{ -\frac{1}{2\sigma^2} D_{8,i} \right\}, \tag{41}$$

which enables rewriting $L_{2,i}$ as

$$\begin{aligned}
L_{2,i} &= \underbrace{\frac{1}{\beta} (2\pi)^{-0.5d\beta} |\hat{\boldsymbol{\Sigma}}_n^{-1} \tilde{\boldsymbol{\Sigma}}_i|^{0.5}}_{=D_{9,i}} \int_{\sigma^2} (\sigma^2)^{-0.5d\beta} \exp \left\{ -\frac{1}{\sigma^2} \cdot \frac{1}{2} [D_{2,i} + D_{8,i}] \right\} \mathcal{IG}^{\text{VB}}(\sigma^2 | \hat{a}_n, \hat{b}_n) d\sigma^2 \\
&\stackrel{I(\mathcal{IG})}{=} \frac{D_{9,i} \cdot \Gamma(\hat{a}_n + 0.5d\beta) \cdot \hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n) \cdot \left[\hat{b}_n + 0.5(D_{2,i} + D_{8,i}) \right]^{(\hat{a}_n + 0.5d\beta)}}, \tag{42}
\end{aligned}$$

finally implying that one may write

$$\begin{aligned}
Q_2 &= \sum_{i=1}^n L_{2,i} - nL_1 \\
&= \sum_{i=1}^n \left\{ \frac{D_{9,i} \cdot \Gamma(\hat{a}_n + 0.5d\beta) \cdot \hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n) \cdot \left[\hat{b}_n + 0.5(D_{2,i} + D_{8,i}) \right]^{(\hat{a}_n + 0.5d\beta)}} \right\} - nD_4 \cdot \frac{\Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n) \hat{b}_n^{0.5d\beta}} \\
&= \sum_{i=1}^n \left\{ \frac{\frac{1}{\beta} (2\pi)^{-0.5d\beta} \left| \hat{\boldsymbol{\Sigma}}_n^{-1} \left[\hat{\boldsymbol{\Sigma}}_n^{-1} + \beta(\mathbf{X}_i \mathbf{X}_i) \right]^{-1} \right|^{0.5} \cdot \Gamma(\hat{a}_n + 0.5d\beta) \cdot \hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n) \cdot \left[\hat{b}_n + 0.5(D_{2,i} + D_{8,i}) \right]^{(\hat{a}_n + 0.5d\beta)}} \right\} \\
&\quad - n \cdot \frac{(2\pi)^{-0.5d\beta} (1 + \beta)^{-0.5d-1} \cdot \Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n) \hat{b}_n^{0.5d\beta}}.
\end{aligned}$$

We further simplify this expression by observing that

$$\begin{aligned}
D_{2,i} + D_{8,i} &= \beta \left[\mathbf{y}'_i \mathbf{y}_i - (\mathbf{y}_i \mathbf{X}'_i) \check{\Sigma}_i (\mathbf{X}_i \mathbf{y}'_i) \right] + \hat{\boldsymbol{\mu}}'_n \hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \check{\boldsymbol{\mu}}'_i \check{\Sigma}_i^{-1} \check{\boldsymbol{\mu}}_i \\
&\quad - \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \check{\Sigma}_i^{-1} \check{\boldsymbol{\mu}}_i \right)' \check{\Sigma}_i \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta \check{\Sigma}_i^{-1} \check{\boldsymbol{\mu}}_i \right) \\
&= \beta \mathbf{y}'_i \mathbf{y}_i - \beta (\mathbf{y}_i \mathbf{X}'_i) \check{\Sigma}_i (\mathbf{X}_i \mathbf{y}'_i) + \hat{\boldsymbol{\mu}}'_n \hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{y}_i \mathbf{X}'_i) \check{\Sigma}_i (\mathbf{X}_i \mathbf{y}'_i) \\
&\quad - \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right)' \check{\Sigma}_i \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right) \\
&= \beta \mathbf{y}'_i \mathbf{y}_i + \hat{\boldsymbol{\mu}}'_n \hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n - \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right)' \left[\hat{\Sigma}_n^{-1} + \beta (\mathbf{X}_i \mathbf{X}_i) \right]^{-1} \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right),
\end{aligned}$$

leaving us with

$$\begin{aligned}
Q_2 &= \frac{\Gamma(\hat{a}_n + 0.5d\beta) \cdot \hat{b}_n^{\hat{a}_n} \cdot |\hat{\Sigma}_n^{-1}|^{0.5}}{\beta(2\pi)^{0.5d\beta} \Gamma(\hat{a}_n)} \times \\
&\sum_{i=1}^n \left\{ \frac{\left| \left[\hat{\Sigma}_n^{-1} + \beta (\mathbf{X}_i \mathbf{X}_i) \right]^{-1} \right|^{0.5}}{\left[\hat{b}_n + 0.5 \left(\beta \mathbf{y}'_i \mathbf{y}_i + \hat{\boldsymbol{\mu}}'_n \hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n - \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right)' \left[\hat{\Sigma}_n^{-1} + \beta (\mathbf{X}_i \mathbf{X}_i) \right]^{-1} \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right) \right) \right]^{\hat{a}_n + 0.5d\beta}} \right\} \\
&\quad - n \cdot \frac{\Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n) \hat{b}_n^{0.5d\beta} (2\pi)^{0.5d\beta} (1 + \beta)^{0.5d+1}}.
\end{aligned}$$

3.3 ELBO

Putting together the results of the two previous sections, the ELBO is obtained as

$$\begin{aligned}
ELBO &= -Q_1 + Q_2 \\
&= -\log \left(\frac{\hat{b}_n^{\hat{a}_n} \Gamma(a_0)}{\hat{b}_0^{a_0} \Gamma(\hat{a}_n)} \right) - 0.5 \log \left| \Sigma_0 \hat{\Sigma}_n^{-1} \right| + \frac{1}{2} \text{tr} \left(I - \Sigma_0^{-1} \hat{\Sigma}_n \right) - (\hat{a}_n - a_0) \left(\Psi(\hat{a}_n) - \log(\hat{b}_n) \right) \\
&\quad - \left[\frac{1}{2} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n)' \Sigma_0^{-1} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n) + (b_0 - \hat{b}_n) \right] \cdot \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)} \\
&\quad + \frac{\Gamma(\hat{a}_n + 0.5d\beta) \cdot \hat{b}_n^{\hat{a}_n} \cdot |\hat{\Sigma}_n^{-1}|^{0.5}}{\beta(2\pi)^{0.5d\beta} \Gamma(\hat{a}_n)} \times \\
&\sum_{i=1}^n \left\{ \frac{\left| \hat{\Sigma}_n^{-1} + \beta (\mathbf{X}'_i \mathbf{X}_i) \right|^{-0.5}}{\left[\hat{b}_n + 0.5 \left(\beta \mathbf{y}'_i \mathbf{y}_i + \hat{\boldsymbol{\mu}}'_n \hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n - \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right)' \left[\hat{\Sigma}_n^{-1} + \beta (\mathbf{X}_i \mathbf{X}_i) \right]^{-1} \left(\hat{\Sigma}_n^{-1} \hat{\boldsymbol{\mu}}_n + \beta (\mathbf{X}'_i \mathbf{y}_i) \right) \right) \right]^{\hat{a}_n + 0.5d\beta}} \right\} \\
&\quad - n \cdot \frac{\Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n) \hat{b}_n^{0.5d\beta} (2\pi)^{0.5d\beta} (1 + \beta)^{0.5d+1}} \tag{43}
\end{aligned}$$

3.4 Differentiation

In this section, we take derivatives of the ELBO with respect to each variational parameter, i.e. $\hat{a}_n, \hat{b}_n, \hat{\boldsymbol{\mu}}_n, \hat{\Sigma}_n$. Observing that differentiation with respect to $\hat{\Sigma}_n^{-1}$ is easier than with respect to $\hat{\Sigma}_n$, parametrize the optimization using the Cholesky decomposition, i.e. $\hat{\Sigma}_n^{-1} = \mathbf{L}\mathbf{L}'$, where \mathbf{L} is a lower triangular matrix and is unique if $\hat{\Sigma}_n$ (equivalently $\hat{\Sigma}_n^{-1}$) is positive definite¹.

¹Note that \mathbf{L} need not be unique if $\hat{\Sigma}_n$ is positive semi-definite, but this is of no concern for us here: Since we implicitly impose that $\hat{\Sigma}_n$ is non-singular (so that $\hat{\Sigma}_n^{-1}$ is unique and well-defined), all covariance matrices $\hat{\Sigma}_n$ considered have to be positive definite.

3.4.1 Derivative with respect to L

In what follows, we differentiate the ELBO term by term with respect to the $p(p-1)\frac{1}{2}$ entries in the lower triangular part of L that can be summarized in the vector $\text{vech}(L)$. To this end, define

$$E_1 = -0.5 \log \left| \Sigma_0 \widehat{\Sigma}_n^{-1} \right| + \frac{1}{2} \text{tr} \left(I - \Sigma_0^{-1} \widehat{\Sigma}_n \right) \quad (44)$$

$$E_2 = \frac{\Gamma(\widehat{a}_n + 0.5d\beta) \cdot \widehat{b}_n^{\widehat{a}_n}}{\underbrace{\beta(2\pi)^{0.5d\beta} \Gamma(\widehat{a}_n)}_{=F}} \left| \widehat{\Sigma}_n^{-1} \right|^{0.5} \quad (45)$$

$$E_{3,i} = \left| \widehat{\Sigma}_n^{-1} + \beta (\mathbf{X}'_i \mathbf{X}_i) \right|^{-0.5} \quad (46)$$

$$E_4 = \widehat{\boldsymbol{\mu}}'_n \widehat{\Sigma}_n^{-1} \widehat{\boldsymbol{\mu}}_n \quad (47)$$

$$E_{5,i} = -\widehat{\boldsymbol{\mu}}'_n \widehat{\Sigma}_n^{-1} \left[\widehat{\Sigma}_n^{-1} + \beta (\mathbf{X}'_i \mathbf{X}_i) \right]^{-1} \widehat{\Sigma}_n^{-1} \widehat{\boldsymbol{\mu}}_n \quad (48)$$

$$E_{6,i} = -\beta^2 (\mathbf{y}'_i \mathbf{X}_i) \left[\widehat{\Sigma}_n^{-1} + \beta (\mathbf{X}'_i \mathbf{X}_i) \right]^{-1} (\mathbf{X}'_i \mathbf{y}_i), \quad (49)$$

$$E_{7,i} = -2\beta \widehat{\boldsymbol{\mu}}'_n \widehat{\Sigma}_n^{-1} \left[\widehat{\Sigma}_n^{-1} + \beta (\mathbf{X}'_i \mathbf{X}_i) \right]^{-1} (\mathbf{X}'_i \mathbf{y}_i). \quad (50)$$

Obtaining the derivative of the ELBO is equivalent to obtaining the derivatives of these newly defined quantities, as

$$\begin{aligned} \frac{\partial}{\partial \text{vech}(L)} \{ELBO\} &= \frac{\partial}{\partial \text{vech}(L)} \{E_1\} \\ &+ \frac{\partial}{\partial \text{vech}(L)} \{E_2\} \cdot \sum_{i=1}^n \left\{ \frac{E_{3,i}}{\left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{\widehat{a}_n + 0.5d\beta}} \right\} \\ &+ E_2 \cdot \sum_{i=1}^n \left\{ \frac{\frac{\partial}{\partial \text{vech}(L)} \{E_{3,i}\}}{\left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{\widehat{a}_n + 0.5d\beta}} \right\} \\ &+ E_2 \cdot \sum_{i=1}^n \left\{ E_{3,i} \cdot \frac{\partial}{\partial \text{vech}(L)} \left\{ \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta} \right\} \right\} \end{aligned} \quad (51)$$

where the chain and sum rule imply that

$$\begin{aligned} &\frac{\partial}{\partial \text{vech}(L)} \left\{ \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta} \right\} \\ &= (-\widehat{a}_n - 0.5d\beta) \cdot \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta - 1} \times \\ &\quad 0.5 \cdot \frac{\partial}{\partial \text{vech}(L)} \{E_4 + E_{5,i} + E_{6,i} + E_{7,i}\}, \end{aligned} \quad (52)$$

For convenience and simplified notation when taking the derivatives of the expressions defined in Eqs (44) - (50), also define the following matrices:

$$\mathbf{R} = \left[\widehat{\Sigma}_n^{-1} + \beta (\mathbf{X}'_i \mathbf{X}_i) \right] \quad (53)$$

$$\mathbf{B} = \widehat{\boldsymbol{\mu}}_n \widehat{\boldsymbol{\mu}}'_n. \quad (54)$$

Define also the following symbols to mark operations used in the derivations:

∂ Switching from differential notation ∂L to the derivative $\frac{\partial}{\partial \text{vech}(L)}$;

tr Properties of the trace like invariance under cyclic permutations, invariance under the transpose, additivity, and the fact that for c a scalar, $\text{tr}(c) = c$.

Note that when the differential operator ∂ is used, its scope is always limited to the next term only, unless brackets are used. Hence $\partial \mathbf{LL}'$ uses the differential only with respect to \mathbf{L} , while $\partial (\mathbf{LL}')^{-1}$ uses it with respect to the entire expression $(\mathbf{LL}')^{-1}$. It is also worth noting that $\partial \mathbf{L}' = (\partial \mathbf{L})'$ for any matrix \mathbf{L} , as this will be used in conjunction with the transpose invariance of the trace throughout to simplify terms. Using these symbols and the differential notation, proceed by noting the following:

$$\partial(\mathbf{LL}') = \partial \mathbf{R} = \partial \mathbf{LL}' + \mathbf{L} \partial \mathbf{L}' = \partial \mathbf{LL}' + \mathbf{L} \partial \mathbf{L}' = \partial \mathbf{LL}' + (\partial \mathbf{LL}')' \quad (55)$$

$$\partial(\mathbf{LL}')^{-1} = -(\mathbf{LL}')^{-1} [\partial(\mathbf{LL}')] (\mathbf{LL}')^{-1} \quad (56)$$

$$\begin{aligned} \partial |\mathbf{LL}'| &= |\mathbf{LL}'| \cdot \text{tr} \left((\mathbf{LL}')^{-1} [\partial \mathbf{LL}' + (\partial \mathbf{LL}')'] \right) \\ &\stackrel{\text{tr}}{=} 2 |\mathbf{LL}'| \cdot \text{tr} (\mathbf{L}' (\mathbf{LL}')^{-1} \partial \mathbf{L}) \end{aligned} \quad (57)$$

$$\partial \mathbf{R}^{-1} = -\mathbf{R}^{-1} \partial \mathbf{R} \mathbf{R}^{-1} = -\mathbf{R}^{-1} \partial \mathbf{LL}' \mathbf{R}^{-1} - [\mathbf{R}^{-1} \partial \mathbf{LL}' \mathbf{R}^{-1}]'. \quad (58)$$

With this in place, the derivatives of the quantities defined before are obtained as

$$\begin{aligned} \partial E_1 &= -\frac{1}{2} \partial \{ \log |\boldsymbol{\Sigma}_0| + \log |\mathbf{LL}'| \} - \frac{1}{2} \partial \left\{ \text{tr} \left(\boldsymbol{\Sigma}_0^{-1} \hat{\boldsymbol{\Sigma}}_n \right) \right\} \\ &= -\frac{1}{2} \cdot |\mathbf{LL}'|^{-1} \cdot \partial |\mathbf{LL}'| - \frac{1}{2} \text{tr} \left(\boldsymbol{\Sigma}_0^{-1} \partial (\mathbf{LL}')^{-1} \right) \\ &= -\frac{1}{2} \text{tr} (\mathbf{L}' (\mathbf{LL}')^{-1} \partial \mathbf{L}) + \frac{1}{2} \text{tr} (\boldsymbol{\Sigma}_0^{-1} (\mathbf{LL}')^{-1} [\partial(\mathbf{LL}')] (\mathbf{LL}')^{-1}) \\ &\stackrel{\text{tr}}{=} -\text{tr} (\mathbf{L}' (\mathbf{LL}')^{-1} \partial \mathbf{L}) + \text{tr} \left(\mathbf{L}' (\mathbf{LL}')^{-1} \boldsymbol{\Sigma}_0^{-1} (\mathbf{LL}')^{-1} \partial \mathbf{L} \right) \\ \partial E_2 &= F \cdot \partial |\mathbf{LL}'|^{0.5} \\ &= \frac{F}{2} \cdot |\mathbf{LL}'|^{-0.5} \cdot 2 |\mathbf{LL}'| \cdot \text{tr} (\mathbf{L}' (\mathbf{LL}')^{-1} \partial \mathbf{L}) \\ &= F \cdot |\mathbf{LL}'|^{0.5} \text{tr} (\mathbf{L}' (\mathbf{LL}')^{-1} \partial \mathbf{L}) \end{aligned} \quad (59)$$

$$\begin{aligned} \partial E_{3,i} &= \partial \mathbf{R}^{-0.5} = -\frac{1}{2} |\mathbf{R}|^{-1.5} \partial \mathbf{R} \\ &= -\frac{1}{2} |\mathbf{R}|^{-0.5} \text{tr} (\mathbf{R}^{-1} \partial (\mathbf{LL}')) \\ &\stackrel{\text{tr}}{=} -|\mathbf{R}|^{-0.5} \text{tr} (\mathbf{L}' \mathbf{R}^{-1} \partial \mathbf{L}) \end{aligned} \quad (60)$$

$$\begin{aligned} \partial E_4 &\stackrel{\text{tr}}{=} \text{tr} (\hat{\boldsymbol{\mu}}_n' \partial (\mathbf{LL}') \hat{\boldsymbol{\mu}}_n) \\ &= \text{tr} (\hat{\boldsymbol{\mu}}_n' [\partial \mathbf{LL}' + \mathbf{L} \partial \mathbf{L}'] \hat{\boldsymbol{\mu}}_n) \\ &\stackrel{\text{tr}}{=} 2 \cdot \text{tr} (\mathbf{L}' \mathbf{B} \partial \mathbf{L}) \end{aligned} \quad (61)$$

$$\begin{aligned} \partial E_{5,i} &\stackrel{\text{tr}}{=} -\text{tr} (\hat{\boldsymbol{\mu}}_n' \partial (\mathbf{LL}') \mathbf{R}^{-1} (\mathbf{LL}') \hat{\boldsymbol{\mu}}_n) - \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \partial \mathbf{R}^{-1} (\mathbf{LL}') \hat{\boldsymbol{\mu}}_n) \\ &\quad - \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \mathbf{R}^{-1} \partial (\mathbf{LL}') \hat{\boldsymbol{\mu}}_n) \\ &\stackrel{\text{tr}}{=} -2 \cdot \text{tr} (\hat{\boldsymbol{\mu}}_n' \partial \mathbf{LL}' \mathbf{R}^{-1} (\mathbf{LL}') \hat{\boldsymbol{\mu}}_n) + 2 \cdot \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \mathbf{R}^{-1} \partial \mathbf{LL}' \mathbf{R}^{-1} (\mathbf{LL}') \hat{\boldsymbol{\mu}}_n) \\ &\quad - 2 \cdot \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \mathbf{R}^{-1} \partial \mathbf{LL}' \hat{\boldsymbol{\mu}}_n) \\ &\stackrel{\text{tr}}{=} -2 \cdot \text{tr} (\mathbf{L}' \mathbf{R}^{-1} (\mathbf{LL}') \mathbf{B} \partial \mathbf{L}) + 2 \cdot \text{tr} (\mathbf{L}' \mathbf{R}^{-1} (\mathbf{LL}') \mathbf{B} (\mathbf{LL}') \mathbf{R}^{-1} \partial \mathbf{L}) \\ &\quad - 2 \cdot \text{tr} (\mathbf{L}' \mathbf{B} (\mathbf{LL}') \mathbf{R}^{-1} \partial \mathbf{L}) \end{aligned}$$

$$\begin{aligned} \partial E_{6,i} &\stackrel{\text{tr}}{=} -\beta^2 \text{tr} ((\mathbf{y}'_i \mathbf{X}_i) \partial \mathbf{R}^{-1} (\mathbf{X}'_i \mathbf{y}_i)) \\ &\stackrel{\text{tr}}{=} 2\beta^2 \text{tr} ((\mathbf{y}'_i \mathbf{X}_i) \mathbf{R}^{-1} \partial \mathbf{LL}' \mathbf{R}^{-1} (\mathbf{X}'_i \mathbf{y}_i)) \\ &\stackrel{\text{tr}}{=} 2\beta^2 \text{tr} (\mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}'_i \mathbf{y}_i) (\mathbf{y}'_i \mathbf{X}_i) \mathbf{R}^{-1} \partial \mathbf{L}) \end{aligned} \quad (62)$$

$$\begin{aligned} \partial E_{7,i} &\stackrel{\text{tr}}{=} -2\beta \cdot \left[\text{tr} (\hat{\boldsymbol{\mu}}_n' \partial (\mathbf{LL}') \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i)) + \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \partial \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i)) \right] \\ &\stackrel{\text{tr}}{=} -2\beta \cdot \left[\text{tr} (\hat{\boldsymbol{\mu}}_n' \partial \mathbf{LL}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i)) + \text{tr} (\hat{\boldsymbol{\mu}}_n' \mathbf{L} \partial \mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i)) \right. \\ &\quad \left. - \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \mathbf{R}^{-1} \partial \mathbf{LL}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i)) - \text{tr} (\hat{\boldsymbol{\mu}}_n' (\mathbf{LL}') \mathbf{R}^{-1} \mathbf{L} \partial \mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i)) \right] \\ &\stackrel{\text{tr}}{=} -2\beta \cdot \left[\text{tr} (\mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i) \hat{\boldsymbol{\mu}}_n' \partial \mathbf{L}) + \text{tr} (\mathbf{L}' \hat{\boldsymbol{\mu}}_n (\mathbf{y}'_i \mathbf{X}_i) \mathbf{R}^{-1} \partial \mathbf{L}) \right] \end{aligned}$$

$$-\text{tr} \left(\mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i) \hat{\boldsymbol{\mu}}_n' (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} \partial \mathbf{L} \right) - \text{tr} \left(\mathbf{L}' \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \hat{\boldsymbol{\mu}}_n (\mathbf{y}_i' \mathbf{X}_i) \mathbf{R}^{-1} \partial \mathbf{L} \right) \quad (63)$$

This can now be converted into derivative notation and simplified. To this end, first note that for any $p \times \partial$ matrix A which is not a function of \mathbf{L} ,

$$\text{tr}(A d\mathbf{L}) = \sum_{i=1}^p A_{1i} dL_{i1} + \sum_{i=2}^p A_{2i} dL_{i2} + \dots = \sum_{j=1}^p \left\{ \sum_{i=j}^p A_{ji} dL_{ji} \right\}, \quad (64)$$

implying in particular that

$$\frac{\partial}{\partial \text{vech}(\mathbf{L})} \text{tr}(A d\mathbf{L}) = \text{vech}(A^T) \quad (65)$$

and use this by defining $\text{vech}^T(A) = \text{vech}(A^T)$ to note that

$$\begin{aligned} \frac{\partial}{\partial \text{vech}(\mathbf{L})} E_1 &\stackrel{\partial}{=} \text{vech}^T \left(-[\mathbf{L}' (\mathbf{L} \mathbf{L}')^{-1}] + [\mathbf{L}' (\mathbf{L} \mathbf{L}')^{-1} \boldsymbol{\Sigma}_0^{-1} (\mathbf{L} \mathbf{L}')^{-1}] \right) \\ &= \text{vech}^T \left(\mathbf{L}' (\mathbf{L} \mathbf{L}')^{-1} \left[\boldsymbol{\Sigma}_0^{-1} (\mathbf{L} \mathbf{L}')^{-1} - I_p \right] \right) \\ &= \text{vech}^T \left(\mathbf{L}^{-1} \left[\boldsymbol{\Sigma}_0^{-1} (\mathbf{L} \mathbf{L}')^{-1} - I_p \right] \right) \\ &= \text{vech} \left(\left[(\mathbf{L} \mathbf{L}')^{-1} \boldsymbol{\Sigma}_0^{-1} - I_p \right] \mathbf{L}^{-T} \right) \end{aligned} \quad (66)$$

$$\begin{aligned} \frac{\partial}{\partial \text{vech}(\mathbf{L})} E_2 &\stackrel{\partial}{=} F \cdot |\mathbf{L} \mathbf{L}'|^{0.5} \cdot \text{vech}^T (\mathbf{L}' (\mathbf{L} \mathbf{L}')^{-1}) \\ &= F \cdot |\mathbf{L} \mathbf{L}'|^{0.5} \cdot \text{vech} (\mathbf{L}^{-T}) \end{aligned} \quad (67)$$

$$\frac{\partial}{\partial \text{vech}(\mathbf{L})} E_{3,i} \stackrel{\partial}{=} -|\mathbf{R}|^{-0.5} \cdot \text{vech} (\mathbf{R}^{-1} \mathbf{L}) \quad (68)$$

$$\frac{\partial}{\partial \text{vech}(\mathbf{L})} E_4 \stackrel{\partial}{=} 2 \cdot \text{vech} (\mathbf{B} \mathbf{L}) \quad (69)$$

$$\begin{aligned} \frac{\partial}{\partial \text{vech}(\mathbf{L})} E_{5,i} &\stackrel{\partial}{=} \text{vech}^T \left(-2\mathbf{L}' \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \mathbf{B} + 2\mathbf{L}' \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \mathbf{B} (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} - 2\mathbf{L}' \mathbf{B} (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} \right) \\ &= 2 \cdot \text{vech}^T \left([\mathbf{L}' \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \mathbf{B} [(\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} - I_p]] - [\mathbf{L}' \mathbf{B} (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1}] \right) \\ &= 2 \cdot \text{vech}^T \left(\mathbf{L}' [\mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \mathbf{B} [(\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} - I_p] - \mathbf{B} (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1}] \right) \\ &= 2 \cdot \text{vech} \left([[\mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') - I_p] \mathbf{B} (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} - \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \mathbf{B}] \mathbf{L} \right) \end{aligned} \quad (70)$$

$$\frac{\partial}{\partial \text{vech}(\mathbf{L})} E_{6,i} \stackrel{\partial}{=} 2\beta^2 \cdot \text{vech} (\mathbf{R}^{-1} (\mathbf{X}_i' \mathbf{y}_i) (\mathbf{y}_i' \mathbf{X}_i) \mathbf{R}^{-1} \mathbf{L}) \quad (71)$$

$$\begin{aligned} \frac{\partial}{\partial \text{vech}(\mathbf{L})} E_{7,i} &\stackrel{\partial}{=} -2\beta \cdot \text{vech}^T \left(\mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i) \hat{\boldsymbol{\mu}}_n' + \mathbf{L}' \hat{\boldsymbol{\mu}}_n (\mathbf{y}_i' \mathbf{X}_i) \mathbf{R}^{-1} \right. \\ &\quad \left. - \mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i) \hat{\boldsymbol{\mu}}_n' (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1} - \mathbf{L}' \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}') \hat{\boldsymbol{\mu}}_n (\mathbf{y}_i' \mathbf{X}_i) \mathbf{R}^{-1} \right) \\ &= -2\beta \cdot \text{vech}^T \left(\mathbf{L}' \mathbf{R}^{-1} (\mathbf{X}_i \mathbf{y}_i) \hat{\boldsymbol{\mu}}_n' [I_p - (\mathbf{L} \mathbf{L}') \mathbf{R}^{-1}] \right. \\ &\quad \left. + [I_p - \mathbf{L}' \mathbf{R}^{-1} \mathbf{L}] \mathbf{L}' \hat{\boldsymbol{\mu}}_n (\mathbf{y}_i' \mathbf{X}_i) \mathbf{R}^{-1} \right) \\ &= -2\beta \cdot \text{vech} \left([I_p - \mathbf{R}^{-1} (\mathbf{L} \mathbf{L}')] \hat{\boldsymbol{\mu}}_n (\mathbf{y}_i' \mathbf{X}_i) \mathbf{R}^{-1} \mathbf{L} \right. \\ &\quad \left. + \mathbf{R}^{-1} (\mathbf{X}_i' \mathbf{y}_i) \hat{\boldsymbol{\mu}}_n' \mathbf{L} [I_p - \mathbf{L}' \mathbf{R}^{-1} \mathbf{L}] \right) \end{aligned} \quad (72)$$

3.4.2 Derivative with respect to $\hat{\boldsymbol{\mu}}_n$

Differentiating with respect to $\hat{\boldsymbol{\mu}}_n$ is trivial. One proceeds by the same logic as in the section before, to which end one additionally needs to define the new term

$$E_8 = -\frac{1}{2} \left[(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n)' \boldsymbol{\Sigma}_0^{-1} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n) + 2(b_0 - \hat{b}_n) \right] \cdot \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)}, \quad (73)$$

allowing us to write

$$\begin{aligned} \frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} \{ELBO\} &= \frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} \{E_8\} + \\ &E_2 \cdot \sum_{i=1}^n \left\{ E_{3,i} \cdot \frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} \left\{ \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta} \right\} \right\} \end{aligned} \quad (74)$$

where

$$\begin{aligned} &\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} \left\{ \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta} \right\} \\ &= (-\widehat{a}_n - 0.5d\beta) \cdot \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta - 1} \times \\ &\quad 0.5 \cdot \frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} \{E_4 + E_{5,i} + E_{7,i}\}, \end{aligned} \quad (75)$$

so that obtaining the derivative is achieved by finding $\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_4$, $\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_{5,i}$, $\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_{7,i}$ and $\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_8$:

$$\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_4 = 2 \cdot \widehat{\boldsymbol{\mu}}'_n \widehat{\boldsymbol{\Sigma}}_n^{-1} \quad (76)$$

$$\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_{5,i} = -2 \cdot \widehat{\boldsymbol{\mu}}'_n \widehat{\boldsymbol{\Sigma}}_n^{-1} \mathbf{R}^{-1} \widehat{\boldsymbol{\Sigma}}_n^{-1} \quad (77)$$

$$\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_{7,i} = -2\beta \cdot (\mathbf{y}'_i \mathbf{X}_i) \mathbf{R}^{-1} \widehat{\boldsymbol{\Sigma}}_n^{-1} \quad (78)$$

$$\begin{aligned} \frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} E_8 &= -\frac{1}{2} \cdot \frac{\Gamma(\widehat{a}_n + 1)}{\widehat{b}_n \Gamma(\widehat{a}_n)} \left[\frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} (\widehat{\boldsymbol{\mu}}'_n \boldsymbol{\Sigma}_0^{-1} \widehat{\boldsymbol{\mu}}_n) - 2 \frac{\partial}{\partial \widehat{\boldsymbol{\mu}}_n} (\widehat{\boldsymbol{\mu}}_n \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0) \right] \\ &= -\frac{1}{2} \cdot \frac{\Gamma(\widehat{a}_n + 1)}{\widehat{b}_n \Gamma(\widehat{a}_n)} [2\widehat{\boldsymbol{\mu}}'_n \boldsymbol{\Sigma}_0^{-1} - 2\boldsymbol{\mu}'_0 \boldsymbol{\Sigma}_0^{-1}] \\ &= \frac{\Gamma(\widehat{a}_n + 1)}{\widehat{b}_n \Gamma(\widehat{a}_n)} [(\boldsymbol{\mu}_0 - \widehat{\boldsymbol{\mu}}_n)' \boldsymbol{\Sigma}_0^{-1}] \end{aligned} \quad (79)$$

3.4.3 Derivative with respect to \widehat{a}_n

We proceed again by the same logic. Define

$$E_9 = -\log \left(\frac{\widehat{b}_n^{\widehat{a}_n} \Gamma(a_0)}{b_0^{a_0} \Gamma(\widehat{a}_n)} \right) \quad (80)$$

$$E_{10} = -(\widehat{a}_n - a_0) \left(\Psi(\widehat{a}_n) - \log(\widehat{b}_n) \right) \quad (81)$$

$$E_{11} = -n \cdot \frac{\Gamma(\widehat{a}_n + 0.5d\beta)}{\Gamma(\widehat{a}_n) \widehat{b}_n^{0.5d\beta} (2\pi)^{0.5d\beta} (1 + \beta)^{0.5d+1}}. \quad (82)$$

Use this to write

$$\begin{aligned} \frac{\partial}{\partial \widehat{a}_n} \{ELBO\} &= \frac{\partial}{\partial \widehat{a}_n} \{E_8\} + \frac{\partial}{\partial \widehat{a}_n} \{E_9\} + \frac{\partial}{\partial \widehat{a}_n} \{E_{10}\} + \frac{\partial}{\partial \widehat{a}_n} \{E_{11}\} + \\ &+ \frac{\partial}{\partial \widehat{a}_n} \{E_2\} \sum_{i=1}^n \left\{ \frac{E_{3,i}}{\left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{\widehat{a}_n + 0.5d\beta}} \right\} \\ &+ E_2 \cdot \sum_{i=1}^n \left\{ E_{3,i} \cdot \frac{\partial}{\partial \widehat{a}_n} \left\{ \left[\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\widehat{a}_n - 0.5d\beta} \right\} \right\} \end{aligned} \quad (83)$$

where for \widehat{a}_n , the inner term equals

$$\frac{\partial}{\partial \widehat{a}_n} \left\{ \left[\underbrace{\widehat{b}_n + 0.5 (\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i})}_{=K} \right]^{-\widehat{a}_n - 0.5d\beta} \right\} = -\log(K) \cdot K^{-\widehat{a}_n - 0.5d\beta}, \quad (84)$$

so that the differentiation with respect to \hat{a}_n requires obtaining the following terms:

$$\begin{aligned} \frac{\partial}{\partial \hat{a}_n} E_2 &= \frac{|\hat{\Sigma}_n^{-1}|^{0.5}}{\beta(2\pi)^{0.5d\beta}} \left[\frac{\frac{\partial}{\partial \hat{a}_n} \{\Gamma(\hat{a}_n + 0.5d\beta)\} \hat{b}_n^{\hat{a}_n}}{\Gamma(\hat{a}_n)} + \frac{\frac{\partial}{\partial \hat{a}_n} \{\hat{b}_n^{\hat{a}_n}\} \Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n)} \right. \\ &\quad \left. + \frac{\partial}{\partial \hat{a}_n} \{\Gamma(\hat{a}_n)^{-1}\} \cdot \hat{b}_n^{\hat{a}_n} \Gamma(\hat{a}_n + 0.5d\beta) \right] \\ &= \frac{|\hat{\Sigma}_n^{-1}|^{0.5} \hat{b}_n^{\hat{a}_n} \Gamma(\hat{a}_n + 0.5d\beta)}{\beta(2\pi)^{0.5d\beta} \Gamma(\hat{a}_n)} \left[\Psi(\hat{a}_n + 0.5d\beta) + \log(\hat{b}_n) - \Psi(\hat{a}_n) \right] \end{aligned} \quad (85)$$

$$\begin{aligned} \frac{\partial}{\partial \hat{a}_n} E_8 &= -\frac{1}{2} \left[(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n)' \boldsymbol{\Sigma}_0^{-1} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n) + 2(b_0 - \hat{b}_n) \right] \cdot \left[\frac{\frac{\partial}{\partial \hat{a}_n} \{\Gamma(\hat{a}_n + 1)\}}{\hat{b}_n \Gamma(\hat{a}_n)} - \frac{\frac{\partial}{\partial \hat{a}_n} \{\Gamma(\hat{a}_n)\} \Gamma(\hat{a}_n + 1)}{\Gamma(\hat{a}_n)^2 \hat{b}_n} \right] \\ &= -\frac{1}{2} \left[(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n)' \boldsymbol{\Sigma}_0^{-1} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n) + 2(b_0 - \hat{b}_n) \right] \cdot \frac{\Gamma(\hat{a}_n + 1)}{\hat{b}_n \Gamma(\hat{a}_n)} \cdot [\Psi(\hat{a}_n + 1) - \Psi(\hat{a}_n)] \end{aligned} \quad (86)$$

$$\begin{aligned} \frac{\partial}{\partial \hat{a}_n} E_9 &= -\frac{\partial}{\partial \hat{a}_n} \{\hat{a}_n \log(\hat{b}_n)\} + \frac{\partial}{\partial \hat{a}_n} \{\log(\Gamma(\hat{a}_n))\} \\ &= -\log(\hat{b}_n) + \Psi(\hat{a}_n) \end{aligned} \quad (87)$$

$$\begin{aligned} \frac{\partial}{\partial \hat{a}_n} E_{10} &= \frac{\partial}{\partial \hat{a}_n} \{\hat{a}_n \log(\hat{b}_n)\} - \frac{\partial}{\partial \hat{a}_n} \{(\hat{a}_n - a_0) \Psi(\hat{a}_n)\} \\ &= \log(\hat{b}_n) - \Psi(\hat{a}_n) - (\hat{a}_n - a_0) \Psi^{(1)}(\hat{a}_n) \end{aligned} \quad (88)$$

$$\begin{aligned} \frac{\partial}{\partial \hat{a}_n} E_{11} &= -\frac{n}{\hat{b}_n^{0.5d\beta} (2\pi)^{0.5d\beta} (1 + \beta)^{0.5d+1}} \cdot \frac{\partial}{\partial \hat{a}_n} \left\{ \frac{\Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n)} \right\} \\ &= -\frac{n}{\hat{b}_n^{0.5d\beta} (2\pi)^{0.5d\beta} (1 + \beta)^{0.5d+1}} \cdot \frac{\Gamma(\hat{a}_n + 0.5d\beta)}{\Gamma(\hat{a}_n)} \cdot [\Psi(\hat{a}_n + 0.5d\beta) - \Psi(\hat{a}_n)], \end{aligned} \quad (89)$$

where $\Psi^{(1)}$ denotes the trigamma function.

3.4.4 Derivative with respect to \hat{b}_n

As for the other variational parameters, note that

$$\begin{aligned} \frac{\partial}{\partial \hat{b}_n} \{ELBO\} &= \frac{\partial}{\partial \hat{b}_n} \{E_8\} + \frac{\partial}{\partial \hat{b}_n} \{E_9\} + \frac{\partial}{\partial \hat{b}_n} \{E_{10}\} + \frac{\partial}{\partial \hat{b}_n} \{E_{11}\} + \\ &\quad + \frac{\partial}{\partial \hat{b}_n} \{E_2\} \sum_{i=1}^n \left\{ \frac{E_{3,i}}{\left[\hat{b}_n + 0.5(\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{\hat{a}_n + 0.5d\beta}} \right\} \\ &\quad + E_2 \cdot \sum_{i=1}^n \left\{ E_{3,i} \cdot \frac{\partial}{\partial \hat{b}_n} \left\{ \left[\hat{b}_n + 0.5(\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i}) \right]^{-\hat{a}_n - 0.5d\beta} \right\} \right\} \end{aligned} \quad (90)$$

where the chain rule implies that

$$\frac{\partial}{\partial \hat{b}_n} \left\{ \left[\underbrace{\hat{b}_n + 0.5(\beta \mathbf{y}'_i \mathbf{y}_i + E_4 + E_{5,i} + E_{6,i} + E_{7,i})}_{=K} \right]^{-\hat{a}_n - 0.5d\beta} \right\} = (-\hat{a}_n - 0.5d\beta) \cdot K^{-\hat{a}_n - 0.5d\beta - 1}. \quad (91)$$

Thus one proceeds by the same logic as before.

$$\frac{\partial}{\partial \hat{b}_n} E_2 = \frac{\hat{a}_n \Gamma(\hat{a}_n + 0.5d\beta) \cdot |\hat{\Sigma}_n^{-1}|^{0.5}}{\beta(2\pi)^{0.5d\beta} \Gamma(\hat{a}_n)} \cdot \hat{b}_n^{\hat{a}_n - 1} \quad (92)$$

$$\frac{\partial}{\partial \hat{b}_n} E_8 = \frac{1}{2} [(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n)' \boldsymbol{\Sigma}_0^{-1} (\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}_n) + 2b_0] \frac{\Gamma(\hat{a}_n + 1)}{\Gamma(\hat{a}_n)} \cdot \frac{1}{\hat{b}_n^2} \quad (93)$$

$$\frac{\partial}{\partial \hat{b}_n} E_9 = -\frac{\hat{a}_n}{\hat{b}_n} \quad (94)$$

$$\frac{\partial}{\partial \hat{b}_n} E_{10} = \frac{\hat{a}_n - a_0}{\hat{b}_n} \quad (95)$$

$$\frac{\partial}{\partial \hat{b}_n} E_{11} = \frac{nd\beta \cdot \Gamma(\hat{a}_n + 0.5d\beta)}{2 \cdot \Gamma(\hat{a}_n) (2\pi)^{0.5d\beta} (1 + \beta)^{0.5d+1}} \cdot \hat{b}_n^{-0.5d\beta - 1} \quad (96)$$

4 Timing and performance comparisons: Markov Chain Monte Carlo vs Structured Variational Bayes

We ran timing comparisons of SVI with MCMC for several subsets of the well-log data set. We ran the β -D BOCPD algorithm implementing an MCMC inference regime using *stan* [Carpenter et al., 2016] and compared this with our SVI inference regime. The two inference schemes were then run on 3 datasets of different time-length; the first 200 observations of the well-log, the first 500 observations of the well log and the full well-log, in order to show the impact changing the number of observations has on the timings for the algorithm. For the SVI used to produce these timings, we perform *full* optimization at every step, which is significantly slower than the SGD-variant that we present in the paper and that can be found in our repository at <https://github.com/alan-turing-institute/rbocpdms>. In spite of this, the SVI is still orders of magnitudes faster.

Table 1: Table of times to run the β -D BOCPD algorithm under the MCMC and SVI with full optimization on the first 200 observations, the first 500 observations and the full well log dataset.

	T=200	T=500	T=4050
MCMC	7615.2	20388.7	106073.0
SVI (full optimization)	102.8	328.5	3240.0

Another question of interest is how much the Stochastic Gradient Descent (SGD) inside our inference procedure provides robustness and how much the β -D itself is responsible for this. To put this question to the test, we ran full vs SGD-based optimization on the well-log data. As shown in Fig. 3, the results are very close to identical: No CPs are declared under one that are not declared in the other, and the run-length distribution’s maximum coincides throughout.

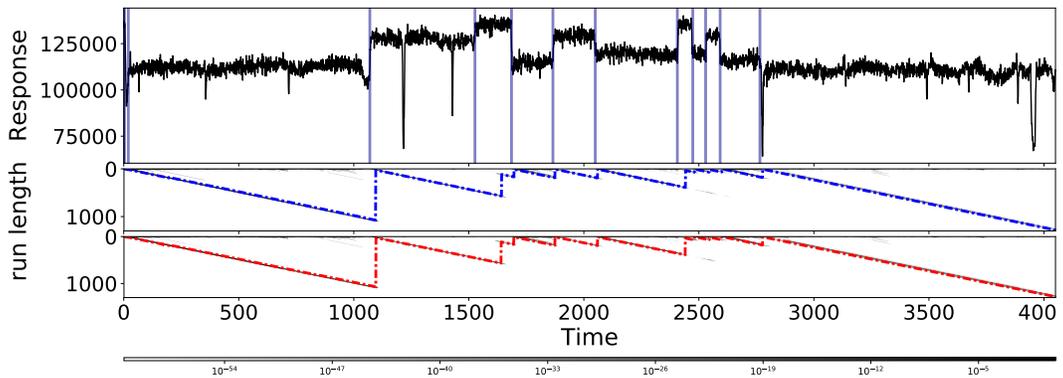


Figure 3: MAP segmentation and run-length distributions of the well-log data. **SGD inference** outcomes in blue, outcomes under **full optimization** in red. The corresponding run-length distributions for SGD (middle) and full (bottom) optimization are shown in grayscale with dashed maximum.

5 Initialization for β_p

The initialization procedure described in the paper is illustrated in Fig. 4. Here, the yellow dashed line gives a standard normal density corresponding to our model for the data. The gray dotted vertical line gives the amount of standard deviations from the posterior mean where one wishes to maximize the influence. We have chosen to maximize the influence at observations with 2.75 standard deviations away from our posterior mean. In the first picture, $\beta_p = 0$ and thus the influence function corresponds to the Kullback-Leibler Divergence. Concordantly, it has no maximum and observations have more influence the further in the tail of our model they occur. Thus, one needs to increase β_p slightly. This is done in the second picture. While observations in the tail get smaller influence now than before, the influence of observations is still increasing beyond 2.75 standard deviations. So one needs to increase β_p two more times, until one finally obtains the desired outcome for $\beta_p = 0.25$ in the fourth picture. Notice that the influence does not immediately drop to 0 for observations further in the tails than 2.75 standard deviations away from the posterior mean, but it does decay. In some sense, we have set β_p such that we think of observations occurring 2.75 standard deviations away from the posterior mean as being *most informative*. This is significantly different from what is implied by the Kullback-Leibler divergence, where an observation is most informative if it agrees least with the fitted model. It is intuitive why this produces good inferences if one is in the M-closed world and similarly intuitive why it does not in the M-open world.

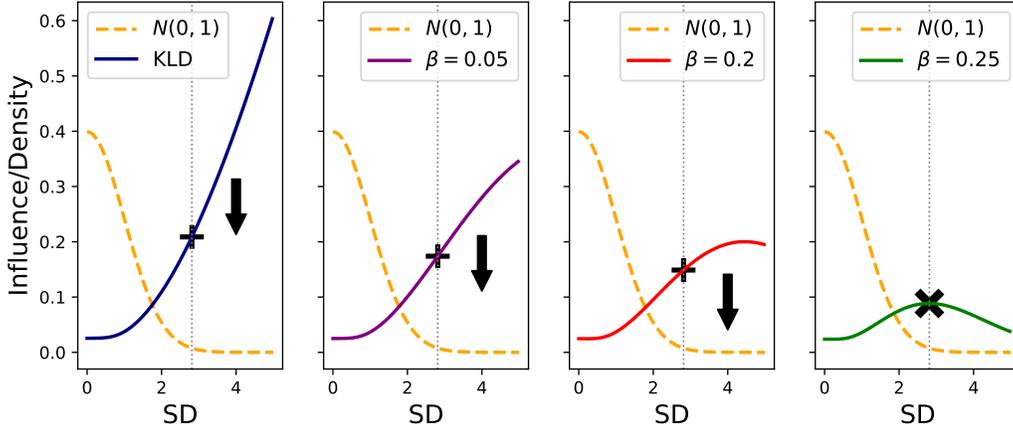


Figure 4: Illustration of the initialization procedure, from left to right.

6 Recursive Optimization for β_{rlm}

Recall that

$$\hat{\mathbf{y}}_t(\boldsymbol{\beta}) = \sum_{r_t, m_t} \mathbb{E}(\mathbf{y}_t | \mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}, \beta_p) p(r_{t-1}, m_{t-1} | \mathbf{y}_{1:(t-1)}, \beta_{rlm}). \quad (97)$$

the issue reduces to finding the partial derivatives $\nabla_{\beta_{rlm}} \hat{\mathbf{y}}_t(\boldsymbol{\beta})$ and $\nabla_{\beta_p} \hat{\mathbf{y}}_t(\boldsymbol{\beta})$. Notice that for $\nabla_{\beta_{rlm}} \hat{\mathbf{y}}_t(\boldsymbol{\beta})$, one finds that

$$\nabla_{\beta_{rlm}} \hat{\mathbf{y}}_t(\boldsymbol{\beta}) = \sum_{r_t, m_t} \mathbb{E}(\mathbf{y}_t | \mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}, \beta_p) \nabla_{\beta_{rlm}} p(r_{t-1}, m_{t-1} | \mathbf{y}_{1:(t-1)}, \beta_{rlm}). \quad (98)$$

Observe now that for $p(\mathbf{y}_{1:t}) = \sum_{r_t, m_t} p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{rlm})$,

$$\begin{aligned} & \nabla_{\beta_{rlm}} p(r_t, m_t | \mathbf{y}_{1:t}, \beta_{rlm}) \\ &= \nabla_{\beta_{rlm}} \left\{ \frac{p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{rlm})}{\sum_{r_t, m_t} p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{rlm})} \right\} \\ &= \frac{\nabla_{\beta_{rlm}} p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{rlm})}{p(\mathbf{y}_{1:t})} - \frac{p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{rlm})}{p(\mathbf{y}_{1:t})^2} \cdot \sum_{r_t, m_t} \nabla_{\beta_{rlm}} p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{rlm}). \end{aligned} \quad (99)$$

Thus we have reduced the problem to finding $\nabla_{\beta_{\text{rlm}}} p(r_t, m_t, \mathbf{y}_{1:t} | \beta_{\text{rlm}})$. Defining for a predictive posterior distribution $f_{m_t}(\mathbf{y}_t | \mathcal{F}_{t-1})$ its β -divergence analogue as

$$f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) = \exp \left\{ \frac{1}{\beta_{\text{rlm}}} f_{m_t}(\mathbf{y}_t | \mathcal{F}_{t-1})^{\beta_{\text{rlm}}} - \frac{1}{1 + \beta_{\text{rlm}}} \int_{\mathcal{Y}} f_{m_t}(\mathbf{y}_t | \mathcal{F}_{t-1})^{1 + \beta_{\text{rlm}}} d\mathbf{y}_t \right\} \quad (100)$$

and uppressing the conditioning on β_{rlm} for convenience, one can using the recursion

$$p(\mathbf{y}_{1:t}, r_t, m_t) = \sum_{m_{t-1}, r_{t-1}} \left\{ f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) q(m_t | \mathcal{F}_{t-1}, m_{t-1}) H(r_t, r_{t-1}) p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}) \right\} \quad (101a)$$

compute $\nabla_{\beta_{\text{rlm}}} p(r_t, m_t, \mathbf{y}_{1:t})$ from $\nabla_{\beta_{\text{rlm}}} p(r_{t-1}, m_{t-1}, \mathbf{y}_{1:(t-1)} | \beta_{\text{rlm}})$ for $r_t = r_{t-1} + 1$ as

$$\begin{aligned} & \nabla_{\beta_{\text{rlm}}} p(\mathbf{y}_{1:t}, r_t, m_t) \\ &= \left\{ \nabla_{\beta_{\text{rlm}}} f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) q(m_t | \mathcal{F}_{t-1}, m_{t-1}) H(r_t, r_{t-1}) p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}) \right\} + \\ & \left\{ f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) \nabla_{\beta_{\text{rlm}}} q(m_t | \mathcal{F}_{t-1}, m_{t-1}) H(r_t, r_{t-1}) p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}) \right\} + \\ & \left\{ f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) q(m_t | \mathcal{F}_{t-1}, m_{t-1}) H(r_t, r_{t-1}) \nabla_{\beta_{\text{rlm}}} p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}) \right\}. \quad (102) \end{aligned}$$

Similarly, for $r_t = 0$ the expression becomes

$$\begin{aligned} & \nabla_{\beta_{\text{rlm}}} p(\mathbf{y}_{1:t}, r_t, m_t) \\ &= \nabla_{\beta_{\text{rlm}}} f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) \cdot q(m_t) \sum_{r_{t-1}, m_{t-1}} H(0, r_{t-1}) p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}) + \\ & f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) \cdot q(m_t) \sum_{r_{t-1}, m_{t-1}} H(0, r_{t-1}) \nabla_{\beta_{\text{rlm}}} p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}). \quad (103) \end{aligned}$$

This implies that if $f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1})$ and $q(m_t | \mathcal{F}_{t-1}, m_{t-1})$ are differentiable with respect to β_{rlm} , then the entire expression can be updated recursively. For most exponential family likelihoods (and in particular the normal likelihood of the Bayesian Linear Regression), $\nabla_{\beta_{\text{rlm}}} f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1})$ is available analytically. In particular, as long as $\int_{\mathcal{Y}} f_{m_t}(\mathbf{y}_t | \mathcal{F}_{t-1})^{1 + \beta_{\text{rlm}}} d\mathbf{y}_t$ has a closed form, $\nabla_{\beta_{\text{rlm}}} f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1})$ can be found in analytic form. In the case of Bayesian Linear Regression where the d -dimensional posterior predictive takes the shape of a student- t distribution with ν degrees of freedom and posterior covariance $\frac{\nu}{\nu-2} \Sigma$, one finds that

$$\begin{aligned} \nabla_{\beta_{\text{rlm}}} f_{m_t}^{\beta_{\text{rlm}}}(\mathbf{y}_t | \mathcal{F}_{t-1}) &= \nabla_{\beta_{\text{rlm}}} g_1(\beta_{\text{rlm}}) g_2(\beta_{\text{rlm}}) g_3(\beta_{\text{rlm}}) + \\ & g_1(\beta_{\text{rlm}}) \nabla_{\beta_{\text{rlm}}} g_2(\beta_{\text{rlm}}) g_3(\beta_{\text{rlm}}) + \\ & g_1(\beta_{\text{rlm}}) g_2(\beta_{\text{rlm}}) \nabla_{\beta_{\text{rlm}}} g_3(\beta_{\text{rlm}}), \quad (104) \end{aligned}$$

where for $\eta = \nu d + d\beta_{\text{rlm}} + \nu$,

$$\begin{aligned} g_1(\beta_{\text{rlm}}) &= \left(\frac{\Gamma(0.5[\nu + d])}{\Gamma(0.5\nu)} \right)^{1 + \beta_{\text{rlm}}} \\ g_2(\beta_{\text{rlm}}) &= \frac{\Gamma(0.5\eta)}{\Gamma(0.5[\eta + p])} \\ g_3(\beta_{\text{rlm}}) &= (\nu\pi)^{-0.5p\beta_{\text{rlm}}} \cdot |\Sigma|^{-\beta_{\text{rlm}}}, \end{aligned}$$

so that their derivatives are given by

$$\begin{aligned} \nabla_{\beta_{\text{rlm}}} g_1(\beta_{\text{rlm}}) &= -(\beta_{\text{rlm}} + 1) \cdot \log(g_1(\beta_{\text{rlm}})) \cdot g_2(\beta_{\text{rlm}}) \\ \nabla_{\beta_{\text{rlm}}} g_2(\beta_{\text{rlm}}) &= 0.5(\nu + p) \left[\frac{\Gamma(0.5\eta) \Psi(0.5\eta)}{\Gamma([0.5[\eta + p])} - \frac{\Gamma(0.5[\eta]) \Psi(0.5[p + \eta])}{\Gamma([0.5[\eta + p])} \right] \\ \nabla_{\beta_{\text{rlm}}} g_3(\beta_{\text{rlm}}) &= -g_3(\beta_{\text{rlm}}) \cdot \log(g_3(\beta_{\text{rlm}})) \cdot \frac{1}{\beta_{\text{rlm}}} \end{aligned} \quad (105)$$

As for $\nabla_{\beta_{\text{rlm}}} q(m_t | \mathcal{F}_{t-1}, m_{t-1})$, one can again obtain it recursively, since for $r_t > 0$,

$$\begin{aligned} & \nabla_{\beta_{\text{rlm}}} q(m_t | \mathcal{F}_{t-1}, m_{t-1}) \\ &= \nabla_{\beta_{\text{rlm}}} \left\{ \frac{p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1})}{\sum_{m_{t-1}} p(\mathbf{y}_{1:(t-1)}, m_{t-1})} \right\} \\ &= \frac{\nabla_{\beta_{\text{rlm}}} p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1})}{\sum_{m_{t-1}} p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1})} - \frac{\sum_{m_{t-1}} \nabla_{\beta_{\text{rlm}}} p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1})}{\left(\sum_{m_{t-1}} p(\mathbf{y}_{1:(t-1)}, r_{t-1}, m_{t-1}) \right)^2}. \quad (106) \end{aligned}$$

7 Proof of Theorem 2

Proof. For ease of notation, we use $\beta_p = \beta$. The model used for the inference is an exponential family model of the form

$$f(x; \theta) = \exp(\eta(\theta)^T T(x)) g(\eta(\theta)) A(x), \quad (107)$$

where $g(\eta(\theta)) := (\int \exp(\eta(\theta)^T T(x)) A(x) dx)^{-1}$. Now under our SVI routine the β -D posterior originating from this model and its conjugate prior is approximated by a member of the conjugate prior family. As a result the conjugate prior and variational posterior to the above model have the form

$$\pi_0(\theta | \nu_0, \mathcal{X}_0) = g(\eta(\theta))^{\nu_0} \exp(\nu_0 \eta(\theta)^T \mathcal{X}_0) h(\mathcal{X}_0, \nu_0) \quad (108)$$

$$\pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n) = g(\eta(\theta))^{\nu_n} \exp(\nu_n \eta(\theta)^T \mathcal{X}_n) h(\mathcal{X}_n, \nu_n), \quad (109)$$

where (ν_0, \mathcal{X}_0) are the prior hyperparameters, (ν_n, \mathcal{X}_n) represent the variational parameters and $h(\mathcal{X}_i, \nu_i) := (\int g(\eta(\theta))^{\nu_i} \exp(\nu_i \eta(\theta)^T \mathcal{X}_i) d\theta)^{-1}$. The resulting ELBO objective function under GBI has the form

$$\begin{aligned} ELBO(\nu_n, \mathcal{X}_n) = & \mathbb{E}_{\pi_n^{VB}} \left[\log \left(\exp \left(\sum_{i=1}^n -\ell^D(x; \theta) \right) \right) \right] - d_{KL}(\pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n), \pi_0(\theta | \nu_0, \mathcal{X}_0)) \end{aligned} \quad (110)$$

where for the β -D posterior

$$\begin{aligned} -\ell^\beta(x; \theta) = & \frac{1}{\beta} (\exp(\eta(\theta)^T T(x)) g(\eta(\theta)) A(x))^\beta - \\ & \frac{1}{\beta + 1} \int (\exp(\eta(\theta)^T T(z)) g(\eta(\theta)) A(z))^{1+\beta} dz \end{aligned} \quad (111)$$

$$\begin{aligned} = & \frac{1}{\beta} \exp(\beta \eta(\theta)^T T(x)) g(\eta(\theta))^\beta A(x)^\beta - \\ & \frac{1}{\beta + 1} \int \exp((1 + \beta) \eta(\theta)^T T(z)) g(\eta(\theta))^{1+\beta} A(z)^{1+\beta} dz. \end{aligned} \quad (112)$$

Therefore the $ELBO(\nu_n, \mathcal{X}_n)$ has three integrals that need evaluating

$$B_1 = \sum_{i=1}^n \int \frac{1}{\beta} \exp(\beta \eta(\theta)^T T(x_i)) g(\eta(\theta))^\beta A(x_i)^\beta \pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n) d\theta \quad (113)$$

$$B_2 = \frac{n}{\beta + 1} \int \left\{ \int \exp((1 + \beta) \eta(\theta)^T T(z)) g(\eta(\theta))^{1+\beta} A(z)^{1+\beta} dz \right\} \pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n) d\theta \quad (114)$$

$$B_3 = d_{KL}(\pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n), \pi_0(\theta | \nu_0, \mathcal{X}_0)). \quad (115)$$

Now firstly for the term B_1 in equation (113)

$$B_1 = \sum_{i=1}^n \int \frac{1}{\beta} \exp(\beta \eta(\theta)^T T(x_i)) g(\eta(\theta))^\beta A(x_i)^\beta g(\eta(\theta))^{\nu_n} \exp(\nu_n \eta(\theta)^T \mathcal{X}_n) h(\mathcal{X}_n, \nu_n) d\theta \quad (116)$$

$$= \sum_{i=1}^n \frac{1}{\beta} A(x_i)^\beta h(\mathcal{X}_n, \nu_n) \int g(\eta(\theta))^{\beta + \nu_n} \exp(\eta(\theta)^T (\beta T(x_i) + \nu_n \mathcal{X}_n)) d\theta \quad (117)$$

$$= \sum_{i=1}^n \frac{1}{\beta} A(x_i)^\beta h(\mathcal{X}_n, \nu_n) \frac{1}{h(\frac{\beta T(x_i) + \nu_n \mathcal{X}_n}{\beta + \nu_n}, \beta + \nu_n)}. \quad (118)$$

Where we know that $h(\frac{\beta T(x_i) + \nu_n \mathcal{X}_n}{\beta + \nu_n}, \beta + \nu_n) = \int g(\eta(\theta))^{\beta + \nu_n} \exp(\eta(\theta)^T (\beta T(x_i) + \nu_n \mathcal{X}_n)) d\theta$ is integrable and closed form as it represents the normalising constant of the same exponential family as the prior and the variational posterior. Next we look at B_2 in equation (114). The whole integral is the product of two densities which must be positive and in order for the $ELBO(\nu_n, \mathcal{X}_n)$ to be defined it must also be integrable. Therefore we can use Fubini's theorem to switch the order of integration

$$\begin{aligned} B_2 &= \frac{n}{\beta + 1} \int \left\{ \int \exp((1 + \beta)\eta(\theta)^T T(z)) g(\eta(\theta))^{1 + \beta} \pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n) d\theta \right\} A(z)^{1 + \beta} dz \quad (119) \\ &= \frac{n}{\beta + 1} h(\mathcal{X}_n, \nu_n) \int \left\{ \int \exp(\eta(\theta)^T ((1 + \beta)T(z) + \nu_n \mathcal{X}_n)) g(\eta(\theta))^{1 + \beta + \nu_n} d\theta \right\} A(z)^{1 + \beta} dz \quad (120) \\ &= \frac{n}{\beta + 1} h(\mathcal{X}_n, \nu_n) \int \frac{A(z)^{1 + \beta}}{h(\frac{(1 + \beta)T(z) + \nu_n \mathcal{X}_n}{1 + \beta + \nu_n}, 1 + \beta + \nu_n)} dz. \quad (121) \end{aligned}$$

once again $h(\frac{(1 + \beta)T(z) + \nu_n \mathcal{X}_n}{1 + \beta + \nu_n}, 1 + \beta + \nu_n) = \int \exp(\eta(\theta)^T ((1 + \beta)T(z) + \nu_n \mathcal{X}_n)) g(\eta(\theta))^{1 + \beta + \nu_n} d\theta$ is the normalising constant of the same exponential family as the prior and the variational posterior and is thus closed form. Lastly we look at B_3 in equation (115)

$$B_3 = \int \pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n) \log \frac{g(\eta(\theta))^{\nu_n} \exp(\nu_n \eta(\theta)^T \mathcal{X}_n) h(\mathcal{X}_n, \nu_n)}{g(\eta(\theta))^{\nu_0} \exp(\nu_0 \eta(\theta)^T \mathcal{X}_0) h(\mathcal{X}_0, \nu_0)} \quad (122)$$

$$= \log \frac{h(\mathcal{X}_n, \nu_n)}{h(\mathcal{X}_0, \nu_0)} \int \pi_n^{VB}(\theta | \nu_n, \mathcal{X}_n) \{(\nu_n - \nu_0) \log g(\eta(\theta)) + (\eta(\theta)^T (\nu_n \mathcal{X}_n - \nu_0 \mathcal{X}_0))\} \quad (123)$$

$$= \log \frac{h(\mathcal{X}_n, \nu_n)}{h(\mathcal{X}_0, \nu_0)} \{(\nu_n - \nu_0) \lambda_n^{VB} + ((\mu_n^{VB})^T (\nu_n \mathcal{X}_n - \nu_0 \mathcal{X}_0))\}, \quad (124)$$

where $\mu_n^{VB} = \mathbb{E}_{\pi_n^{VB}}[\eta(\theta)]$ and $\lambda_n^{VB} = \mathbb{E}_{\pi_n^{VB}}[\log g(\eta(\theta))]$.

As a result we get that

$$ELBO(\nu_n, \mathcal{X}_n) = B_1 - B_2 - B_3 \quad (125)$$

$$\begin{aligned} &= \sum_{i=1}^n \frac{1}{\beta} A(x_i)^\beta h(\mathcal{X}_n, \nu_n) \frac{1}{h(\frac{\beta T(x_i) + \nu_n \mathcal{X}_n}{\beta + \nu_n}, \beta + \nu_n)} \\ &\quad - \frac{n}{\beta + 1} h(\mathcal{X}_n, \nu_n) \int \frac{A(z)^{1 + \beta}}{h(\frac{(1 + \beta)T(z) + \nu_n \mathcal{X}_n}{1 + \beta + \nu_n}, 1 + \beta + \nu_n)} dz \quad (126) \\ &\quad - \log \frac{h(\mathcal{X}_n, \nu_n)}{h(\mathcal{X}_0, \nu_0)} \{(\nu_n - \nu_0) \lambda_n^{VB} + ((\mu_n^{VB})^T (\nu_n \mathcal{X}_n - \nu_0 \mathcal{X}_0))\}. \end{aligned}$$

□

8 Complexity Analysis of Inference

Time complexity: Our SVRG method crucially hinges on the complexity of the gradient evaluations. For BLR, we note that evaluating the complete ELBO gradient derived above for n observations has complexity $\mathcal{O}(np^3)$, where p is the number of regressors. We proceed by defining g as the (generic) complexity of a gradient evaluation, so for BLR $g = p^3$. Clearly, an SGD step using b observations is of order $\mathcal{O}(bg)$. Similarly, the computation of the anchors is $\mathcal{O}(Bg)$. Next, let the optimization routine used for full optimization have complexity $\mathcal{O}(m(n, \dim(\theta)))$. Most standard (quasi-) Newton optimization routines such as BFGS or LBFGSB (used in our implementation) are polynomial in n and $\dim(\theta)$. For such methods, since it holds that at most $W \geq n$ observations are evaluated in the full optimization, and since $\dim(\theta)$ is time-constant, $m(n, \dim(\theta))$ is also constant in time. Thus, though these constants can be substantial, all optimization steps (whether SVRG steps

or full optimization steps) are $\mathcal{O}(1)$ in time. Since one performs T of them for T observations, the computational complexity (in time) is $\mathcal{O}(T)$.

Space complexity: One needs to store observations \mathbf{y}_t as well as gradient evaluations. Storing one of them takes $\mathcal{O}(d)$ and $\mathcal{O}(\dim(\boldsymbol{\theta}))$ space, respectively. Since we only keep a window W of the most recent observations (and gradients), this means that the space requirement is of order $\mathcal{O}(W(d + \dim(\boldsymbol{\theta})))$ and in particular constant in time.

9 Additional Details on Experiments

For all experiment, constrained Limited Memory Broyden–Fletcher–Goldfarb–Shannon is used for the full optimization step, where the constraints are $\hat{a}_n > 1, \hat{b}_n > 1$. We use Python’s `scipy.optimize` wrapper, which calls a Fortran implementation. We also tested whether inference is sensitive to different initializations of β_p and found that it is fairly stable as long as β_p is chosen reasonably. For example, for the Air Pollution data, we could recover the same changepoint (± 5 days) for initializations of β_p ranging from 0.005 up to 0.1. All experiments were performed on a 2017 MacBook Pro with 16 GB 2133 MHz LPDDR3 and 3.1 GHz Intel Core i7.

9.1 Well-log data

Hyperparameters: We set the hyperparameters for standard Bayesian On-line Changepoint Detection slightly differently, the reason being that due to the robustness guarantee of Theorem 1, we can use much less informative priors with the robust version than we can with the standard version: If priors are too flat, the standard version declares far too many changepoints. Thus, for the standard version, we use a constant CP prior (hazard) $H(r_t = r_{t-1} + 1 | r_{t-1}) = 0.01, a_0 = 1, b_0 = 10^4, \Sigma_0 = 0.25, \boldsymbol{\mu}_0 = 1.15 \cdot 10^4$, while for the robust version we can use a less informative prior by instead setting $b_0 = 10^7$. By virtue of our initialization procedure for β_p , this implies setting $\beta_{p,0} \approx 0.05$. To start out close to the KLD, we initialize $\beta_{\text{rid},0} = 0.0001$.

Inferential procedure: For the robust version, we set $W = 360, B = 25, b = 10, m = 20, K = 1$. For both versions, only the 50 most likely run-lengths are kept. For the robust version, the average processing time was 0.487 per observation.

9.2 Air Pollution data

Preprocessing & Model Setup: The air pollution data is observed every 15 minutes across 29 stations for 365 days. We average the 96 observations made over 24 hours. This is done to move the observed data closer to a normal distribution, as the measurements have significant daily volatility variations. To account for weekly cycles, we also calculate for each station the mean for each weekday and subtract it from the raw data.. Yearly seasonality is not accounted for. Afterwards, the data is normalized station-wise. This is done only for numerical stability, because the internal mechanisms of the used VAR models perform matrix operations (QR-decompositions and matrix multiplications in particular) that can adversely affect numerical stability for observations with large absolute value. Fig. 5 shows some of the station’s data after these preprocessing steps have been taken.

The autoregressive models and spatially structured vector autoregressive models (VARs) are chosen to have lag lengths 1, 2, 3. These short lag lengths are chosen to explicitly disadvantage the robust model universe: The non-robust run we compare against uses more than 20 models, with lag lengths 1, 5, 6, 7, meaning that it is much more expressive and should be able to cope with outliers better. In spite of this, it not only declares more CPs, but also does worse than the robust version in terms of predictive performance. For both the robust and non-robust model, two spatially structured VARs are included as in Knoblauch and Damoulas [2018].

Hyperparameters: We set $H(r_t = r_{t-1} + 1 | r_{t-1}) = 0.001, a_0 = 1, b_0 = 25, \boldsymbol{\mu}_0 = \mathbf{0}, \Sigma_0 = I \cdot 20$, which yields initialization $\beta_p \approx 0.005, \beta_{\text{rlm}} = 0.1$. The non-robust results are directly taken from Knoblauch and Damoulas [2018] and can be replicated running the code available from <https://github.com/alan-turing-institute/bocpdms/>

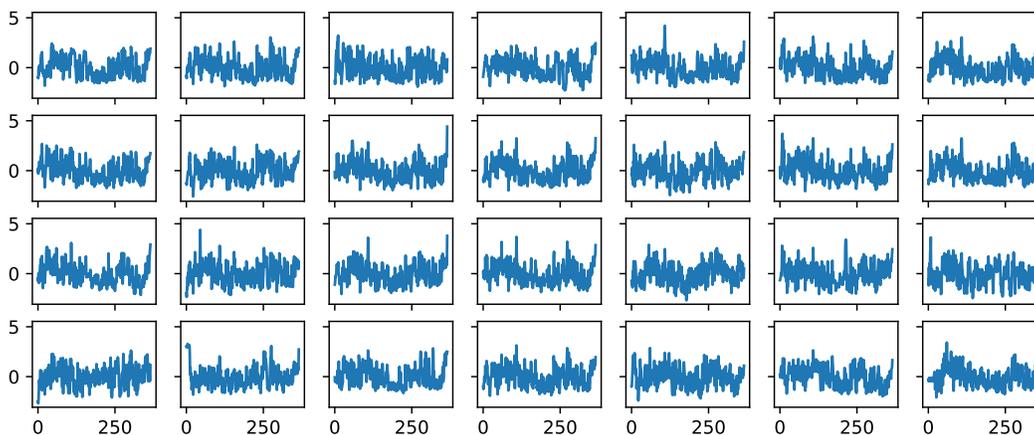


Figure 5: Some of the stations after preprocessing steps. x -axis gives NOx level, y -axis the day.

Inferential procedure: We set $W = 300$, $m = 50$, $B = 20$ and $b = 10$, $K = 25$ and retain the 50 most likely run-lengths. Processing times are more volatile than for the well-log because the full optimization procedure is significantly more expensive to perform. Most observations take significantly less than 20 seconds to process, but some take over a minute (depending on how many of the retained run-lengths are divisible by m at each time point).

9.3 Optimizing β

Lastly, we investigate the trajectories for β as it is being optimized. For all trajectories, a bounded predictive absolute loss was used with threshold τ , i.e. $L(x) = \max\{|x|, \tau\}$. For β_{rid} , $\tau = 5/T$ (where T is the length of the time series) while for β_{p} , $\tau = 0.1$. The results are not sensitive to these thresholds, and they are picked with the intent that (1) a single observation should not affect β_{p} by more than 0.1 and (2) that overall, β_{rid} should not change by more than 5 in absolute magnitude. As the initialization procedure for β_{p} works very well for predictive performance, the on-line optimization never even comes close to making a step with size τ . The picture is rather different for β_{rid} , which reaches τ rather often. We note that this is because the estimated gradients for β_{rid} can be very extreme, which is why the implementation averages 50 consecutive gradients before performing a step. Overall, we note that for the well log data whose trajectories are depicted in Fig. 6, the degrees of robustness do not change much relative to their starting points at $\beta_{\text{p}} = 0.05$ and $\beta_{\text{rid}} = 0.001$. In particular, the absolute change over more than 4,000 observations is < 0.002 for β_{p} and < 0.015 for β_{rid} . Step sizes are $1/t$ at time t .

For the Air Pollution Data, the story is slightly different: Here, β_{p} does not change after the first iteration, where it jumps from 0.005 directly to 10^{-10} . While this seems odd, it is mainly due to the fact that for numerical stability reasons², one needs to ensure that $\beta_{\text{p}} > \varepsilon$ for some $\varepsilon > 0$; and in our implementation, $\varepsilon = 10^{-10}$. The interpretation of the trace graph is thus that the optimization continuously suggests less robust values for β_{p} , but that we cannot admit them due to numerical stability. The downward trend also holds for β_{rid} , which is big enough to not endanger numerical stability and hence can drift downwards.

Fig. 6 also shows that the optimization technique used for β needs further investigation and research. For starters, the outcomes suggest that a second order method could yield better results than using a first-order SGD technique. In the future, we would like to explore this in greater detail and also

²In particular, working with the β -D implies that one takes the exponential of a density, i.e. e^{f^β} . So even working on a log scale now means working with the densities f^β directly. It should be clear that these quantities become numerically unstable for β too large or too small.

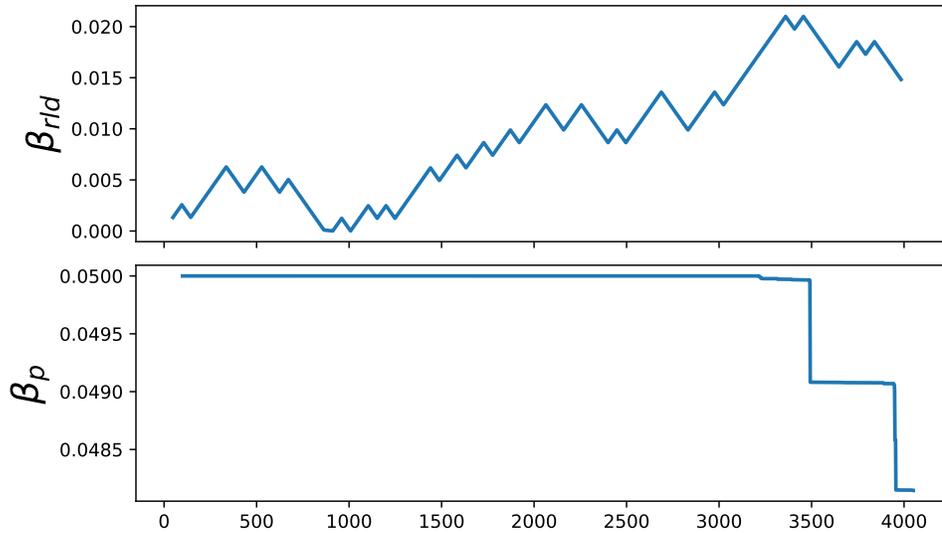


Figure 6: β trajectories for the well-log data. For β_{rld} , steps are only taken every 50 observations to average gradient noise

explore more advanced optimization methods like line search or trust region optimization methods for this problem.

References

- Bob Carpenter, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 20, 2016.
- Jeremias Knoblauch and Theodoros Damoulas. Spatio-temporal Bayesian on-line changepoint detection with model selection. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, 2018.