

## Supplementary Material

### Bayesian linear regression

Bayesian linear regression is an approach to ordinary linear regression (OLR) within a Bayesian framework. The OLR model is typically denoted as:

$$y = \beta_0 + \beta_1 x + \varepsilon \quad (\text{A1})$$

where  $\mathbf{x}$  are the explanatory variables,  $\mathbf{y}$  the dependent variables,  $\beta_0$  and  $\beta_1$  the regression coefficients, and  $\varepsilon$  are the regression error, assumed to be independent and identically distributed, expressed as  $\sim N(0, \sigma^2)$ . Another assumption of OLR is that the  $\mathbf{x}$  are independent of  $\varepsilon$ . The objective of OLR is to calculate values of  $\beta_0$  and  $\beta_1$  that minimize (in the least-square sense) the error between the observed values of  $\mathbf{y}$  and  $\beta_0$  and  $\beta_1 \mathbf{x}$ . In a Bayesian framework, this model structure can also be represented as a Normal probability model:

$$y \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I}) \quad (\text{A2})$$

where  $\mathbf{X}$  is matrix of dependent variables, limited to  $[\mathbf{x}_1 \ \mathbf{x}_2]$  here with  $\mathbf{x}_1 = \mathbf{I}$  to allow for an intercept, and  $\mathbf{x}_2$  are the observations, the  $\boldsymbol{\beta}$  are the corresponding regression coefficients,  $\sigma^2$  is the variance of the dependent variable  $\mathbf{y}$  given the observations (*i.e.*  $\mathbf{y}|\boldsymbol{\beta}, \mathbf{X}$ ), and lastly,  $\mathbf{I}$  is the identity matrix. In a Bayesian framework, this function ( $\mathbf{y}|\boldsymbol{\beta}, \sigma^2, \mathbf{X}$ ) is the likelihood function. Note that Bayes rule states that the posterior distribution is proportional to the product of the likelihood function and the prior distribution. Thus, OLR when viewed through a Bayesian framework can be written as:

$$p(\boldsymbol{\beta}, \sigma^2 | \mathbf{y}, \mathbf{X}) \propto p(\mathbf{y} | \boldsymbol{\beta}, \sigma^2, \mathbf{X}) p(\boldsymbol{\beta}, \sigma^2) \quad (\text{A3})$$

where the term on the left hand side represents the posterior, and the right hand side is the product of the likelihood and the prior. The likelihood in this case is a multivariate Normal, and is often more convenient to be parameterized using the following terms (the derivation is not shown here but available in many econometric textbooks including Zellner, 1971; Birkes and Dodge, 1993; Koop, 2003; Greenberg, 2008; Gelman et al., 2014):

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y} \quad (\text{A4})$$

$$s^2 = (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})^T (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}) / \nu \quad (\text{A5})$$

These values are identical to the least-squares estimates of degrees of freedom ( $\nu = N - 2$ ) (with  $N$  the total number of observations),  $\boldsymbol{\beta}$  and the standard error,  $s^2$ .

#### Natural conjugate priors

In Bayesian analysis, the prior is meant to reflect any previously available data or information about the parameters of interest. The choice of prior used determines whether or not there is an analytical solution for the posterior distribution. In the case where no analytical solution is possible, numerical integration techniques, using Monte Carlo simulations, such as the Gibb's sampler are needed. A particular class of priors, known as *natural conjugate priors*, are typically used because when they are combined with the likelihood function, these yield a posterior of the same class of distributions (Birkes and Dodge, 1993; Koop, 2003). This ensures that an analytical solution for the posterior is possible. For the Normal regression model, the natural conjugate prior can be stated as:

$$p(\boldsymbol{\beta}, \sigma^2) = p(\boldsymbol{\beta} | \sigma^2) p(\sigma^2) \quad (\text{A6})$$

or:

$$p(\boldsymbol{\beta}, \tau) = p(\boldsymbol{\beta} | \tau) p(\tau) \quad (\text{A7})$$

where  $\tau = 1/\sigma^2$ , and  $(\boldsymbol{\beta}|\tau) \sim N(\underline{\boldsymbol{\beta}}, \tau^{-1}\underline{\mathbf{V}})$  and  $\tau \sim \Gamma(\underline{s}^{-2}, \underline{\nu})$ , and the underlined terms represent the prior hyper-parameters. This means that the joint prior is Normal-Gamma:  $(\boldsymbol{\beta}, \tau) \sim \text{NG}(\underline{\boldsymbol{\beta}}, \underline{\mathbf{V}}, \underline{s}^{-2}, \underline{\nu})$ . Multiplying this prior with the likelihood function results in the joint posterior of  $\boldsymbol{\beta}$  and  $\tau$  that is also Normal-Gamma (as expected because of the choice of *natural conjugate priors*):

$$(\boldsymbol{\beta}, \tau | \mathbf{y}) \sim \text{NG}(\bar{\boldsymbol{\beta}}, \bar{\mathbf{V}}, \bar{s}^{-2}, \bar{\nu}) \quad (\text{A8})$$

where the bar above parameter denotes the "posterior" value of the hyper-parameters. These hyper-parameters can then be estimated as follows:

$$\bar{\mathbf{V}} = (\underline{\mathbf{V}}^{-1} + \mathbf{X}^T \mathbf{X})^{-1} \quad (\text{A9})$$

$$\bar{\boldsymbol{\beta}} = \bar{\mathbf{V}}(\underline{\mathbf{V}}^{-1} \underline{\boldsymbol{\beta}} + \mathbf{X}^T \mathbf{X} \boldsymbol{\beta}) \quad (\text{A10})$$

$$\bar{\mathbf{v}} = \underline{\mathbf{v}} + N \quad (\text{A11})$$

And  $\bar{s}^{-2}$  can be calculated from:

$$\bar{v}s^{-2} = \underline{v}s^{-2} + \nu s^{-2} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})^T [\underline{\mathbf{V}} + (\mathbf{X}^T \mathbf{X})^{-1}]^{-1} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \quad (\text{A12})$$

In essence, this form of BLR results in model estimates (of  $\boldsymbol{\beta}$  and  $\boldsymbol{\tau}$ ) that are weighted averages of the provided priors and the least-square estimates (West, 1984). This formulation is the standard BLR form with a *natural conjugate prior*.

#### Non-informative priors

It can be readily seen from the above formulation that the impact of the selection of priors can be extremely significant, and can severely influence the posterior values of the model parameters. Thus, often a single value for a prior is rarely used; instead an upper or lower limit is used to get an estimate of parameters. However, in many cases there is no prior information available, and thus this type of prior is unrealistic or subjective. Thus, an objective prior is needed that can represent complete uncertainty regarding the model parameters – this type of prior is known as a *non-informative* prior. To represent this in the formulation above, the values of  $\underline{\mathbf{v}}$  and  $\underline{\mathbf{V}}^{-1}$  are both set at 0. This results in posterior estimates of the hyper-parameters as follows:

$$\bar{\mathbf{V}} = (\mathbf{X}^T \mathbf{X})^{-1} \quad (\text{A13})$$

$$\bar{\boldsymbol{\beta}} = \boldsymbol{\beta} \quad (\text{A14})$$

$$\bar{\mathbf{v}} = N \quad (\text{A15})$$

And  $\bar{s}^{-2}$  can be calculated from:

$$\bar{v}s^{-2} = \nu s^{-2} \quad (\text{A16})$$

Note that when using non-informative priors, the posterior estimates of the hyper-parameters reduce to the familiar least-squares parameter estimates  $\boldsymbol{\beta}$  and  $s^{-2}$ . Using this formulation the predictive distribution of  $\mathbf{y}$  can be found analytically and is multivariate  $t$ :

$$(\mathbf{y}^* | \mathbf{y}) \sim t(\mathbf{X}^* \bar{\boldsymbol{\beta}}, \bar{s}^2 [\mathbf{I} + \mathbf{X}^* \bar{\mathbf{V}} \mathbf{X}^{*T}], \bar{\mathbf{v}}) \quad (\text{A17})$$

where  $\mathbf{X}^*$  are new observation used for prediction, and  $\mathbf{y}^*$  are the predicted values.

#### Independent priors

Apart from the natural conjugate prior, and the non-informative prior described above, another option for the form of prior used in BLR is the independent prior. Under this prior the parameters  $\boldsymbol{\beta}$  and  $\boldsymbol{\tau}$  are independent, *i.e.*  $p(\boldsymbol{\beta}, \boldsymbol{\tau}) = p(\boldsymbol{\beta}) p(\boldsymbol{\tau})$ , which compared to the original relationship between the two parameters removes the conditional on  $\boldsymbol{\beta}$  (*n.b.* the variance of  $\boldsymbol{\beta}$  was listed as  $\boldsymbol{\tau}^{-1} \underline{\mathbf{V}}$ , again showing the dependence of  $\boldsymbol{\beta}$  on  $\boldsymbol{\tau}$ ). By removing this conditional, the independent priors can be expressed as:  $\boldsymbol{\beta} \sim N(\underline{\boldsymbol{\beta}}, \underline{\mathbf{V}})$  and  $\boldsymbol{\tau} \sim \Gamma(\underline{\boldsymbol{\tau}}^{-2}, \underline{\mathbf{v}})$ . Unlike the first case using the natural conjugate priors, multiplying these priors with the likelihood function results in an expression for  $(\boldsymbol{\beta}, \boldsymbol{\tau} | \mathbf{y})$  that is no longer NF, or in the form of a known density function. Thus, an analytical solution for the joint distribution for both posterior parameters and the predictive distribution for  $\mathbf{y}^*$  is not possible. However, a numerical solution using a Monte Carlo integration method along with the conditional posterior distributions of the parameters is possible. The conditional posterior distribution for  $\boldsymbol{\beta}$  and  $\boldsymbol{\tau}$  can be calculated as:

$$(\boldsymbol{\beta} | \boldsymbol{\tau}, \mathbf{y}) \sim N(\bar{\boldsymbol{\beta}}, \bar{\mathbf{V}}) \quad (\text{A18})$$

$$(\boldsymbol{\tau} | \boldsymbol{\beta}, \mathbf{y}) \sim \Gamma(\bar{s}^{-2}, \bar{\mathbf{v}}) \quad (\text{A19})$$

where the hyper-parameters can be calculated as follows:

$$\bar{\mathbf{V}} = (\underline{\mathbf{V}}^{-1} + \boldsymbol{\tau} \mathbf{X}^T \mathbf{X})^{-1} \quad (\text{A20})$$

$$\bar{\beta} = \bar{V}(\underline{V}^{-1} \underline{\beta} + \tau X^T y) \quad (\text{A21})$$

$$\bar{v} = \underline{v} + N \quad (\text{A22})$$

$$\bar{s}^{-2} = \frac{(y - X \bar{\beta})^T (y - X \bar{\beta}) + \underline{v} \bar{s}^{-2}}{\bar{v}} \quad (\text{A23})$$

A Gibb's sampling algorithm can be used to sample from the posterior conditionals to find the joint posterior distribution  $(\beta, \tau)$  numerically (adapted from Greenberg, 2008). First a random value of  $\tau$  is selected; this value is used to calculate one realisation of  $\bar{V}$  then  $\bar{\beta}$ . A value of  $\beta$  is then sampled from  $(\beta | \tau, y) \sim N(\bar{\beta}, \bar{V})$ , and this value is used to sample  $(\tau | \beta, y) \sim \Gamma(\bar{s}^{-2}, \bar{v})$ . This process is repeated until convergence and typically a number of samples considered to be "burn-in" are discarded (selected as 10%, or 100 samples for this research). Each pair of sampled values of  $\beta$  and  $\tau$  represents an approximation of the joint posterior. To predictive probability of  $y$  can be estimate numerically; for each pair of  $\beta$  and  $\tau$ , a sample for  $y^*$  can be drawn from the predictive distribution:

$$(y^*) \sim N(X^* \beta, \frac{1}{\tau}) \quad (\text{A24})$$

where  $X^*$  are new observation used for prediction and  $y^*$  are the predicted values. The posterior conditional distributions from the Gibb's sampling routine can then be used as the priors for any subsequent models.

## Summary of performance metrics for all models

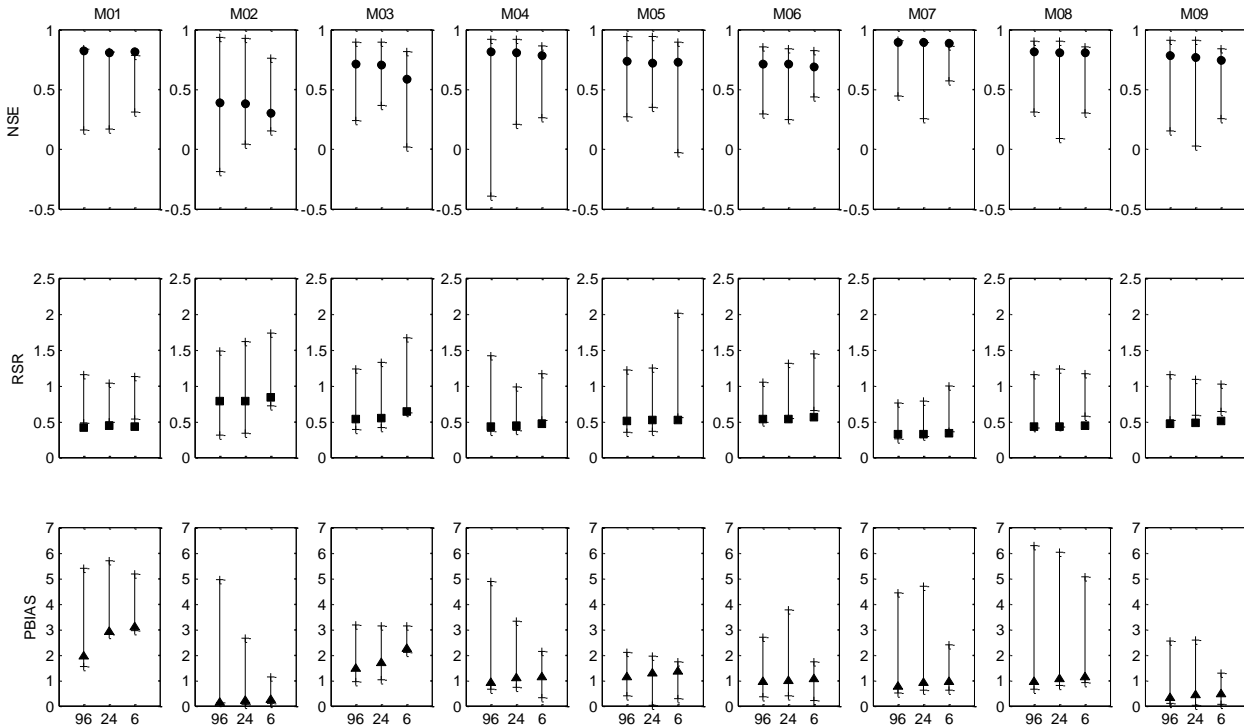


Figure A1. NSE, RSR, and PBIAS results for M01 to M09 using BLR and FLR models for the 1 day lag at all resolutions for the validation dataset. The black markers are the results for the mean Bayesian predictions, whereas the solid black lines are the upper and lower limits of the fuzzy prediction.

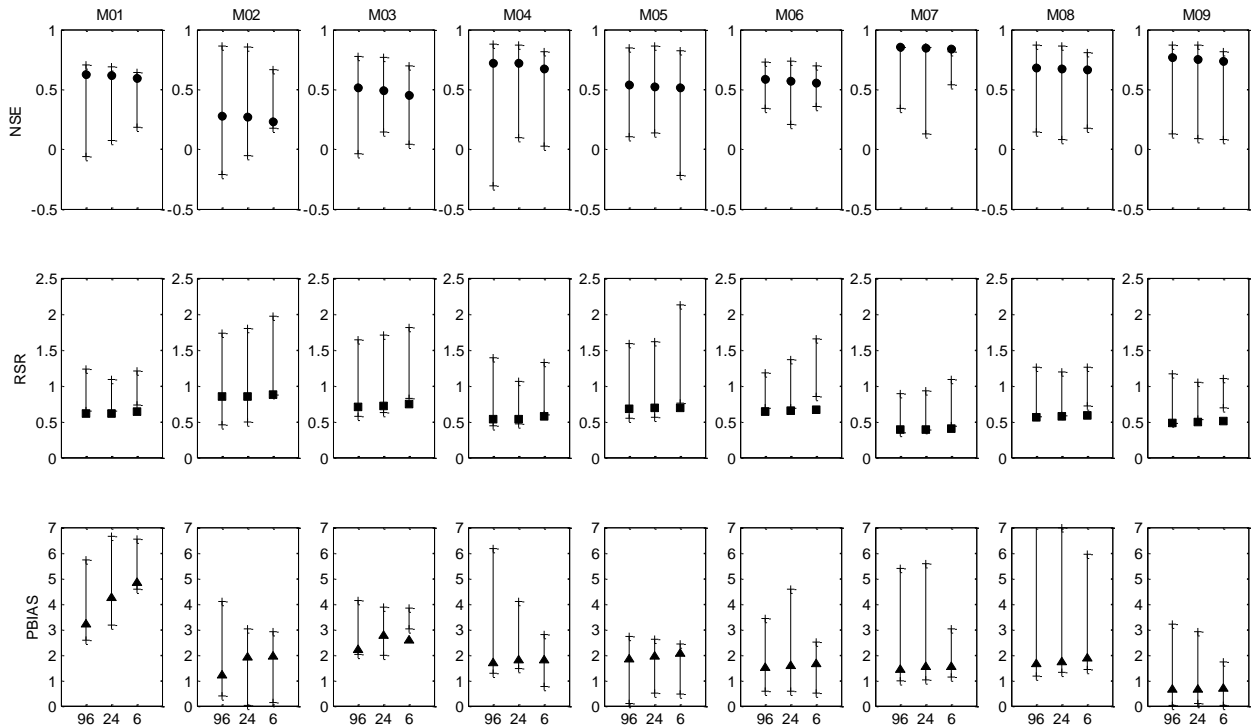


Figure A2. NSE, RSR, and PBIAS results for M01 to M09 using BLR and FLR models for the 2 day lag at all resolutions for the validation dataset. The black markers are the results for the mean Bayesian predictions, whereas the solid black lines are the upper and lower limits of the fuzzy prediction.

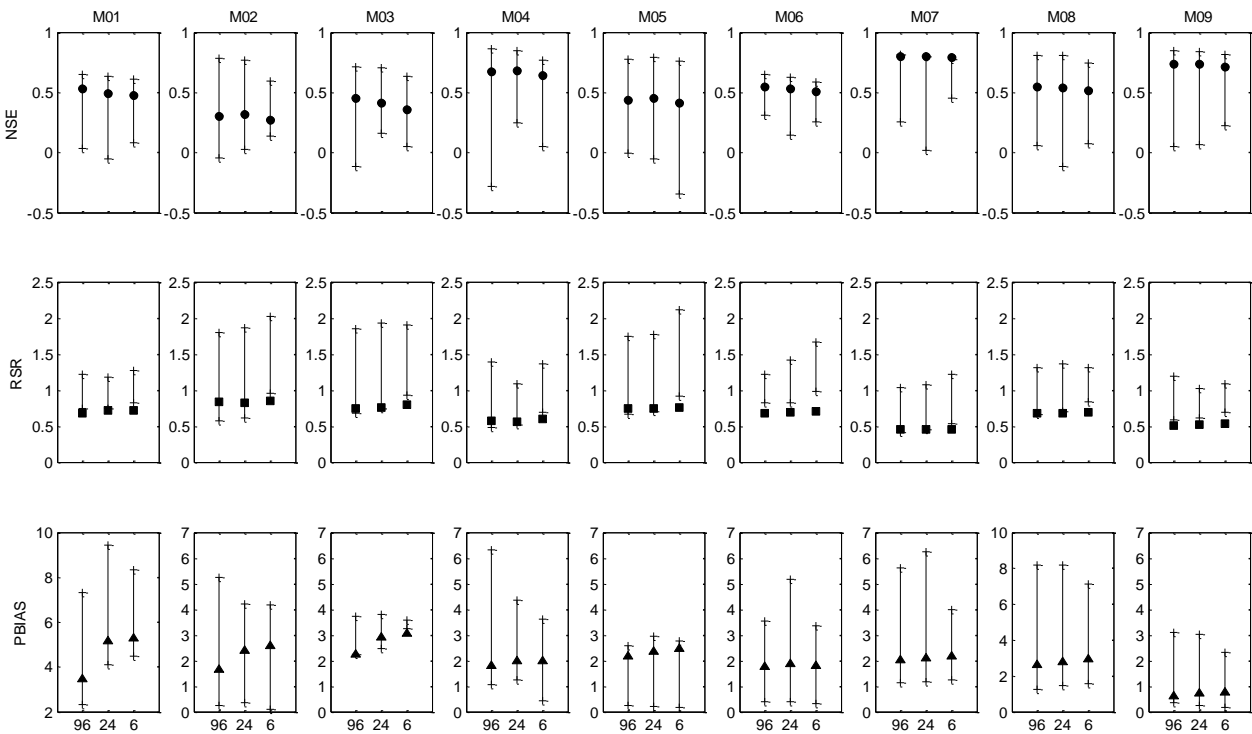


Figure A3. NSE, RSR, and PBIAS results for M01 to M09 using BLR and FLR models for the 3 day lag at all resolutions for the validation dataset. The black markers are the results for the mean Bayesian predictions, whereas the solid black lines are the upper and lower limits of the fuzzy prediction.

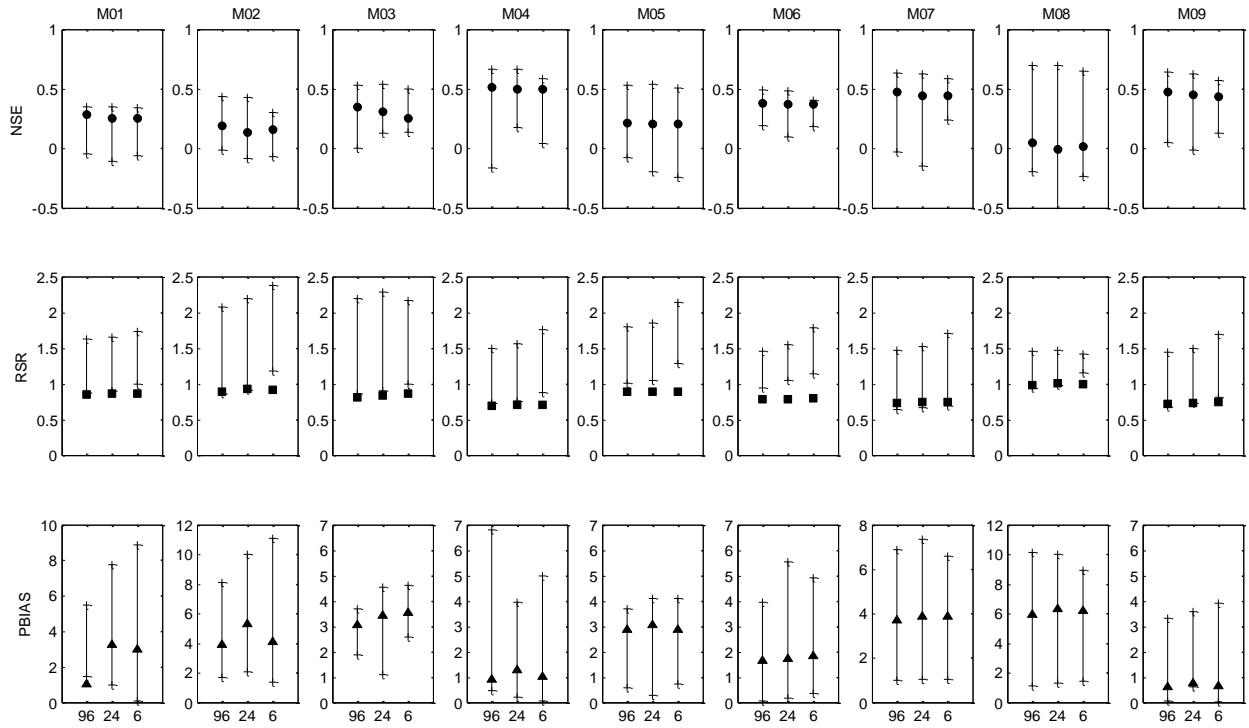


Figure A4. NSE, RSR, and PBIAS results for M01 to M09 using BLR and FLR models for the 7 day lag at all resolutions for the validation dataset. The black markers are the results for the mean Bayesian predictions, whereas the solid black lines are the upper and lower limits of the fuzzy prediction.