

Parameter estimation and inference in dynamic systems described by linear partial differential equations

Gianluca Frasso, Jonathan Jaeger & Philippe Lambert

September 12, 2014

Abstract

Differential equations (DEs) are commonly used to describe dynamic systems evolving in one (ordinary differential equations or ODEs) or in more than one dimensions (partial differential equations or PDEs). In real data applications the parameters involved in the DE models are usually unknown and need to be estimated from the available measurements together with the state function. In this paper, we present frequentist and Bayesian approaches for the joint estimation of the parameters and of the state functions involved in linear PDEs. We also propose two strategies to include state (initial and/or boundary) conditions in the estimation procedure. We evaluate the performances of the proposed strategy through simulated examples and a real data analysis involving (known or necessary) state conditions.

Keywords: Linear partial differential equations; Parameter estimation; Penalized tensor B-spline smoothing; State conditions

1 Introduction

Dynamic systems are commonly described by differential equations (DEs). In real data applications the parameters involved in the DE models are usually unknown and need to be estimated from the available measurements together with the state functions. For one dimensional applications this estimation task has been extensively discussed in the statistical literature. The most popular procedures rely on nonlinear least squares (Biegler et al., 1986). These approaches are computationally intensive and often poorly suited for statistical inference. An attractive alternative is represented by the penalized smoothing framework introduced by Ramsay et al. (2007). This approach can be viewed as a generalization of the L-spline theory (see, for example, Schultz and Varga (1967) among others) where the flexibility acquired by a high dimensional B-spline basis expansion of the state function is counterbalanced by a penalty term expressed using the (set of) ODE(s) that needs to be solved. The fidelity of the final fit to the hypothesized differential model is then tuned using an ODE-compliance parameter. Jaeger and Lambert (2013) adapt the latter approach to a full Bayesian framework. This Bayesian alternative offers two major advantages with respect to the frequentist one: the selection of the ODE-compliance parameter becomes automatic and uncertainty measures about the parameters can readily be obtained using MCMC.

Dealing with dynamics evolving in more than one dimension, partial differential equations (or PDEs) are usually invoked. Inference in models specified by PDEs has not received a large attention in the literature. Here, we discuss statistical approaches to deal with dynamics described by linear PDEs. These kind of equations is used to model large classes of phenomena (e.g. diffusion, heat transfer, price of financial instrument, etc). More precisely, the general form of PDE we consider is given by

$$\mathcal{F}\left(x_1, \dots, x_p, u, \frac{\partial u}{\partial x_1}, \dots, \frac{\partial u}{\partial x_p}, \frac{\partial^2 u}{\partial x_1^2}, \dots, \frac{\partial^2 u}{\partial x_1 \partial x_p}, \dots, \boldsymbol{\theta}\right) = 0, \quad (1)$$

where $u(\mathbf{x})$ is the state function evaluated in $\mathbf{x} = (x_1, \dots, x_p)^\top \in \mathbb{R}^p$ and $\boldsymbol{\theta}$ is the vector of PDE parameters. In this paper, we consider \mathcal{F} as a linear function in u and its derivatives and with coefficients $\boldsymbol{\theta}$. Examples of this kind of PDE are given in Eqs. (13) and (14) of this paper. In particular, we are interested in estimating the vector of unknown $\boldsymbol{\theta}$ using a set of measurements $\boldsymbol{\zeta} = \zeta(\mathbf{x}) = u(\mathbf{x}) + \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon}$ can be considered as a vector of independent normally distributed measurement errors.

Xun et al. (2013) propose both frequentist and Bayesian approaches for the estimation of parameters in model driven by PDEs. The frequentist approach could be viewed as a multidimensional generalization of the proposal made by Ramsay et al. (2007). The state function solving the differential problem is approximated using tensor product B-spline. The flexibility of the approximation is then counterbalanced by a penalty related to the PDE. In the frequentist

proposal, the PDE-penalty is represented by the integral of the PDE operator evaluated at the B-spline approximation of the solution. In their Bayesian approach, they combine a PDE-based penalty with one defined by finite differences (roughness penalty in the spirit of the framework introduced by Eilers and Marx (1996)).

In this paper we present both frequentist and Bayesian frameworks for the joint estimation of the parameters and the state function involved in PDEs. Our methods generalize those proposed in one-dimensional applications by exploiting a PDE-based penalized tensor product B-spline smoothing approach. Contrary to Xun et al. (2013), our Bayesian proposal is directly related to the frequentist one. In Section 2, we introduce our estimation procedures. Both frequentist and Bayesian approaches can be adapted in order to include the state (initial, boundary or mixed) conditions in the estimation process (see Section 3). In Section 4, we evaluate the performance of our approach and compare them with the competitors using simulations. In Section 5 we illustrate our proposals dealing with a real data analysis based on the Black and Scholes (1973) model. Our application represents one of those situations in which the state condition of the differential model (in the specific case the necessary no-arbitrage conditions) are known and need to be satisfied in order to obtain reasonable results. We conclude the paper with a discussion in Section 6.

2 Estimation of the unknown differential parameters for unknown state conditions

In this section, we introduce the PDE-based penalized tensor product B-spline smoothing approach both in frequentist and Bayesian frameworks without taking the state conditions (initial and/or boundary) into account.

2.1 Frequentist approach

Assume that one observes $\{(\zeta; \mathbf{x}) = (\zeta_n; x_{1,n}, \dots, x_{p,n}); n = 1, \dots, N\}$ describing the evolution of the state function u whose dynamics is driven by (1). Our task is to jointly estimate the state function u and the vector of parameters $\boldsymbol{\theta}$. Following Ramsay et al. (2007), we suggest to exploit the DE-based-penalized smoothing approach. More precisely, we denote by $\tilde{u}(\mathbf{x})$ the B-spline tensor product approximation of $u(\mathbf{x})$:

$$\begin{aligned}\tilde{u}(\mathbf{x}) &= (\mathbf{B}_{x_p} \otimes \dots \otimes \mathbf{B}_{x_1}) \mathbf{c} \\ &= \mathbf{B}\mathbf{c},\end{aligned}\tag{2}$$

where \mathbf{c} is the M^p -vector of spline coefficients and \mathbf{B}_{x_p} is the $N \times M$ B-spline matrix defined on the p -th direction taking a generous numbers of internal knots. The placement of the knots on x_p -grid can be either regular or not. To simplify the notation, we suppose that the number of B-spline in each direction is the same. The spline coefficients are shrunk by adding a penalty expressed using the PDE model (1):

$$PEN(\mathbf{c}|\boldsymbol{\theta}) = \int \mathcal{F}(\mathbf{x}, \tilde{u}, \boldsymbol{\theta})^2 d\mathbf{x}.\tag{3}$$

This penalty term will be close to zero, for given PDE parameters $\boldsymbol{\theta}$, if the approximation to the solution of (1) is consistent with the PDE model. As we are considering linear PDEs, this penalty is just a homogeneous polynomial of degree 2 in the spline coefficients:

$$PEN(\mathbf{c}|\boldsymbol{\theta}) = \mathbf{c}^\top \mathbf{R}(\boldsymbol{\theta}) \mathbf{c} + 2\mathbf{c}^\top \mathbf{r}(\boldsymbol{\theta}) + l(\boldsymbol{\theta}),$$

where $\mathbf{R}(\boldsymbol{\theta})$ is the penalty matrix, $\mathbf{r}(\boldsymbol{\theta})$ is the penalty vector and $l(\boldsymbol{\theta})$ is a constant not depending on the spline coefficients. A fast computation of these penalty components is given in Appendix C.

The compromise between data fitting and fidelity to the differential model can be assessed using the penalized least square criterion

$$J(\mathbf{c}|\boldsymbol{\theta}, \gamma, \zeta) = \frac{N}{2} \log(\tau) - \frac{\tau}{2} \|\zeta - \tilde{u}(\mathbf{x})\|^2 - \frac{\gamma}{2} PEN(\mathbf{c}|\boldsymbol{\theta}).\tag{4}$$

Similarly to Ramsay et al. (2007), the estimation process for $\boldsymbol{\theta}$ and \mathbf{c} implies the iteration of two profiling steps. First, for a given value of $\boldsymbol{\theta}$, τ and γ , the spline coefficients \mathbf{c} are estimated as the maximizer of J :

$$\begin{aligned}\hat{\mathbf{c}} &= \underset{\mathbf{c}}{\operatorname{argmax}} J(\mathbf{c}|\boldsymbol{\theta}, \tau, \gamma, \zeta) \\ &= \left(\tau \mathbf{B}^\top \mathbf{B} + \gamma \mathbf{R}(\boldsymbol{\theta}) \right)^{-1} \left(\tau \mathbf{B}^\top \zeta - \gamma \mathbf{r}(\boldsymbol{\theta}) \right).\end{aligned}\tag{5}$$

Then, given the last available $\hat{\mathbf{c}}$, $\boldsymbol{\theta}$ and τ are estimated by maximizing:

$$H(\boldsymbol{\theta}, \tau | \hat{\mathbf{c}}, \boldsymbol{\zeta}) = \frac{N}{2} \log(\tau) - \frac{\tau}{2} \|\boldsymbol{\zeta} - \mathbf{B}\hat{\mathbf{c}}\|^2. \quad (6)$$

The γ parameter weights and controls the relative emphasis on goodness-of-fit and solving the partial differential equation. This parameter measures the confidence that one has in the PDE to describe the dynamics in the system. Therefore, we suggest to christen it PDE-compliance parameter. As usual in penalized smoothing approach, the parameter γ has to be selected in a higher optimization level (e.g. using cross-validation, AIC, etc). As suggested by Schall (1991), γ can be seen as the ratio of the variance of the residuals and the variance of the penalty. Following this definition, we recommend to use an EM-type procedure to select it. Exact confidence intervals for the PDE parameters cannot easily be obtained. If one is interested in them, we suggest to use bootstrap-type confidence intervals.

2.2 Bayesian approach

We propose to adapt the frequentist proposal into a Bayesian framework. Contrary to Xun et al. (2013), we will use the same PDE-penalty than in the frequentist approach. The penalty term in Eq. (4) appears as a term subtracted from the log-likelihood. The log of the joint posterior yields the same fitting criterion using the following model specification:

$$\begin{cases} \boldsymbol{\zeta} | \mathbf{c}, \tau \sim \mathcal{N}_N(\mathbf{B}\mathbf{c}; \tau^{-1} \mathbf{I}_N), \\ p(\mathbf{c} | \boldsymbol{\theta}, \gamma) \propto \exp\left(-\frac{\gamma}{2} PEN\right). \end{cases}$$

The corresponding prior distribution for the spline coefficients is a multivariate normal distribution with mean $\mathbf{V}_1^{-1} \mathbf{v}_1$ and variance-covariance \mathbf{V}_1^{-1} where $\mathbf{v}_1 = -\gamma \mathbf{r}(\boldsymbol{\theta})$ and $\mathbf{V}_1 = \gamma \mathbf{R}(\boldsymbol{\theta})$.

Further prior distributions have to be given. For the precision of measurement τ , it is convenient to take a gamma distribution $\mathcal{G}(a_\tau, b_\tau)$ with mean a_τ/b_τ . If no prior information is available on this precision parameter, we recommend either to set $a_\tau = b_\tau$ equal to a small value (e.g. 10^{-6}) or to set a_τ equal to one and b_τ to a small value (e.g. 10^{-6}).

The PDE-compliance parameter appears as a precision parameter in the specification of the prior distribution for the spline coefficients. It is therefore convenient to take a gamma distribution $\mathcal{G}(a_\gamma, b_\gamma)$ as prior. In order to put prior confidence on the PDE model, we recommend to set a_γ to one and b_γ to a small value (e.g. 10^{-8}). Indeed, such a prior for γ is rather flat although it puts slightly more weight on value of $\log_{10}(\gamma)$ around $-\log_{10}(b_\gamma)$.

For the PDE parameter $\boldsymbol{\theta}$, the prior distribution depends on the problem and is denoted by $p(\boldsymbol{\theta})$.

The log joint posterior distribution for $(\mathbf{c}, \boldsymbol{\theta}, \gamma, \tau | \boldsymbol{\zeta})$ can be shown to be:

$$\begin{aligned} \log(p(\mathbf{c}, \boldsymbol{\theta}, \gamma, \tau | \boldsymbol{\zeta})) &= \frac{N}{2} \log(\tau) - \frac{\tau}{2} \|\boldsymbol{\zeta} - \mathbf{B}\mathbf{c}\|^2 \\ &+ \frac{1}{2} \log(\det(\mathbf{V}_1)) - \frac{1}{2} \{\mathbf{c}^\top \mathbf{V}_1 \mathbf{c} - 2\mathbf{c}^\top \mathbf{v}_1 + \mathbf{v}_1^\top \mathbf{V}_1^{-1} \mathbf{v}_1\} \\ &+ (a_\tau - 1) \log(\tau) - b_\tau \tau \\ &+ (a_\gamma - 1) \log(\gamma) - b_\gamma \gamma \\ &+ \log(p(\boldsymbol{\theta})). \end{aligned} \quad (7)$$

It can be easily shown that the conditional posterior distribution for the spline coefficients is a multivariate normal distribution with mean $\mathbf{V}_2^{-1} \mathbf{v}_2$ and variance-covariance \mathbf{V}_2^{-1} where $\mathbf{V}_2 = \tau \mathbf{B}^\top \mathbf{B} + \mathbf{V}_1$ and $\mathbf{v}_2 = \tau \mathbf{B}^\top \boldsymbol{\zeta} + \mathbf{v}_1$:

$$\mathbf{c} | \boldsymbol{\theta}, \tau, \gamma, \boldsymbol{\zeta} \sim \mathcal{N}_{M^p}(\mathbf{V}_2^{-1} \mathbf{v}_2; \mathbf{V}_2^{-1}).$$

The conditional posterior distribution for the precision of measurements and for the PDE-compliance parameter are gamma distributed:

$$\begin{aligned} \tau | \mathbf{c}, \boldsymbol{\zeta} &\sim \mathcal{G}\left(\frac{N}{2} + a_\tau; \frac{\|\boldsymbol{\zeta} - \mathbf{B}\mathbf{c}\|^2}{2} + b_\tau\right), \\ \gamma | \boldsymbol{\theta}, \mathbf{c} &\sim \mathcal{G}\left(\frac{M^p}{2} + a_\gamma; \frac{\mathbf{c}^\top \mathbf{R}(\boldsymbol{\theta}) \mathbf{c} + 2\mathbf{c}^\top \mathbf{r}(\boldsymbol{\theta}) + \mathbf{r}(\boldsymbol{\theta})^\top \mathbf{R}(\boldsymbol{\theta})^{-1} \mathbf{r}(\boldsymbol{\theta})}{2} + b_\gamma\right). \end{aligned}$$

The conditional posterior distribution for the PDE parameters $\boldsymbol{\theta}$ is not necessarily of a familiar type. Furthermore, high posterior correlation could occur between the spline coefficients \mathbf{c} and the elements of $\boldsymbol{\theta}$. Therefore, we recommend as in Jaeger and Lambert (2013) to marginalize the joint posterior distribution with respect to the spline coefficients, as given in Appendix A. This marginalization allows us to sample the PDE parameters without the need for updating the spline coefficients during MCMC. The coefficients can then be sampled in a second moment.

3 Estimation of the unknown differential parameters including state conditions

In many applications state (initial and/or boundary) conditions arise naturally or are implicitly defined by the observed dynamics (a clear real example is the analysis discussed in Section 5). Such information can be included in the statistical framework introduced in Section 2.

Consider the general partial problem given by:

$$\begin{cases} \mathcal{F}(\mathbf{x}, u, \boldsymbol{\theta}) = 0, \\ \frac{\partial^{(i)} u}{\partial x_0^{(i)}}(\mathbf{x}_0) = v(\mathbf{x}_0), \end{cases} \quad (8)$$

where \mathbf{x}_0 is a part of the domain where the state function u and/or its i th order derivatives are forced to be equal to the function v . Using the B-spline approximation, these conditions can be translated into :

$$\mathbf{H}\mathbf{c} = v(\mathbf{x}_0), \quad (9)$$

where \mathbf{H} is a tensor product B-spline matrix evaluated on the desired support points. In the frequentist framework, we propose two alternatives in order to introduce the state conditions in the penalized smoothing approach: least squares and Lagrange multipliers strategies. In the Bayesian settings, we only include the conditions using the least squares strategy.

3.1 Frequentist approach

One way to take into account the conditions related to a generic PDE is to consider them as an extra penalty. In the frequentist approach, the fitting criterion J is modified as follows:

$$J(\mathbf{c}|\boldsymbol{\theta}, \gamma, \kappa, \boldsymbol{\zeta}) = \frac{N}{2} \log(\tau) - \frac{\tau}{2} \|\boldsymbol{\zeta} - \tilde{u}(\mathbf{x})\|^2 - \frac{\gamma}{2} PEN(\boldsymbol{\theta}|\mathbf{c}) - \frac{\kappa}{2} (\mathbf{H}\mathbf{c} - v(\mathbf{x}_0))^\top (\mathbf{H}\mathbf{c} - v(\mathbf{x}_0)). \quad (10)$$

Note that with such a least square penalty, the approximated state function is not forced to be exactly equal to the conditions. With this fitting criterion, the optimal spline coefficients are computed as:

$$\begin{aligned} \hat{\mathbf{c}} &= \underset{\mathbf{c}}{\operatorname{argmax}} J(\mathbf{c}|\boldsymbol{\theta}, \tau, \gamma, \kappa, \boldsymbol{\zeta}) \\ &= \left(\tau \mathbf{B}^\top \mathbf{B} + \gamma \mathbf{R}(\boldsymbol{\theta}) + \kappa \mathbf{H}^\top \mathbf{H} \right)^{-1} \left(\tau \mathbf{B}^\top \boldsymbol{\zeta} - \gamma \mathbf{r}(\boldsymbol{\theta}) + \kappa \mathbf{H}^\top v(\mathbf{x}_0) \right). \end{aligned} \quad (11)$$

If κ tends to infinity, we simply force the state function to be equal to the state conditions at the prescribed points. On the other hand, if κ is equal to zero, then we just go back to the case considered in Section 2. Two alternatives exist to choose the κ parameter. The first one consists to fix it to a large value (10^6 , say). It is a reasonable choice if one is totally confident that the state conditions in (8) should be checked. On the other hand, this parameter can also be selected using standard criteria (such as cross-validation, AIC, etc). The selection of λ and κ requires to compute these criteria on a grid which can be time consuming. Therefore we advise to fix κ to a large value.

Using Lagrange multipliers in a frequentist framework, we can force the smoothing function to be exactly consistent with the conditions.

The Lagrange function for our constrained maximization problem is:

$$\mathcal{L}(\mathbf{c}, \boldsymbol{\omega}|\boldsymbol{\theta}, \tau, \gamma, \boldsymbol{\zeta}) = \frac{N}{2} \log(\tau) - \frac{\tau}{2} \|\boldsymbol{\zeta} - \tilde{u}(\mathbf{x})\|^2 - \frac{\gamma}{2} PEN(\boldsymbol{\theta}|\mathbf{c}) - \frac{1}{2} \boldsymbol{\omega}^\top (\mathbf{H}\mathbf{c} - v(\mathbf{x}_0)), \quad (12)$$

where $\boldsymbol{\omega}$ is the vector of Lagrange multipliers. As in Currie (2013), the maximization with respect to $(\mathbf{c}, \boldsymbol{\omega})$ follows from the solution of:

$$\begin{pmatrix} \tau \mathbf{B}^\top \mathbf{B} + \gamma \mathbf{R}(\boldsymbol{\theta}) & \mathbf{H}^\top \\ \mathbf{H} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{c} \\ \boldsymbol{\omega} \end{pmatrix} = \begin{pmatrix} \tau \mathbf{B}^\top \boldsymbol{\zeta} - \gamma \mathbf{r}(\boldsymbol{\theta}) \\ v(\mathbf{x}_0) \end{pmatrix}.$$

The precision of measurement τ and the PDE parameters $\boldsymbol{\theta}$ are then estimated in the same way as in Eq. (6).

3.2 Bayesian approach

In a Bayesian setting, the frequentist least squares strategy can be translated into the following prior for the spline coefficients:

$$p(\mathbf{c}|\boldsymbol{\theta}, \gamma, \kappa) \propto \exp\left(\frac{\gamma}{2}PEN - \frac{\kappa}{2}(\mathbf{H}\mathbf{c} - v(\mathbf{x}_0))^\top (\mathbf{H}\mathbf{c} - v(\mathbf{x}_0))\right).$$

This prior distribution corresponds to a multivariate normal distribution with mean $\mathbf{V}_1^{-1}\mathbf{v}_1$ and variance-covariance \mathbf{V}_1^{-1} where $\mathbf{v}_1 = -\gamma\mathbf{r}(\boldsymbol{\theta}) + \kappa\mathbf{H}^\top v(\mathbf{x}_0)$ and $\mathbf{V}_1 = \gamma\mathbf{R}(\boldsymbol{\theta}) + \kappa\mathbf{H}^\top\mathbf{H}$.

As in the frequentist approach, two alternatives are available for κ . One can either fix it to a large value or consider it as random with a gamma prior distribution $\mathcal{G}(a_\kappa; b_\kappa)$ where a_κ is equal to one and b_κ is equal to a small value (e.g. 10^{-6}). This choice for the prior distribution for κ translates prior confidence in the state conditions. The log joint posterior distribution, the log marginalized posterior distribution and all the conditional posterior distributions can be found in Appendix B. For the Lagrange multipliers strategy, the possible Bayesian translation would require to project the spline coefficients on a parameter sub-space where the constraints are met: we have no practical implementation of such a strategy for the moment.

4 Simulation

An intensive simulation study was set up to study the properties of the proposed estimation strategies. We consider the following advection-decay partial differential equation:

$$\mathcal{F}\left(x_1, x_2, u, \frac{\partial u}{\partial x_1}, \frac{\partial u}{\partial x_2}, \theta_1, \theta_2\right) = u_{x_1} + \theta_1 u_{x_2} + \theta_2 u = 0, \text{ with } x_1 \in \mathbb{R}, x_2 \in \mathbb{R}^+, \quad (13)$$

where u_{x_i} denotes the first derivative of the state function with respect to x_i . The general solution of (13) is:

$$u(x_1, x_2) = g(\theta_1 x_2 - x_1) \exp(-0.5\theta_2 x_1),$$

with $g(\cdot)$ an unknown function to be determined according to the state conditions. For simulation purposes, we take $\theta_1 = 0.5$ and $\theta_2 = 1.5$ and consider the condition $u(x_1, 0) = \frac{1}{1+x_1^2}$ (note that it is, in a certain sense, analogous to the terminal condition of the equation discussed in Section 5). Applying this state condition the differential problem has closed form solution:

$$u(x_1, x_2) = \frac{\exp(-3x_2)}{1 + 4x_2^2 - 4x_1x_2 + x_1^2}.$$

This expression contaminated with Gaussian noise is used to generate the data. This PDE appears particularly attractive as it has a closed form solution, enabling an objective comparison of the merits of the proposed constrained approaches with the nonlinear least squares one (see Section 4.1.2). Furthermore, the estimation abilities of our proposals are compared to the methods presented by Xun et al. (2013) (see Section 4.2).

4.1 Simulation using the PDE-P-splines approaches

The data used for the simulations of this section have been obtained by adding a Gaussian noise component to the analytic solution computed on the grid of equidistant points $x_{1,i} \in [-3, 3]$, $x_{2,j} \in [0, 1]$ with $i = 1, \dots, 50$ and $j = 1, \dots, 50$. Different levels for the precision of measurements $\tau = \{10000, 400, 100\}$ have been used to simulate the data in such a way to obtain a low, medium and large level of noise in the measurements. Figure 1 shows three possible simulated data clouds. We are interested in evaluating the efficiency of the proposed estimation procedures. For this reason we consider 500 simulated datasets per noise level and compute, for the frequentist approaches, the relative bias (in percent), the relative root mean squared error and the standard deviation of the estimated differential equation parameters together with the precision of measurements and the optimal compliance parameter. On the other hand, with the Bayesian approaches, we compute the relative bias (based on the posterior mean), the relative root mean squared error (based on the posterior mean), the mean posterior standard deviation and the 80% and 95% coverage probabilities (based on the HPD intervals).

In this simulation study we stress our proposal with and without taking into account the state conditions. For the frequentist procedures we include the state conditions using either the least squares or the Lagrange multiplier strategies. For the Bayesian approach, we only investigate the least squares strategy as it enables to marginalize the joint posterior distribution with respect to the spline coefficients.

In this simulation study we consider unequidistant knots. We use a set of 28 spline basis in the x_1 direction and 13 spline basis in the x_2 direction. In both direction, more knots have been located in the part of the domain where variation of

Figure 1: Three simulated data sets (dots) obtained considering an advection-decay dynamics and different levels for the precision of measurements. The surfaces represent the signal used in the simulation. The contour lines on the top of the perspective plots give a 1D representation of the simulated data (solid gray lines) and of the signal (solid black lines).

the signal appears more evident. The selection of knots location is not mandatory and fine grid of equidistant knots produces equivalent estimate but requires higher computational efforts. For the degree of the B-splines we found that a useful rule of thumb is to set it equal to the degree of the differential equation plus two.

4.1.1 Simulation results without considering state conditions

In Table 1 the results of the simulation study described above are shown. In this case the state conditions have not been taken into account in the estimation process. The cited table shows some interesting features. The relative bias (in percent) for the estimated parameters is small both for the frequentist and Bayesian approaches. On the other hand, looking at the estimated PDE parameters, the relative bias obtained within the Bayesian framework is positive and larger (but still very small) than for the frequentist estimates. This could be explained by the choice of the posterior mean for the point estimate while the estimates for the precision parameter look similar for the two approaches. The RMSEs seem to be larger for the PDE parameters in the Bayesian approach and larger for the precision of measurement in the frequentist approach. This can be explained looking at the bias. On the other hand, based on the STDs, the estimation procedure seems quite robust against the noise level. For the Bayesian estimates, the mean posterior standard deviations tends to increase with the level of the noise for the PDE parameters and to decrease when the noise level increases for the τ parameter. Only some estimated coverage probabilities (at 80% and 95% level) are in agreement with the nominal values. On the other hand, for a high level of noise, there is evidence of an under coverage effect that could be explained by the fact that we are estimating the state conditions without considering them. Finally, the (average) optimal compliance parameters γ shown at the bottom of each sub-table are fairly large indicating a strong (posterior) confidence in the appropriateness of the differential model used in the penalty to describe the observed dynamics. Figure 2 shows three possible smoothing surfaces coming from this simulation study. As expected, if we don't include conditions in the estimation procedure, the fitted surface tends to be sensitive to noise variability showing wiggly part for large values of x_1 and moderate value of x_2 (see first column of Figure 2).

4.1.2 Simulation results considering the state conditions

The results of the simulation study taking into account the conditions using least squares (both frequentist and Bayesian approaches with fixed $\kappa = 10^6$) or Lagrange multipliers (only frequentist approach) are shown in Table 2. The simulation performances obtained using a nonlinear least squares solution of the state function are shown in Table 3 (note that the knowledge of the state conditions is necessary to apply the nonlinear least squares estimation procedure). It appears that the relative bias for the estimated parameters is small both for the frequentist and Bayesian approaches. One can notice that the bias decreases in absolute value when the state conditions are introduced in the estimation procedure in the case of moderate and high noise variability. As before, for the frequentist estimates, the RMSEs seem to increase together with the level of measurement noise. On the other hand, based on the STDs, the estimation procedure seems quite robust against the noise level. The estimated standard deviations for the PDE parameters seem to be lower than what we observed without considering the conditions. Looking at the frequentist estimates the results obtained imposing conditions by least squares and using Lagrange multipliers look really similar. For the Bayesian estimates,

Table 1: Frequentist and Bayesian simulation for advection-decay PDE smoothing example without considering the state conditions. At the end of each section in the table the average optimal compliance parameters are listed.

		True parameter values				
		$\theta_1 = 0.5$	$\theta_2 = 1.5$	$\tau = 1/\sigma^2$		
PDE-P-splines Frequentist	R-BIAS	$\sigma^2 = 0.01^2$	8.18E-02%	8.76E-02%	2.55E+00%	
		$\sigma^2 = 0.05^2$	4.47E-01%	4.74E-01%	2.22E+00%	
		$\sigma^2 = 0.1^2$	1.16E+00%	1.21E+00%	6.48E-01%	
	R-RMSE	$\sigma^2 = 0.01^2$	4.49E-03	5.00E-03	3.91E-02	
		$\sigma^2 = 0.05^2$	2.27E-02	2.52E-02	3.85E-02	
		$\sigma^2 = 0.1^2$	4.67E-02	6.81E-01	2.99E-02	
	R-STD	$\sigma^2 = 0.01^2$	4.42E-03	4.93E-03	2.98E-02	
		$\sigma^2 = 0.05^2$	2.23E-02	2.48E-02	3.15E-02	
		$\sigma^2 = 0.1^2$	4.53E-02	5.04E-02	2.92E-02	
	$E(\gamma) = \{2.36E+08, 5.20E+07, 4.30E+07\}$					
	PDE-P-splines Bayesian	R-BIAS	$\sigma^2 = 0.01^2$	6.45E-02 %	6.89E-02 %	3.25E-01%
			$\sigma^2 = 0.05^2$	1.14E+00%	1.35E+00%	1.48E-01%
$\sigma^2 = 0.1^2$			4.02E+00%	4.79E+00%	9.00E-02%	
R-RMSE		$\sigma^2 = 0.01^2$	4.46E-03	4.96E-03	2.93E-02	
		$\sigma^2 = 0.05^2$	2.54E-02	2.86E-02	2.90E-02	
		$\sigma^2 = 0.1^2$	6.27E-02	7.20E-02	2.90E-02	
MPSD		$\sigma^2 = 0.01^2$	2.19E-03	7.31E-03	2.85E+02	
		$\sigma^2 = 0.05^2$	1.10E-02	3.66E-02	1.14E+01	
		$\sigma^2 = 0.1^2$	2.35E-02	7.84E-02	2.84E+00	
CP80		$\sigma^2 = 0.01^2$	80.2	80.6	80.6	
		$\sigma^2 = 0.05^2$	76.6	73.4	80.0	
		$\sigma^2 = 0.1^2$	69.4	68.6	79.6	
CP95		$\sigma^2 = 0.01^2$	94.4	93.4	92.6	
		$\sigma^2 = 0.05^2$	91.0	93.0	92.2	
		$\sigma^2 = 0.1^2$	89.2	87.0	93.2	
$E(\gamma) = \{1.48E+07, 2.02E+07, 3.01E+07\}$						

the mean posterior standard deviations tends to increase with the level of the noise for the PDE parameters and to decrease when the noise level increases for the τ parameter. As expected, including the differential condition, the mean standard deviation for the PDE parameters tends to decrease. The estimate coverage probabilities are in agreement with the nominal values in almost all the simulation settings. This is probably due to the fact that the conditions are explicitly specified and not estimated. The average optimal compliance parameters γ given in each sub-tables, indicate a stronger (posterior) confidence in the PDE model as descriptor of the dynamics of the observed state. The middle and the third column of Figure 2 show three smoothing surfaces extracted from the simulation runs. Compared to the first column, the inclusion of the state conditions forces the smoothing surface to be adherent to the closed form solution over the entire domain even for lower precision of the measurements. Finally, the performance of the proposed frequentist and Bayesian estimation procedures are compatible with those obtained through nonlinear least squares estimation of the parameters involved in the analytic solution.

4.2 Comparison of the PDE-P-splines approaches with the one by Xun et al. (2013)

In the Bayesian framework, Xun et al. (2013) proposes to use a penalty that combines the requirement for PDE compliance (in a L^2 sense) and the one for the smoothness of the estimated surface (using extra finite difference penalties). Our interest is to compare this proposal with our Bayesian approaches (that, as shown above, can be exploited with known and unknown state conditions). Furthermore, we also consider a modified version of the penalty proposed by Xun et al. (2013) not including the extra finite difference penalties in order to evaluate their impact on the quality of the estimates. The two main differences between our proposals and the one by Xun et al. (2013) lay in the construction of the penalty term and in computation of the appropriate constant of normalization for the prior distribution for the splines coefficients. For the construction of the penalty term, we used an integrated penalty that ensure the fidelity of the B-spline approximation to the proposed PDE. Note that our Bayesian approach is a direct “Bayesianization” of the frequentist ones. As appears in Eq. (7), we compute the constant of normalization of the prior distribution for the spline coefficients as the determinant of the PDE-penalty matrix whereas Xun et al. (2013) just consider a power of the smoothing parameters as constant of normalization for that prior.

The data used for the comparison have been obtained by adding a Gaussian noise component ($\tau = 400$) to the analytic solution of the PDE given in Eq. (13) computed on the grid of equidistant points $x_{1,i} \in [-4, 4]$, $x_{2,j} \in [0, 1]$ with $i = 1, \dots, n$, where $n \times m = 40 \times 20, 100 \times 50, 200 \times 100$ (low, medium and large sample sizes).

Figure 3 shows the estimated surfaces obtained using the four approaches. For the large sample sizes, all the estimated smoothing surfaces appear adherent to the state function, except for the approach of Xun et al. (2013) that seems less efficient in reproducing the peak of the signal. When the sample size decreases, our approaches seem to perform better than the competitors in recovering the underlying signal. In the case where the sample is limited, the approach of Xun et al. (2013) tends to oversmooth the data and seems to put not enough weight on the PDE-model penalty. On the other hand, for the same simulation settings, the modified version of the model by Xun et al. (2013) without the extra finite difference penalty seems to better catch the signal. The posterior densities for the PDE parameters in the modified Xun et al. approach appears clearly left shifted (see Figure 4). The extra finite difference penalty seems to mitigate the introduced bias. The complete constant of normalization in the prior distribution for the spline coefficients centers the posterior distributions around the true PDE parameters. Finally, whatever the sample size, the variability of the posteriors decreases when the state conditions are introduced in the model.

5 Application

In this section we evaluate the performance of the proposed approaches dealing with a real data application. We consider as example the analysis of a series of option prices traded for different strike prices and maturities. Quoting Cox, Ross, and Rubinstein (1979), an option is a contract that gives the right, but not the obligation, to buy or sell a risky asset at a predetermined fixed price within (American style) or at (European style) a given date (maturity of the contract). This kind of financial instruments allow to bet about the future evolution of the price of the underlying asset that can be another contract (a stock for example), a market index or a commodity. Options and other related instruments are often called “derivatives” given that their price is derived from the market value (S) of the underlying risky asset.

Two different kinds of option contracts are distinguished according to the right that they give to the holder. “Call” options give the right to buy the underlying asset while “put” give the right to sell it. These characteristics determine the payoff functions at maturity (T):

$$C(S, T) = \begin{cases} 0 & \text{if } S_T \leq E \\ S_T - E & \text{if } S_T > E, \end{cases}$$

Table 2: Frequentist and Bayesian simulation for advection-decay PDE smoothing example considering the state conditions. At the end of each section in the table the average optimal compliance parameters are listed.

		True parameter values				
		$\theta_1 = 0.5$	$\theta_2 = 1.5$	$\tau = 1/\sigma^2$		
PDE-P-splines Frequentist (Least squares)	R-BIAS	$\sigma^2 = 0.01^2$	1.06E-01 %	7.27E-02 %	1.51E+00 %	
		$\sigma^2 = 0.05^2$	1.46E-01%	5.08E-02%	2.12E+00%	
		$\sigma^2 = 0.1^2$	1.62E-01%	2.02E-02 %	5.84E-01%	
	R-RMSE	$\sigma^2 = 0.01^2$	3.56E-03	3.83E-03	6.20E-02	
		$\sigma^2 = 0.05^2$	1.62E-02	1.77E-02	3.69E-02	
		$\sigma^2 = 0.1^2$	3.25E-02	6.68E-01	2.98E-02	
	R-STD	$\sigma^2 = 0.01^2$	3.40E-03	3.77E-03	6.02E-02	
		$\sigma^2 = 0.05^2$	1.61E-02	1.77E-02	3.03E-02	
		$\sigma^2 = 0.1^2$	3.25E-02	3.57E-02	2.92E-02	
	$E(\gamma) = \{2.10E+08, 2.94E+07, 1.21E+07\}$					
	PDE-P-splines Frequentist (Lagrange)	R-BIAS	$\sigma^2 = 0.01^2$	1.12E-01%	6.14E-02%	2.24E+00%
			$\sigma^2 = 0.05^2$	1.44E-01%	4.73E-02%	2.13E+00%
$\sigma^2 = 0.1^2$			1.59E-01%	1.37E-02%	5.78E-01%	
R-RMSE		$\sigma^2 = 0.01^2$	3.40E-03	3.58E-03	3.72E-02	
		$\sigma^2 = 0.05^2$	1.62E-02	1.77E-02	3.70E-02	
		$\sigma^2 = 0.1^2$	3.25E-02	6.68E-01	2.98E-02	
R-STD		$\sigma^2 = 0.01^2$	3.22E-03	3.53E-03	2.97E-02	
		$\sigma^2 = 0.05^2$	1.62E-02	1.77E-02	3.03E-02	
		$\sigma^2 = 0.1^2$	3.25E-02	3.57E-02	2.93E-02	
$E(\gamma) = \{2.11E+08, 3.62E+07, 2.08E+07\}$						
PDE-P-splines Bayesian		R-BIAS	$\sigma^2 = 0.01^2$	1.56E-02%	8.29E-03%	2.78E-01%
			$\sigma^2 = 0.05^2$	2.42E-01%	2.24E-01%	2.33E-01%
	$\sigma^2 = 0.1^2$		8.14E-01%	8.21E-01%	2.28E-01%	
	R-RMSE	$\sigma^2 = 0.01^2$	3.22E-03	3.54E-03	2.91E-02	
		$\sigma^2 = 0.05^2$	1.64E-02	1.79E-02	2.90E-02	
		$\sigma^2 = 0.1^2$	3.39E-02	3.71E-02	2.91E-02	
	MPSD	$\sigma^2 = 0.01^2$	1.61E-03	5.38E-03	2.84E+02	
		$\sigma^2 = 0.05^2$	7.70E-03	2.57E-02	1.14E+01	
		$\sigma^2 = 0.1^2$	1.56E-02	5.21E-02	2.84E+00	
	CP80	$\sigma^2 = 0.01^2$	80.2	80.2	80.6	
		$\sigma^2 = 0.05^2$	78.4	78.2	80.6	
		$\sigma^2 = 0.1^2$	78.0	79.0	79.8	
	CP95	$\sigma^2 = 0.01^2$	95.0	94.2	93.2	
		$\sigma^2 = 0.05^2$	92.4	93.4	93.0	
		$\sigma^2 = 0.1^2$	93.2	93.2	93.0	
$E(\gamma) = \{1.05E+08, 1.48E+08, 1.37E+08\}$						

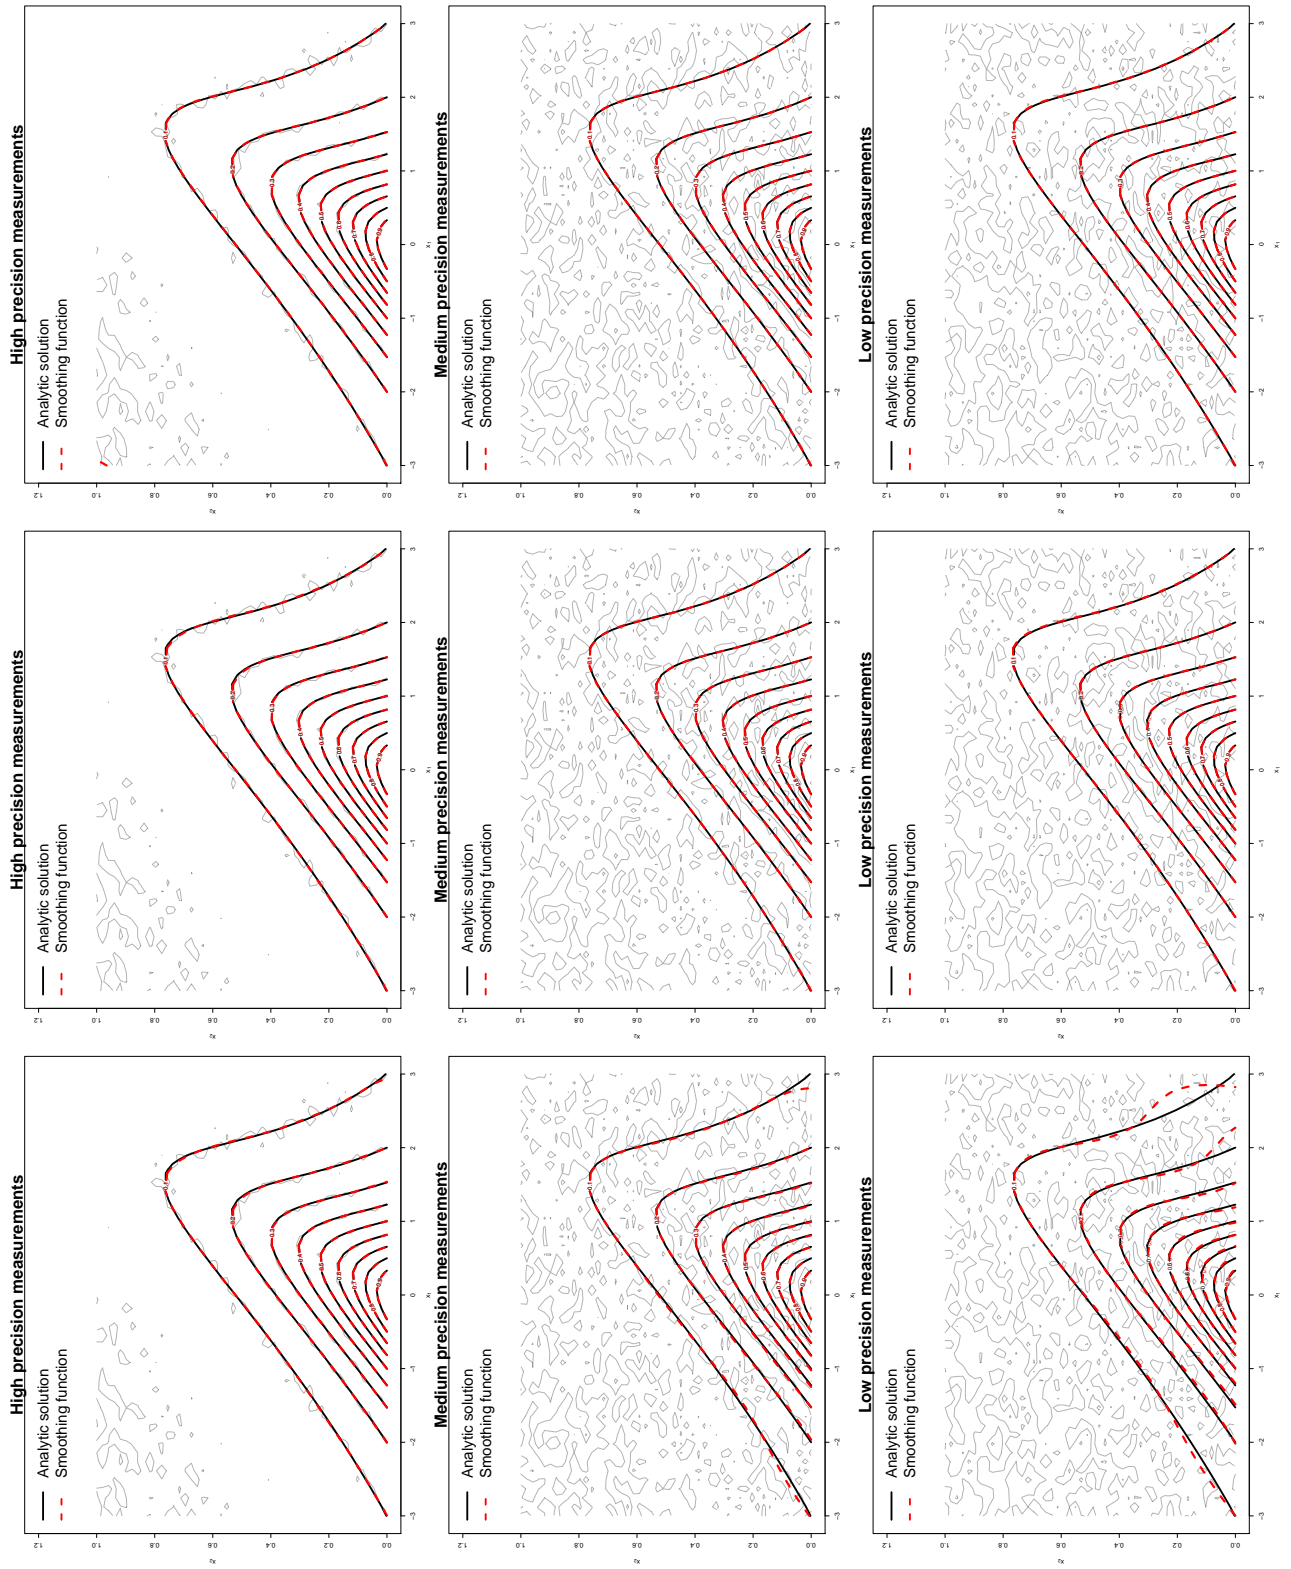


Figure 2: Typical profiles for the frequentist approaches. The first column shows the smoothing surfaces estimated for unknown conditions. The middle column shows the smoothing surfaces estimated using the state conditions imposed by least squares, while the third column shows the smoothing surface using the Lagrange multipliers approach. The gray level lines indicate the contour of the simulated observations.

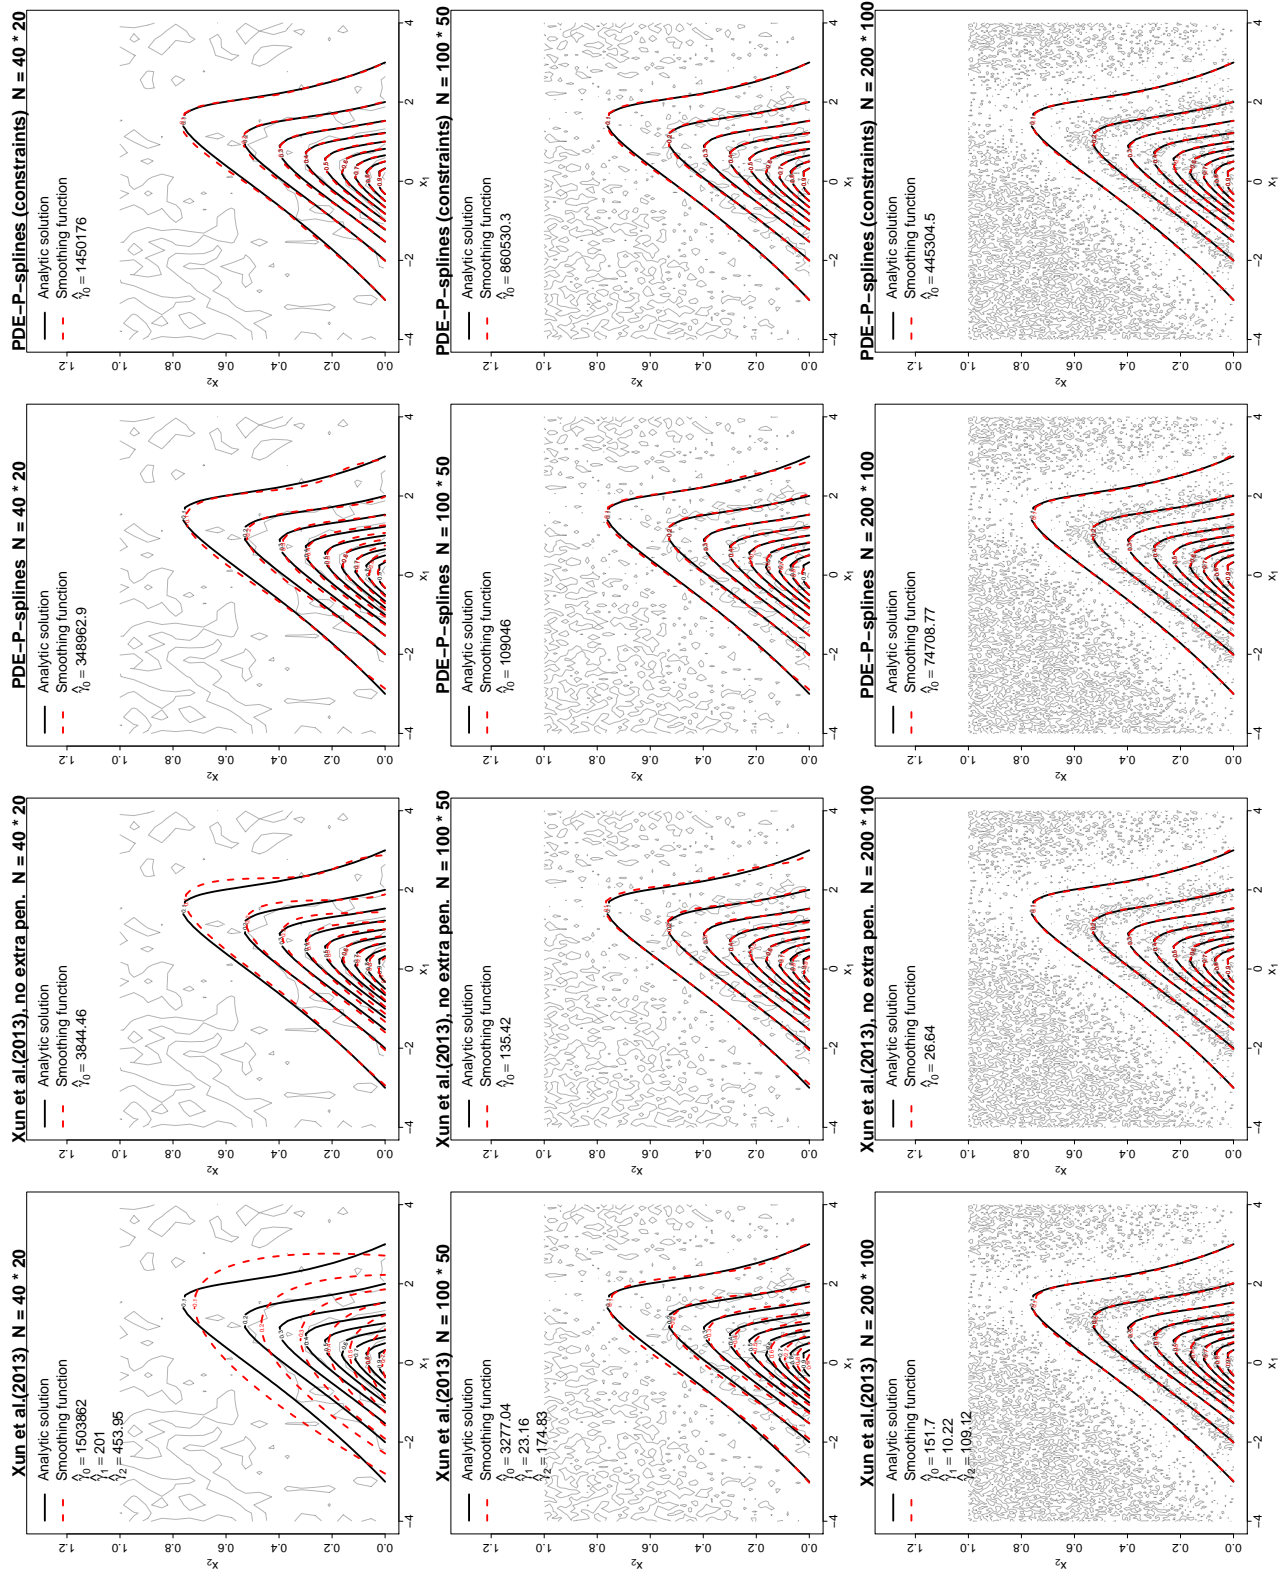


Figure 3: Posterior pointwise smoothing surfaces estimated with Xun et al. (2013) approach with and without the extra finite difference penalty (first and second rows) and smoothing surfaces estimated using the PDE-P-splines approaches without and with the boundary conditions imposed by least squares (third and fourth rows). In each column, the surfaces estimated for different sample sizes are depicted ($n = 40 \times 20$ in the left column, $n = 100 \times 50$ in the middle column and $n = 200 \times 100$ in the right column). The gray level lines indicates the contour of the simulated observations.

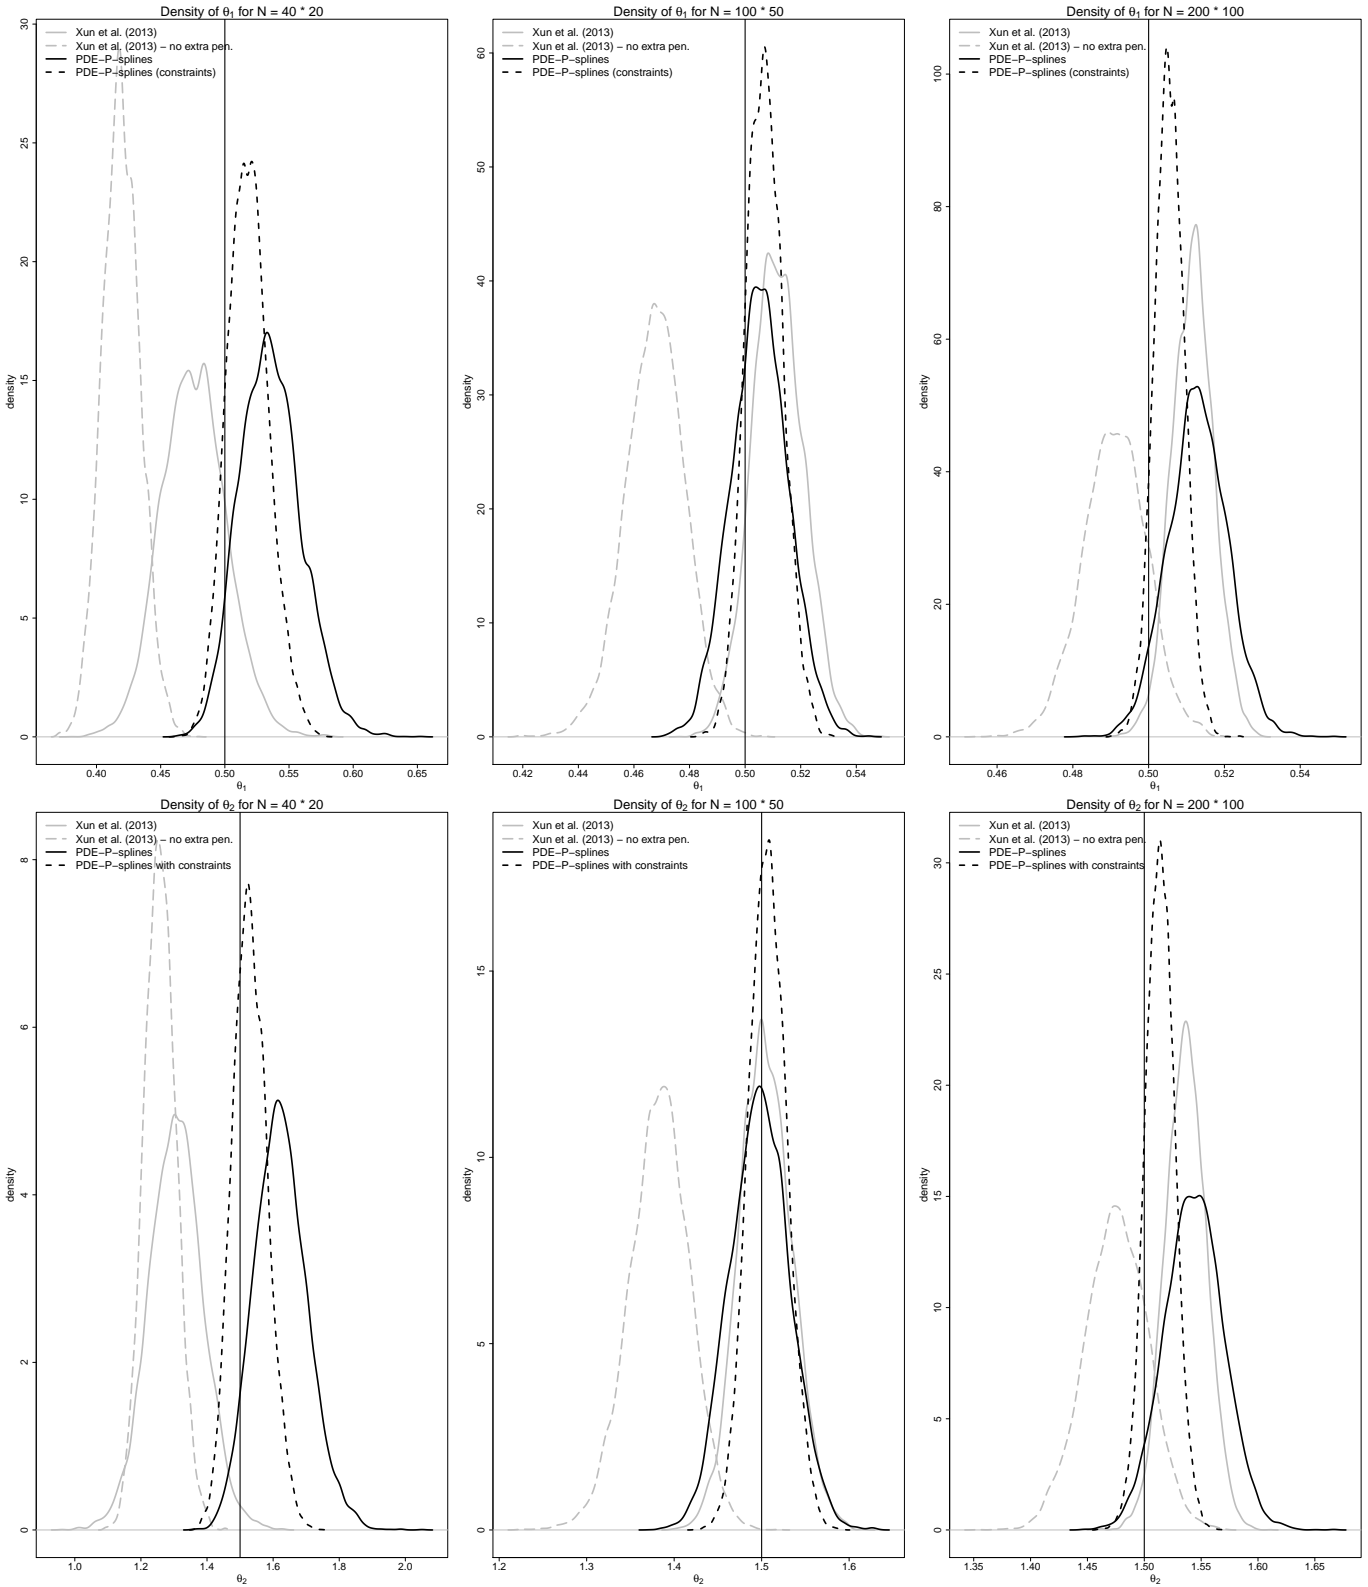


Figure 4: Posterior densities for $\theta_1 = 0.5$ and $\theta_2 = 1.5$, respectively. Densities in gray are those obtained using the approaches proposed by Xun et al. (2013) (solid curve with extra finite difference penalty and dotted curve without it). Densities in black are those obtained using the PDE-P-splines approaches (solid curve without the differential condition and dashed one with it). In each column, the posterior densities for different sample sizes are depicted ($n = 40 \times 20$ in the left column, $n = 100 \times 50$ in the middle column and $n = 200 \times 100$ in the right column).

Table 3: Frequentist nonlinear least squares simulation for advection-decay PDE smoothing example.

		True parameter values			
		$\theta_1 = 0.5$	$\theta_2 = 1.5$	$\tau = 1/\sigma^2$	
Frequentist (NLS)	R-BIAS	$\sigma^2 = 0.01^2$	1.21E-02 %	-1.47E-03 %	-1.10E-01 %
		$\sigma^2 = 0.05^2$	4.16E-02 %	7.12E-02 %	1.96E-01 %
		$\sigma^2 = 0.1^2$	-7.05E-02 %	-5.27E-02 %	7.83E-01 %
	R-RMSE	$\sigma^2 = 0.01^2$	1.39E-03	2.90E-03	1.40E-04
		$\sigma^2 = 0.05^2$	6.70E-03	1.38E-02	6.95E-04
		$\sigma^2 = 0.1^2$	1.26E-02	2.71E-02	1.42E-03
	R-STD	$\sigma^2 = 0.01^2$	1.39E-03	2.90E-03	1.40E-04
		$\sigma^2 = 0.05^2$	6.71E-03	1.37E-02	6.96E-04
		$\sigma^2 = 0.1^2$	1.26E-02	2.71E-02	1.42E-03

$$P(S, T) = \begin{cases} 0 & \text{if } S_T \geq E \\ E - S_T & \text{if } S_T < E, \end{cases}$$

where C and P indicate call and put prices respectively, E is the strike price and S_T represents the market price of the underlying asset at maturity. The current price of the contract is then computed as the present value of the payoff at maturity. For this reason it is necessary to describe the dynamics of the underlying asset in order to price the derivative contract “written” on it.

The Black and Scholes (1973) framework is the most famous pricing model for financial derivatives. The B&S model comes from the assumption that the price of the underlying risky asset follows a geometric Brownian motion in the instantaneous time with a constant standard deviation (volatility). This assumption leads to the well known second order partial differential equation:

$$\mathcal{F} \left(S, t, C, \frac{\partial C}{\partial S}, \frac{\partial C}{\partial t}, \frac{\partial^2 C}{\partial S^2}, \sigma \right) = \frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} - rC = 0, \quad (14)$$

where the parameter σ identifies the “implied volatility” and r is the observed risk free rate. For European options (exercise allowed only at a given date) a close form solution can be found introducing the following no-arbitrage terminal (related to the expiration of the contract) and boundary conditions:

$$\begin{aligned} C(0, t) &= 0, \\ C(S, t) &\approx S \text{ for } S \rightarrow \infty, \\ C(S, T) &= \max(S - E, 0) = (S - E)^+. \end{aligned} \quad (15)$$

This framework provides a concise and easy to interpret description of the price dynamics. Furthermore, this model allows to compute the theoretical option prices for unobserved strikes and maturities and, inverting the solution of the PDE above, to recover the volatility (σ) implied in the observed prices (the volatility parameter that solves the equation for the observed prices). Moving from these considerations, we suggest to achieve these goals including the PDE in (14) and the no-arbitrage state conditions in (15) in a PDE-P-spline approach. Note that in this application we share with the theoretical model the unrealistic assumption of constant volatility over strikes and maturities. Besides the limitations implicit in the B&S model, we believe that the proposed estimation procedure is flexible enough to ensure appropriate estimates even if the hypotheses of the differential model do not completely match with the data evidence.

We conduct our analysis using the data proposed by Carmona (2004). The dataset counts 2800 prices of European call options written on the SP-500 index and traded in 1993 with expiration between 1993 and 1994. The six variables included in the dataset indicate the value of the index, the strike prices of the option, the time to expiration (in fraction of year), the spot interest rate, the observed implied volatilities and the prices of the options. This application is, in our opinion, particularly interesting because it represents one of those cases in which the state conditions are known a priori (given by theoretical arguments) and need to be taken into account in order to obtain valid estimates.

The raw data and the estimated smoothing surface (obtained through a soft constrained frequentist approach) are shown in figure 5. These results have been obtained using tensor product B-splines of degree four built on 25 equidistant internal knots for both the maturity and moneyness directions. The moneyness parametrization of the PDE allows for a simple formulation of the penalty and of the necessary state conditions as discussed by Piché and Kannianen

(2007). The estimated implied volatilities are equal to 1.003E-01 and 1.014E-01 for the frequentist and Bayesian procedures respectively and are in agreement with the median of the ones observed in the dataset. The frequentist 95% confidence interval (obtained through 1000 bootstrap replicates) is (9.940E-01; 1.012E-01) while the corresponding credibility interval has been found equal to (1.007E-01; 1.020E-01). The lower left panel of Figure 5 shows the posterior distribution of the estimated parameter.

The estimates obtained using our approach shows some desirable properties. First of all, we are able to extrapolate the prices associated with strikes and maturities not traded, and those estimates are consistent with the no arbitrage constraints (due to the compliance to the Black and Scholes model). Furthermore, it is particularly appreciable in real data analysis to have confidence/credible bounds for prices associated to unobserved strikes and/or maturities. On the other hand, the PDE-P-spline approach allows here to obtain point and intervals estimates of the implied volatility consistent with the Black and Scholes equation ensuring a good performance in terms of data fitting. Frequentist and Bayesian approaches suggest compatible volatility parameters with narrow confidence/credibility intervals. These results are consistent with the ones obtained in our simulations. The smoothing surface (frequentist estimates) in Figure 5 (left panel) is compatible with the observed cloud of prices. The Bayesian analysis gives results similar really results as expected (not shown).

6 Discussion

In this paper, we propose PDE-based penalized tensor product B-spline smoothing approach in order to solve and estimate unknown parameters in partial differential equations. Our aim was to introduce frequentist and a Bayesian procedures to analyze data which dynamics is defined over more than one dimension. Both approaches exploit a tensor product B-spline approximation of the state function while the consistency of the final estimates to the PDE is ensured by PDE-based soft constraints. The compromise between the data fitting and consistency to the PDE is tuned by an compliance parameter. As in Xun et al. (2013), we estimate the spline coefficients, the PDE parameters and the precision of measurement using the information implied in the observation. If state conditions are not included in the model, our approach and the one by Xun et al. (2013) are equivalent except that the PDE-compliance parameter is selected through an EM-like algorithm (Schall, 1991). For the Bayesian approach, we use the same PDE based penalty as the one we introduced in the frequentist approach. In particular, our Bayesian formulation takes into account the functional dependence of the normalizing constant of the spline coefficient prior on PDE parameters. Note that prior confidence in the PDE model is introduced by a specific choice for the prior distribution of the compliance parameter. Even if their knowledge is not strictly required for the application of the proposed methods, we are able to improve the quality of the estimates by introducing state conditions as extra constraints using either a L^2 penalty or Lagrange multipliers. On the other hand, the introduction of these conditions arises sometimes naturally from theoretical arguments and cannot be omitted during the estimation process.

We demonstrate through simulations that our PDE-P-spline approaches show desirable properties and that the inclusion of the state conditions has an influence on the estimation performances, reducing the variability of the PDE parameters and improving their coverages. We also compare our Bayesian proposals with the one suggested by Xun et al. (2013). We found that, within their Bayesian approach, the introduction of an additional finite difference penalty help to reduce the bias in estimating the PDE parameters. Indeed, we think that this additional penalty tends to mitigate the effect of the miss-specified constant of normalization in the prior of the spline coefficients. On the other hand, the same extra term tends to produce smoothing surface not adherent to the state function for small sample sizes. These undesirable features are not shared by our Bayesian proposal. Indeed, even for small sample sizes and/or large level of noise, the use of the proper prior for the spline coefficients ensures good fitting and estimating performances.

As real data application, we analyze the SP-500 call option prices discussed in Carmona (2004). The proposed example represents one of those cases where the state conditions play a role in the estimation process ensuring that the smoothing surface is consistent with the no-arbitrage requirements and so cannot be ignored. We use the unrealistic Black and Scholes model to define a suitable PDE-based penalty even if, paying a price in terms of computational complexity and interpretability of the final results, it is possible to consider more sophisticated and realistic pricing models.

Our further research will focus on possible extension of the PDE-P-spline approach. The volatility in the Black and Scholes model is supposed to be constant. This assumption is known as unrealistic. The first solution in order to deal with the time-dependent volatility could be to consider a more realistic differential pricing model (e.g. the Heston model (Heston, 1993)). In our opinion, the same aim can be achieved by allowing the implied volatility parameter in the Black and Scholes equation to vary nonparametrically with the time (e.g. using P-spline representation).

In this paper, we deal with linear partial differential equations. Moving to nonlinear partial differential equation would lead to additional challenges. First of all, the PDE-penalty term would not be a second order homogeneous polynomial in the spline coefficients. In the frequentist approach, this would require the use of the implicit function theorem to estimate the PDE parameters (Ramsay et al., 2007). On the other hand, in a Bayesian setting, the constant of

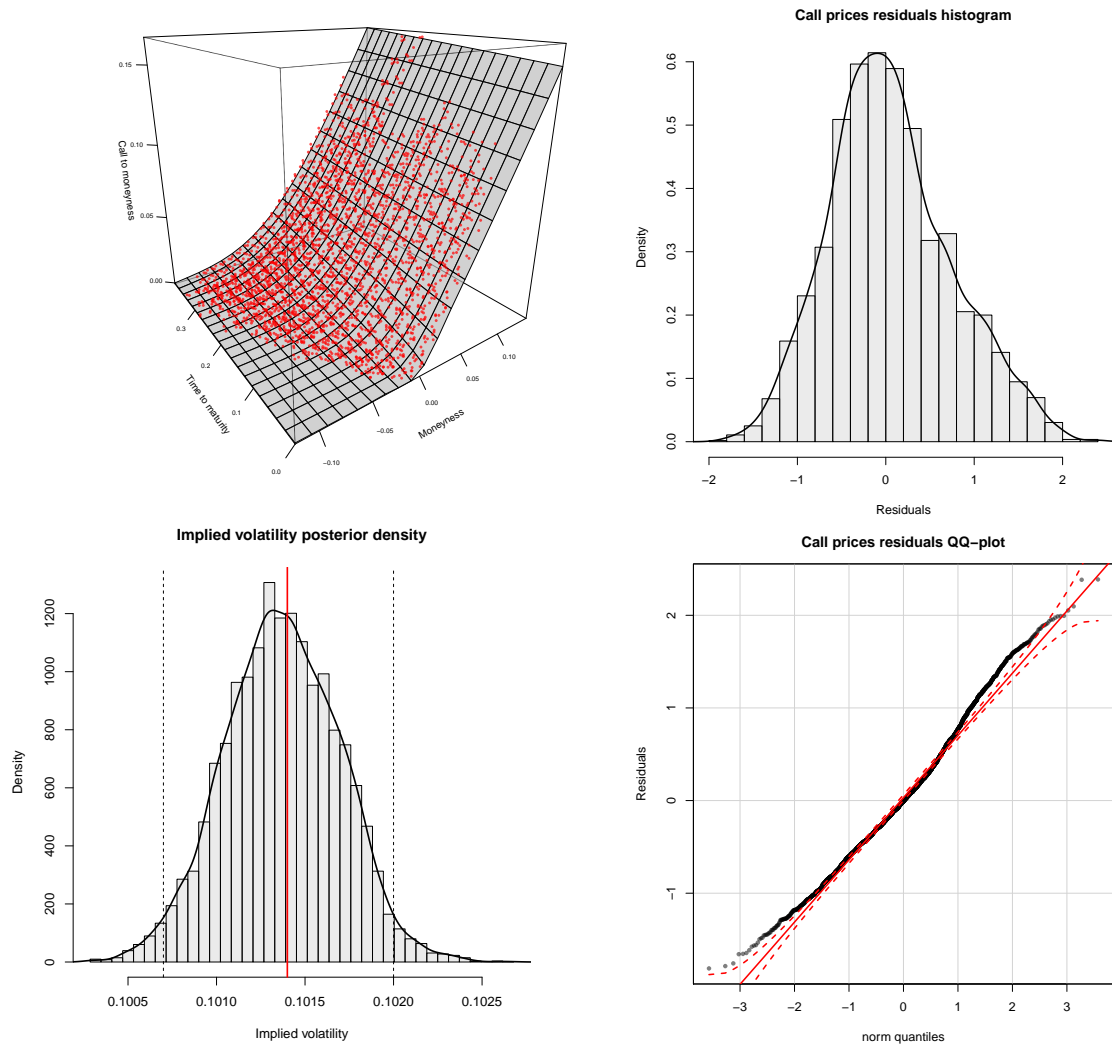


Figure 5: Left upper panel: observations (red dots) and smoothing surface obtained applying a frequentist constrained PDE-P-spline estimation procedure to the SP-500 European call option prices analyzed by Carmona (2004). The prices in the plot are scaled by the observed strikes and the moneyness is defined as the log of the ratio between the observed spot and strike prices. The estimated implied volatility was $\hat{\sigma} = 1.003E - 01$ with an optimal compliance parameter equal to $\hat{\gamma} = 47385275$. Left lower panel: posterior density of the implied volatility parameter (Bayesian estimates). Right panels: histogram and qq-plot of the residuals (call price scale) between the smoothing surface and the observed SP-500 European call option prices.

normalization for the prior distribution for the spline coefficients will not have an explicit form anymore.

Acknowledgment

Support from the IAP Research Network P7/06 of the Belgian State (Belgian Science Policy) is gratefully acknowledged.

References

- L Biegler, J Damiano, and G Blau. Nonlinear parameter estimation: A case study comparison. *AIChE Journal*, 32(1): 29–45, 1986.
- F Black and M Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- R Carmona. *Statistical Analysis of Financial Data in S-Plus*. Springer Text in Statistics, 2004.
- J Cox, S Ross, and M Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3): 229–263, 1979.
- I D Currie. Smoothing constrained generalized linear models with an application to the lee-carter model. *Statistical Modelling*, 13(1):69–93, 2013.
- P Eilers and B Marx. Flexible smoothing with B-splines and penalties. *Statistical Science*, 11:89–121, 1996.
- S Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343, 1993.
- J Jaeger and P Lambert. Bayesian P-spline estimation in hierarchical models specified by systems of affine differential equations. *Statistical Modelling*, 13(1):3–40, 2013.
- R Piché and J Kannianen. Solving financial differential equations using differentiation matrices. In *World Congress on Engineering*, Lecture Notes in Engineering and Computer Science. Newswood Limited, 2007.
- J O Ramsay, G Hooker, D Campbell, and J Cao. Parameter estimation for differential equations: a generalized smoothing approach. *Journal of the Royal Statistical Society, Series B*, 69:741–796, 2007.
- R Schall. Estimation in generalized linear models with random effects. *Biometrika*, 78(4):719–727, 1991.
- M H Schultz and R S Varga. L-splines. *Numerische Mathematik*, 10(4):345–369, 1967.
- X Xun, J Cao, B Mallick, R Carroll, and A Maity. Parameter estimation of partial differential equation models. *Journal of the American Statistical Association*, 108(503):1009–1020, 2013.

A Logarithm of the marginalized posterior distribution

The joint posterior distribution given in Eq. (7) is marginalized with respect to the spline coefficients. The log of the marginalized posterior distribution can be shown to be:

$$\begin{aligned}
 \log(p(\boldsymbol{\theta}, \gamma, \tau | \boldsymbol{\zeta})) &= \frac{N}{2} \log(\tau) - \frac{\tau}{2} \boldsymbol{\zeta}^\top \boldsymbol{\zeta} \\
 &+ \frac{1}{2} \log(\det(\mathbf{V}_1)) - \frac{1}{2} \mathbf{v}_1^\top \mathbf{V}_1^{-1} \mathbf{v}_1 \\
 &- \frac{1}{2} \log(\det(\mathbf{V}_2)) + \frac{1}{2} \mathbf{v}_2^\top \mathbf{V}_2^{-1} \mathbf{v}_2 \\
 &+ (a_\tau - 1) \log(\tau) - b_\tau \tau \\
 &+ (a_\gamma - 1) \log(\gamma) - b_\gamma \gamma \\
 &+ \log(p(\boldsymbol{\theta})).
 \end{aligned}$$

B Posterior distributions including the DE conditions by least squares

The log joint posterior distribution can be shown to be:

$$\begin{aligned}
\log(p(\mathbf{c}, \boldsymbol{\theta}, \gamma, \kappa, \tau | \boldsymbol{\zeta})) &= \frac{N}{2} \log(\tau) - \frac{\tau}{2} \|\boldsymbol{\zeta} - \mathbf{B}\mathbf{c}\|^2 \\
&+ \frac{1}{2} \log(\det(V_1)) - \frac{1}{2} \{\mathbf{c}^\top \mathbf{V}_1 \mathbf{c} - 2\mathbf{c}^\top \mathbf{v}_1 + \mathbf{v}_1^\top \mathbf{V}_1^{-1} \mathbf{v}_1\} \\
&+ (a_\tau - 1) \log(\tau) - b_\tau \tau \\
&+ (a_\gamma - 1) \log(\gamma) - b_\gamma \gamma \\
&+ (a_\kappa - 1) \log(\kappa) - b_\kappa \kappa \\
&+ \log(p(\boldsymbol{\theta})).
\end{aligned}$$

From this joint posterior distribution, only two conditional posterior distribution can be identified. The conditional posterior distribution for the spline coefficients is a multivariate normal distribution:

$$\mathbf{c} | \boldsymbol{\theta}, \tau, \gamma, \boldsymbol{\zeta} \sim \mathcal{N}_{M^p}(\mathbf{V}_2^{-1} \mathbf{v}_2; \mathbf{V}_2^{-1}).$$

where $\mathbf{V}_2 = \tau \mathbf{B}^\top \mathbf{B} + \mathbf{V}_1$ and $\mathbf{v}_2 = \tau \mathbf{B}^\top \boldsymbol{\zeta} + \mathbf{v}_1$. For the precision of measurement, the conditional posterior distribution is a gamma distribution:

$$\tau | \mathbf{c}, \boldsymbol{\zeta} \sim \mathcal{G}\left(\frac{N}{2} + a_\tau; \frac{\|\boldsymbol{\zeta} - \mathbf{B}\mathbf{c}\|^2}{2} + b_\tau\right).$$

The joint posterior distribution can be marginalized with respect to the spline coefficients. The log of the marginalized posterior distribution can be shown to be:

$$\begin{aligned}
\log(p(\boldsymbol{\theta}, \gamma, \tau | \boldsymbol{\zeta})) &= \frac{N}{2} \log(\tau) - \frac{\tau}{2} \boldsymbol{\zeta}^\top \boldsymbol{\zeta} \\
&+ \frac{1}{2} \log(\det(V_1)) - \frac{1}{2} \mathbf{v}_1^\top \mathbf{V}_1^{-1} \mathbf{v}_1 \\
&- \frac{1}{2} \log(\det(V_2)) + \frac{1}{2} \mathbf{v}_2^\top \mathbf{V}_2^{-1} \mathbf{v}_2 \\
&+ (a_\tau - 1) \log(\tau) - b_\tau \tau \\
&+ (a_\gamma - 1) \log(\gamma) - b_\gamma \gamma \\
&+ (a_\kappa - 1) \log(\kappa) - b_\kappa \kappa \\
&+ \log(p(\boldsymbol{\theta})).
\end{aligned}$$

C Elementary penalty elements

The penalty term, dealing with linear PDEs, can be seen as a homogeneous polynomial of second degree in the spline coefficients. The computations needed to construct the penalty term are:

$$PEN(\mathbf{c} | \boldsymbol{\theta}) = \mathbf{c}^\top \mathbf{R}(\boldsymbol{\theta}) \mathbf{c} + 2\mathbf{c}^\top \mathbf{r}(\boldsymbol{\theta}) + l(\boldsymbol{\theta}).$$

The \mathbf{R} matrix is given by the sum of a series of elementary penalty matrices:

$$\begin{aligned}
\mathbf{P}_{i_p, \dots, i_1}^{j_p, \dots, j_1} &= \int \left(\mathbf{B}_p^{(i_p)}(x_p) \otimes \dots \otimes \mathbf{B}_1^{(i_1)}(x_1) \right) \left(\mathbf{B}_p^{(j_p)}(x_p) \otimes \dots \otimes \mathbf{B}_1^{(j_1)}(x_1) \right)^\top dx_p \dots dx_1, \\
&= \int \mathbf{B}_p^{(i_p)}(x_p) \left(\mathbf{B}_p^{(j_p)}(x_p) \right)^\top dx_p \otimes \dots \otimes \int \mathbf{B}_1^{(i_1)}(x_1) \left(\mathbf{B}_1^{(j_1)}(x_1) \right)^\top dx_1, \\
&= \mathcal{S}_p^{(i_p, j_p)} \otimes \dots \otimes \mathcal{S}_1^{(i_1, j_1)},
\end{aligned}$$

where the matrix $\mathcal{S}_h^{(i_h, j_h)}$ can be approximated using a trapezoidal rule. The \mathbf{r} vector can be constructed as the sum of elementary penalty vectors:

$$\begin{aligned}
\mathbf{p}_{i_p, \dots, i_1} &= \int \mathbf{B}_p^{(i_p)}(x_p) \otimes \dots \otimes \mathbf{B}_1^{(i_1)}(x_1) dx_p \dots dx_1, \\
&= \int \mathbf{B}_p^{(i_p)}(x_p) dx_p \otimes \dots \otimes \int \mathbf{B}_1^{(i_1)}(x_1) dx_1, \\
&= \mathbf{s}_p^{(i_p)} \otimes \dots \otimes \mathbf{s}_1^{(i_1)},
\end{aligned}$$

as before, a trapezoidal rule can be applied to approximate the vector $\mathbf{s}_h^{(i_h)}$.