

Han Lin Shang* and Xibin Zhang

Bayesian bandwidth estimation for local linear fitting in nonparametric regression models

<https://doi.org/10.1515/snde-2018-0050>

Received May 24, 2018; accepted November 18, 2020; published online November 30, 2020

Abstract: This paper presents a Bayesian sampling approach to bandwidth estimation for the local linear estimator of the regression function in a nonparametric regression model. In the Bayesian sampling approach, the error density is approximated by a location-mixture density of Gaussian densities with means the individual errors and variance a constant parameter. This mixture density has the form of a kernel density estimator of errors and is referred to as the kernel-form error density (c.f. Zhang, X., M. L. King, and H. L. Shang. 2014. “A Sampling Algorithm for Bandwidth Estimation in a Nonparametric Regression Model with a Flexible Error Density.” *Computational Statistics & Data Analysis* 78: 218–34.). While (Zhang, X., M. L. King, and H. L. Shang. 2014. “A Sampling Algorithm for Bandwidth Estimation in a Nonparametric Regression Model with a Flexible Error Density.” *Computational Statistics & Data Analysis* 78: 218–34) use the local constant (also known as the Nadaraya-Watson) estimator to estimate the regression function, we extend this to the local linear estimator, which produces more accurate estimation. The proposed investigation is motivated by the lack of data-driven methods for simultaneously choosing bandwidths in the local linear estimator of the regression function and kernel-form error density. Treating bandwidths as parameters, we derive an approximate (pseudo) likelihood and a posterior. A simulation study shows that the proposed bandwidth estimation outperforms the rule-of-thumb and cross-validation methods under the criterion of integrated squared errors. The proposed bandwidth estimation method is validated through a nonparametric regression model involving firm ownership concentration, and a model involving state-price density estimation.

Keywords: kernel-form error density; Markov chain Monte Carlo; ownership concentration; state-price density.

JEL Classification: C53; C58.

1 Introduction

Local linear fitting is an important and widely used nonparametric tool for exploring an unknown relationship between a response variable and a group of explanatory variables. The local linear estimator reduces estimation error compared to the local constant estimator. While the choice of kernel function is not important, the performance of the local linear estimator is mainly determined by choice of bandwidth (see, e.g., Fan and Gijbels 1996, p. 76). Similar to a previous study by Zhang, King, and Shang (2014), we present a Bayesian sampling approach to bandwidth estimation for a nonparametric regression model with unknown error density. However, in contrast to Zhang, King, and Shang (2014, 2016), we consider local linear fitting rather than local constant fitting, to improve the estimation accuracy of the regression function.

*Corresponding author: Han Lin Shang, Department of Actuarial Studies and Business Analytics, Macquarie University, Sydney, New South Wales, Australia; and Research School of Finance, Actuarial Studies and Statistics, Australian National University, Canberra, Australian Capital Territory, Australia, E-mail: hanlin.shang@mq.edu.au. <https://orcid.org/0000-0003-1769-6430>
Xibin Zhang, Department of Econometrics and Business Statistics, Monash University, 900 Dandenong Road, Caulfield East, Victoria, Australia

A large body of literature discusses methods for bandwidth selection, such as the rule-of-thumb (ROT) and cross-validation (CV) discussed by Härdle (1990) and Scott (1992), the plug-in method discussed by Herrmann et al. (1995), and bootstrapping as proposed by Hall, Lahiri, and Polzehl (1995). In terms of the local constant estimator, a particular version of the local linear estimator, Zhang, King, and Shang (2014) present a simulation study, finding that CV outperforms ROT and bootstrapping. As evidenced by the simulation results presented in this paper, CV also outperforms ROT for local linear fitting. This finding partly explains why CV is often used in a wide range of application studies. However, even if CV is asymptotically unbiased, it has a relatively large finite-sample bias. Faraway and Jhun (1990), Padgett and Thombs (1986) and Romano (1988) have established the same behavior studying nonparametric estimates for the density, quantiles and the mode, respectively. This problem arises in nonparametric-based models, as exemplified by Härdle and Mammen (1993). In addition, CV sometimes fails to produce meaningful results. Li and Zhou (2005) show that in terms of the local constant estimator with one regressor, a necessary and sufficient condition has to be met by the CV to guarantee the existence of a unique optimal bandwidth. The conditions required to ensure the uniqueness of a cross-validated bandwidth vector become complicated where they are multiple regressors and/or local linear fitting. Moreover, in application studies, it is difficult to verify such conditions. Therefore, the lack of an applicable data-driven bandwidth selection method for local linear fitting motivates the investigation carried out in this paper.

Existing bandwidth selection methods do not take error density into account. Although these methods do not require any information about the error distribution, the distribution of the response variable is often of great interest to error-based techniques, such as value-at-risk estimation (e.g., Martins-Filho, Yao, and Torero 2018). Such a distribution is characterized by the error density, estimation of which is fundamental in statistical inference for any regression model. A straightforward approach to error density estimation is to use residuals as proxies for errors and derive a kernel estimator. However, bandwidth selection is yet to be investigated when the regression function is estimated through local linear fitting. Thus, the proposed investigation is also motivated by the lack of a data-driven procedure for simultaneously choosing bandwidths in the local linear estimator and the kernel estimator of error density.

Let y denote the response and $\mathbf{x} = (x_1, x_2, \dots, x_d)^\top$ a set of explanatory variables. Given observations (y_i, \mathbf{x}_i) , the multivariate nonparametric regression model is expressed as

$$y_i = m(\mathbf{x}_i) + \epsilon_i, \quad i = 1, 2, \dots, n,$$

where $(\epsilon_1, \epsilon_2, \dots, \epsilon_n)$ are assumed to be independent and identically distributed (iid) with an unknown error density, denoted as $f(\epsilon)$. It is also assumed that the errors are independent of the explanatory variables.

The regression function $m(\mathbf{x})$ can be estimated by the local linear estimator denoted by $\hat{m}(\mathbf{x}, \mathbf{h})$, where $\mathbf{h} = (h_1, h_2, \dots, h_d)^\top$ is a vector of bandwidths corresponding to the d explanatory variables. We propose to approximate the unknown error density by a mixture density given by

$$f(\epsilon; b) = \frac{1}{n} \sum_{i=1}^n \frac{1}{b} \phi\left(\frac{\epsilon - \epsilon_i}{b}\right), \quad (1)$$

where $\phi(\cdot)$ is the probability density of the standard Gaussian distribution. It is a location-mixture density with Gaussian components having a common variance b^2 and individual means at the errors. The validity of this mixture density as an error density has been investigated in Zhang, King, and Shang (2014) for the local constant estimator in nonparametric regression models. As $f(\epsilon; b)$ has the form of a kernel density estimator of errors, it can approximate the true error density very well when the sample size is sufficiently large.

There is a growing literature on the estimation of the error density in a nonparametric regression model. Efromovich (2005) presented the so-called Efromovich–Pinsker estimator of the error density and show that this estimator is asymptotically as accurate as an oracle that knows the underlying errors. Cheng (2004) demonstrates that the kernel density estimator of residuals is uniformly, weakly, and strongly consistent.

While the regression function is estimated by the Nadaraya–Watson estimator, the error density is estimated by the kernel estimator of residuals. Samb (2010, 2011) proved the asymptotic normality of the bandwidths in both estimators, deriving the optimal convergence rates of the two bandwidths. Linton and Xiao (2007) propose a kernel estimator based on residuals obtained through the local polynomial fitting of the unknown regression function. They showed that the estimator is adaptive and conclude that adaptive estimation is possible in the local polynomial fitting. In a class of nonlinear regression models, Yuan and de Gooijer (2007) constructed an approximate (pseudo) likelihood through the kernel density estimator of pre-fitted residuals with its bandwidth pre-chosen by the rule-of-thumb. They prove that under some regularity conditions, the resulting maximum likelihood estimates of parameters are consistent, asymptotically normal, and efficient. Jaki and West (2008, 2011) propose using the kernel density estimator of the pre-fitted residuals to construct an approximate likelihood, which they call “kernel likelihood”. In all these papers, residuals are commonly used as proxies of errors, and the bandwidth for the kernel density estimator of residuals is pre-chosen or chosen in a way that is unrelated to the bandwidth in the regression function. To our knowledge, no method can simultaneously estimate the bandwidths for the local linear estimator of the regression function and kernel-form error density.

The main contribution of this paper is to construct an approximate likelihood and, therefore, the posterior of the bandwidth parameters for the nonparametric regression model with the kernel-form error density given by (1), where the unknown regression function is estimated through the local linear fitting. In contrast to Zhang, King, and Shang (2014, 2016), the local linear fitting improves the estimation accuracy of the regression function. A Bayesian sampling algorithm is presented to sample the bandwidth parameters from their posterior; thus, the two types of bandwidths can be estimated simultaneously. We conduct a Monte Carlo simulation study, which shows that our proposed method outperforms the ROT and CV methods where there are multiple regressors under the criterion of integrated squared errors (ISE). The proposed bandwidth estimation method is also validated through a nonparametric regression model involving firm ownership concentration, and a model involving state-price density (SPD) estimation.

The remainder of this paper is organized as follows. In Section 2, we present the likelihood and posterior of bandwidth parameters based on local linear fitting. Section 3 presents Monte Carlo simulation studies to examine the accuracy of the proposed bandwidth estimation method. The proposed Bayesian method is validated through a nonparametric regression model involving firm ownership concentration in Section 4. In Section 5, we apply our method to bandwidth estimation in a nonparametric regression model involved in SPD estimation. Section 6 concludes the paper.

2 Bayesian bandwidth estimation

The bandwidths in the local linear estimator of the regression function and kernel-form error density estimator play an important role in controlling the smoothness of the regression function and error density estimator, respectively. In the case of the local constant estimator, bandwidths are treated as parameters by Zhang, Brooks, and King (2009), Zhang, King, and Shang (2014).

In this paper, the error density of the nonparametric regression model given by (1) is approximated by the kernel-form density $\phi(\cdot)$ in (1). When the error density of a regression model is non-Gaussian, a scale-mixture density of several Gaussian components is often used, where each Gaussian density may have mean zero and different variance parameters. However, the use of a scale-mixture density of Gaussian components is at the cost of dramatically increasing the number of parameters. In contrast, the kernel-form error density has the advantage of having only one parameter, the common variance parameter, which also acts as the bandwidth for this kernel-form density.

2.1 Cross-validation

2.1.1 Consider the nonparametric regression model given by

$$y_j = m(\mathbf{x}_j) + \epsilon_j, \quad j = 1, 2, \dots, n, \quad (2)$$

where $\mathbf{x}_j = (x_{j,1}, \dots, x_{j,d})$ is a random vector of dimension d . Define the derivative of $m(\mathbf{x})$ as $\beta(\mathbf{x}) = \partial m(\mathbf{x})/\partial \mathbf{x}$.

Let $\delta(\mathbf{x}) = (m(\mathbf{x}), \beta(\mathbf{x})^\top)^\top$ be a vector-valued function with dimension $d+1$. Taking a Taylor series expansion of $m(\mathbf{x}_j)$ at \mathbf{x}_i , we obtain

$$m(\mathbf{x}_j) = m(\mathbf{x}_i) + (\mathbf{x}_j - \mathbf{x}_i)^\top \beta(\mathbf{x}_i) + o_{ij}, \quad (3)$$

where o_{ij} represents the remaining term of the expansion. Substituting (3) into (2), we obtain

$$y_j = \left(1, (\mathbf{x}_j - \mathbf{x}_i)^\top\right) \delta(\mathbf{x}_i) + o_{ij} + \epsilon_j.$$

As pointed out by Li and Racine (2004), the leave-one-out local linear estimator of $\delta(\mathbf{x}_i) = (m(\mathbf{x}_i), \beta(\mathbf{x}_i)^\top)^\top$ is obtained by a kernel weighted regression of y_j on $(1, (\mathbf{x}_j - \mathbf{x}_i)^\top)$ given by

$$\widehat{\delta}_{-j}(\mathbf{x}_j; \mathbf{h}) = \left[\sum_{\substack{i=1 \\ i \neq j}}^n K_{\mathbf{h}}(\mathbf{x}_j; \mathbf{x}_i) \begin{pmatrix} 1 & (\mathbf{x}_j - \mathbf{x}_i)^\top \\ \mathbf{x}_j - \mathbf{x}_i & (\mathbf{x}_j - \mathbf{x}_i)(\mathbf{x}_j - \mathbf{x}_i)^\top \end{pmatrix} \right]^{-1} \sum_{\substack{i=1 \\ i \neq j}}^n K_{\mathbf{h}}(\mathbf{x}_j; \mathbf{x}_i) \begin{pmatrix} 1 \\ \mathbf{x}_j - \mathbf{x}_i \end{pmatrix} y_i,$$

for $j = 1, 2, \dots, n$, where $K_{\mathbf{h}}(\mathbf{x}_j; \mathbf{x}_i) = \prod_{k=1}^d h_k^{-1} K((x_{j,k} - x_{i,k})/h_k)$ with h_k being the bandwidth corresponding to the k th component of \mathbf{x} . Note that $x_{j,k}$ is the k th element of \mathbf{x}_j , for $k = 1, 2, \dots, d$.

The leave-one-out estimator of $\widehat{m}(\mathbf{x}_j)$ is given by

$$\widehat{m}_{-j}(\mathbf{x}_j; \mathbf{h}) = \mathbf{I}_1^\top \widehat{\delta}_{-j}(\mathbf{x}_j; \mathbf{h}),$$

where \mathbf{I}_1 is a $(d+1) \times 1$ vector, the first element of which is one and other elements zero. An optimal \mathbf{h} can be chosen by minimizing the CV function given by

$$\text{CV}(\mathbf{h}) = \sum_{j=1}^n [y_j - \widehat{m}_{-j}(\mathbf{x}_j; \mathbf{h})]^2. \quad (4)$$

The CV method is considered the benchmark method for comparison with our proposed method (to follow).

2.2 Likelihood

Zhang, King, and Shang (2014) propose approximation of the unknown error density by a location-mixture density of n Gaussian densities, where the component Gaussian densities have a common variance and different means at individual errors. The errors were approximated by the residuals from a nonparametric regression model with the regression function estimated by the local constant estimator. In this paper, we follow the same idea for the nonparametric model with a local linear estimator of the regression function, where the error density is assumed to be approximated by (1), and, therefore, the density of y_i is approximated by $f(\{y_i - \widehat{m}_{-i}(\mathbf{x}_i; \mathbf{h})\}; b)$ expressed as

$$\widehat{f}(\{y_i - \widehat{m}_{-i}(\mathbf{x}_i; \mathbf{h})\}; b) = \frac{1}{n} \sum_{j=1}^n \frac{1}{b} \phi\left(\frac{\{y_i - \widehat{m}_{-i}(\mathbf{x}_i; \mathbf{h})\} - \{y_j - \widehat{m}_{-j}(\mathbf{x}_j; \mathbf{h})\}}{b}\right), \quad (5)$$

for $i = 1, 2, \dots, n$. It is likely that $y_j - \widehat{m}_{-j}(\mathbf{x}_j; \mathbf{h}) = y_i - \widehat{m}_{-i}(\mathbf{x}_i; \mathbf{h})$, for some j , resulting in an unwanted term $\phi(0)/b$. A remedy is to exclude the j th term that makes $y_j - \widehat{m}_{-j}(\mathbf{x}_j; \mathbf{h}) = y_i - \widehat{m}_{-i}(\mathbf{x}_i; \mathbf{h})$ from the summation in (5). Therefore, the density of y_i is approximated by

$$\hat{f}(\{y_i - \hat{m}_{-i}(\mathbf{x}_i; \mathbf{h})\}; b) = \frac{1}{n - n_i} \sum_{j \in J_i} \frac{1}{b} \phi \left(\frac{\{y_i - \hat{m}_{-i}(\mathbf{x}_i; \mathbf{h})\} - \{y_j - \hat{m}_{-j}(\mathbf{x}_j; \mathbf{h})\}}{b} \right), \quad (6)$$

where $J_i = \{j : y_j - \hat{m}_{-j}(\mathbf{x}_j; \mathbf{h}) \neq y_i - \hat{m}_{-i}(\mathbf{x}_i; \mathbf{h}), \text{ for } j = 1, 2, \dots, n\}$, and n_i is the number of terms excluded from the summation in (5).

Given \mathbf{h}^2 and b^2 , the likelihood of $\mathbf{y} = (y_1, y_2, \dots, y_n)^\top$ is the product of the density of y_i , for $i = 1, 2, \dots, n$, which are approximated by (6). The likelihood is expressed explicitly as

$$L(\mathbf{y} | \mathbf{h}^2, b^2) = \prod_{i=1}^n \left\{ \frac{1}{n - n_i} \sum_{j \in J_i} \frac{1}{b} \phi \left(\frac{\{y_i - \hat{m}_{-i}(\mathbf{x}_i; \mathbf{h})\} - \{y_j - \hat{m}_{-j}(\mathbf{x}_j; \mathbf{h})\}}{b} \right) \right\}.$$

This likelihood function is not a proper likelihood since some terms are left out. Instead, the likelihood is a pseudo-likelihood. The pseudo-likelihood is a likelihood function associated with a family of probability distributions, which does not necessarily contain the true distribution (Gourieroux, Monfort, and Trognon 1984). As a consequence, the resulting Bayesian estimators, while consistent, may have an inaccurate posterior variance, and subsequent credible sets constructed by this posterior may not be asymptotically valid. Note that similar issues are discussed by, for example, Chernozhukov and Hong (2003) and Müller (2013). There is no satisfactory solution to such a practical problem because on the one hand, dropping the terms $i = j$ invalidates posterior inference, and on the other hand, keeping all terms in likelihood would result in corner solutions for the bandwidths.

Nonetheless, the dropped terms may have a negligible impact on the resulting likelihood. If the dropped terms were included back to (6), they would contribute to the would-be likelihood by $\phi(0)/b$, which is a constant and is unaffected by observations. This constant has no impact on the would-be likelihood when \mathbf{h} is considered. Only when b is considered, will the dropped terms have an impact on the would-be likelihood. However, this might be regarded as “a cost” to avoid any possible corner solution for b . Any possible theoretical justification on this issue is left for future research.

2.3 Priors

As studied by Nadaraya (1964), Watson (1964) and Härdle and Marron (1985), for consistency of the nonparametric estimator, the bandwidth parameters h_k and b should go to infinity, as the sample size n converges to infinity. In most cases, the optimal bandwidth rate can easily be derived up to a constant factor. We acknowledge that inverse Gamma prior density with constant hyperparameters may not always be appropriate for different sample sizes. As such, we choose prior densities by re-expressing bandwidths as the product of an unknown constant and a known order convergence rate:

$$\begin{aligned} b &= b_0 \times n^{-1/5}, \\ h_k &= h_{0,k} \times n^{-1/(d+4)}, \quad k = 1, \dots, d, \end{aligned}$$

where b_0 and $h_{0,k}$ do not depend on sample size n . Therefore, bandwidth estimation is equivalent to the estimation of b_0 and $h_{0,k}$. The priors of b_0^2 and $h_{0,k}^2$ can be chosen as the inverse Gamma prior with constant hyperparameters (see, e.g., Zhang, King, and Hyndman 2006, 2009, for discussion of other prior choices). Thus, the priors of b_0^2 and $h_{0,k}^2$ are given by

$$\begin{aligned} \pi(b_0^2) &= \frac{(\beta_b)^{\alpha_b}}{\Gamma(\alpha_b)} \left(\frac{1}{b_0^2} \right)^{\alpha_b+1} \exp \left\{ -\frac{\beta_b}{b_0^2} \right\}, \\ \pi(h_{0,k}^2) &= \frac{(\beta_h)^{\alpha_h}}{\Gamma(\alpha_h)} \left(\frac{1}{h_{0,k}^2} \right)^{\alpha_h+1} \exp \left\{ -\frac{\beta_h}{h_{0,k}^2} \right\}, \end{aligned}$$

respectively, where (α_b, β_b) and (α_h, β_h) are hyperparameters. The hyperparameters were set at $\alpha_b = \alpha_h = 1$ and $\beta_b = \beta_h = 0.05$. These values are often used as the parameter values of an inverse Gamma density chosen as the prior of a variance parameter (see, e.g., Geweke 2010).

As pointed out by a reviewer, we may also consider bandwidths that are of the form $b = n^{-\omega}$ and $h = n^{-\eta}$, and choose to work with beta priors on the parameters $(\omega, \eta) \in (0, 1)$. We assume that the priors of ω and η are the beta density with hyperparameters denoted by $\text{Beta}(\psi_b, \kappa_b)$ and $\text{Beta}(\psi_{h_k}, \kappa_{h_k})$, respectively. Thus, the priors of b and h_k are given by

$$b = n^{-\omega}, \quad \pi_{\omega}(\psi_b, \kappa_b) = \frac{\omega^{\psi_b-1} (1-\omega)^{\kappa_b-1}}{B(\psi_b, \kappa_b)}, \quad B(\psi_b, \kappa_b) = \frac{\Gamma(\psi_b)\Gamma(\kappa_b)}{\Gamma(\psi_b + \kappa_b)}$$

$$h_k = n^{-\eta_k}, \quad \pi_{\eta_k}(\psi_{h_k}, \kappa_{h_k}) = \frac{\eta_k^{\psi_{h_k}-1} (1-\eta_k)^{\kappa_{h_k}-1}}{B(\psi_{h_k}, \kappa_{h_k})}, \quad B(\psi_{h_k}, \kappa_{h_k}) = \frac{\Gamma(\psi_{h_k})\Gamma(\kappa_{h_k})}{\Gamma(\psi_{h_k} + \kappa_{h_k})},$$

respectively, where (ψ_b, κ_b) and $(\psi_{h_k}, \kappa_{h_k})$ are hyperparameters.

In simulation studies to be presented in Section 3, we choose our prior density to be the inverse Gamma prior density for b_0 and $h_{0,k}$. Because of the time-consuming nature of our simulation, we tend to leave the other prior choice for future research.

2.4 Posterior

According to Bayes' theorem, the posterior of b^2 and \mathbf{h}^2 is approximated by (up to a normalizing constant)

$$\pi(b^2, \mathbf{h}^2 | \mathbf{y}) \propto L(\mathbf{y} | b^2, \mathbf{h}^2) \pi(b^2) \prod_{k=1}^d \pi(h_k^2),$$

from which we use an adaptive random-walk Metropolis algorithm of Garthwaite, Fan, and Sisson (2016) to sample b^2 and the elements of \mathbf{h}^2 . As studied by Gelman, Roberts, and Gilks (1996), Roberts, Gelman, and Gilks (1997), Roberts and Rosenthal (2001), Roberts and Rosenthal (2009) and Lee and Neal (2018), the acceptance rate in the adaptive random-walk Metropolis algorithm was controlled to be around 0.234 for multivariate parameters (i.e., bandwidths of predictors in the regression mean) and 0.440 for a univariate parameter (i.e., bandwidth in the error density). As Markov chain Monte Carlo iterations increase, the adaptive algorithm can automatically update the step size in a random-walk Metropolis algorithm, to achieve the optimal acceptance rates. If the proposal parameters are accepted in an iteration, then the step size is reduced in the next iteration to search the optimal parameters in greater detail. If the proposal parameters are rejected in an iteration, then the step size is increased in the next iteration to explore other possible values. The adaptive random-walk Metropolis algorithm is computationally fast with guaranteed acceptance rates, in comparison to a fixed step size in the classical Metropolis algorithm.

When the sampling procedure is completed, the ergodic average of each simulated chain is used as an estimate of each element of \mathbf{h}^2 and b^2 . Once bandwidths are estimated, the kernel-form error density can be derived according to (1).

3 Monte Carlo simulation studies

The purpose of this simulation study is to examine the performance of the proposed Bayesian bandwidth estimation method by assessing the accuracy of the resulting local linear estimator and kernel density estimator. For comparison purposes, we also choose bandwidths through the ROT and CV methods described in Härdle (1990) and Scott (1992). With bandwidths for the local linear estimator estimated through these three methods, we calculate the residuals, based on which we apply the ROT and likelihood CV methods to obtain

the bandwidth for the kernel density estimator of errors. This is a two-step procedure, whereas the proposed Bayesian algorithm selects the bandwidths of the local linear estimator and kernel-form error density simultaneously.

3.1 Simulation setup

Consider the relationships between y and $\mathbf{x} = (x_1, x_2, x_3)^\top$ given by

$$m_1(\mathbf{x}) = \cos(2\pi x_1) + \sin(2\pi x_1), \quad (7)$$

$$m_2(\mathbf{x}) = \sin(2\pi x_1) + \cos(2\pi x_2) + 4(1 - x_3^2). \quad (8)$$

A sample was generated by drawing x_1, x_2, x_3 independently from the uniform density on $(0, 1)$. The errors were simulated from either $N(0, 0.5^2)$ or a mixture of two Gaussian densities defined as $0.7N(0, 0.4^2) + 0.3N(0, 0.8^2)$. The response was calculated through (7) and (8) with the generated errors added. We consider three sample sizes: $n = 200, 1000$ and 5000 .

The relationships between the responses and explanatory variables were modeled by nonparametric regression models given (respectively) by

$$y_j = m(x_{j,1}) + \epsilon_j, \quad j = 1, 2, \dots, n, \quad (9)$$

$$y_j = m(x_{j,1}, x_{j,2}, x_{j,3}) + \epsilon_j, \quad (10)$$

where the errors are assumed to be iid with an unknown density.

3.2 Simulation result

The error density of (9) and (10) is assumed to be approximated by the kernel-form density given by (1). With each pseudo sample generated from each model, we implemented the proposed Bayesian sampling algorithm to estimate bandwidths for local linear estimators of (9) and (10). The random-walk Metropolis algorithm was used to sample parameters, where the burn-in period contains the first 1000 draws, and the successive 10,000 draws were recorded.

The posterior averages of bandwidths for each model are presented in Table 1. The mixing performance of this posterior sampler is examined by the simulation inefficiency factor (SIF), which can be loosely interpreted as the number of draws needed to obtain independent draws from the simulated chain. For example, a SIF value of 20 indicates that, approximately, we should keep one draw for every 20 draws to derive independent draws (see, e.g., Kim, Shephard, and Chib 1998; Roberts 1996).

The standard deviation of the posterior mean is approximated by the batch-mean standard deviation. It becomes smaller and smaller as the number of iterations increases and the sampler achieves a reasonable mixing performance. We used the SIF and batch-mean standard deviation to monitor the mixing performance. Table 1 presents the values of these two indicators, which show that the sampler has mixed very well.

3.3 Comparison of bandwidth estimation methods

The accuracy of estimated bandwidth parameters is compared via the ISE of the resulting local linear estimator. The ISE is defined as

$$\text{ISE}(\hat{m}) = \int_{a_0}^{a_1} [\hat{m}(x; \mathbf{h}) - m(x)]^2 dx,$$

Table 1: Posterior estimates and statistics for Bayesian bandwidth estimation of one and three regressors under (mixture) Gaussian error density with $n = 1000$ in each sample.

Error density	Parameter estimate	95% Bayesian credible interval	Standard deviation	Batch-mean standard dev	SIF
$d = 1$					
Gaussian	$b = 0.2698$	(0.1858, 0.3689)	0.0474	0.0010	4.4
	$h_1 = 0.0607$	(0.0514, 0.0706)	0.0049	0.0001	4.5
Mixture Gaussian	$b = 0.2958$	(0.2196, 0.3835)	0.0423	0.0009	4.3
	$h_1 = 0.0615$	(0.0521, 0.0717)	0.0050	0.0001	4.4
$d = 3$					
Gaussian	$b = 0.2653$	(0.1790, 0.3660)	0.0485	0.0011	5.4
	$h_1 = 0.0784$	(0.0675, 0.0901)	0.0058	0.0002	12.6
	$h_2 = 0.1608$	(0.1308, 0.1930)	0.0160	0.0006	15.0
	$h_3 = 0.0924$	(0.0778, 0.1082)	0.0078	0.0003	13.5
Mixture Gaussian	$b = 0.2957$	(0.2156, 0.3869)	0.0442	0.0010	5.2
	$h_1 = 0.0802$	(0.0688, 0.0923)	0.0060	0.0002	12.3
	$h_2 = 0.1659$	(0.1348, 0.1994)	0.0166	0.0006	14.8
	$h_3 = 0.0951$	(0.0800, 0.1114)	0.0081	0.0003	13.3

where $[a_0, a_1]$ is the support of $m(x)$. Given a set of grids, $a_0 = x_1 < x_2 < \dots < x_{1000} = a_1$, we approximate the ISE by

$$\tilde{\text{ISE}}(\hat{m}) = \frac{a_1 - a_0}{1000} \sum_{j=1}^{1000} [\hat{m}(x_j) - m(x_j)]^2.$$

The CV method aims to derive an optimal \mathbf{h} that minimizes ISE (see, e.g., Härdle 1991; Sain, Baggerly, and Scott 1994). However, the difficulty in numerically deriving an optimal bandwidth increases with the dimension.

Another criterion for bandwidth selection is the mean ISE (MISE) given by

$$\text{MISE}(\hat{m}) = \mathbb{E} \left\{ \int_{a_0}^{a_1} [\hat{m}(x; \mathbf{h}) - m(x)]^2 dx \right\},$$

where $\mathbb{E}\{\cdot\}$ denotes expectation. It is well known that the optimal bandwidth that minimizes MISE does not have a closed-form (Zhang, Brooks, and King 2009); in practice, it is usually approximated asymptotically. When data are observed from the multivariate normal density, and the kernel function is the standard Gaussian density, the optimal bandwidth that minimizes MISE can be approximated by

$$h_k = \sigma_k \left[\frac{4}{(d+2)n} \right]^{1/(d+4)},$$

where $\sigma_k = \min\{s_k, q_k/1.34\}$ with s_k the sample standard deviation and q_k the interquartile range for the k th predictor, for $k = 1, 2, \dots, d$. This method is often used in practice because of its simplicity and a lack of alternatives (see, e.g., Bowman and Azzalini 1997; Scott 1992; Silverman 1986). However, we might be misled by the bandwidths selected through ROT when the true density is far from Gaussian distribution. The CV method given in (4) has been extensively discussed for bandwidth selection (see, e.g., Härdle and Müller 2000; Härdle and Marron 1985; Wahba and Wold 1975). In this simulation study, we consider the ROT and CV as competing methods for the Bayesian method.

We calculate the ISE of the local linear estimator of the regression function with bandwidths estimated by the ROT, CV and Bayesian methods. The boxplots of ISE values are presented in Figure 1. When $d = 1$, the CV has the best performance, with MISE performing worse than the other methods considered. However, as d increases to three, the proposed Bayesian method performs the best.




Figure 1: Boxplots of 1000 ISE derived through each bandwidth estimation method over $n = 200, 1000$ and 5000 generated samples for estimating the regression function, where errors were simulated from the Gaussian density $N(0, 0.5^2)$ and mixture Gaussian density $0.7N(0, 0.4^2) + 0.3N(0, 0.8^2)$ for both one and three regressors.

After estimating the regression function, we apply the ROT and likelihood CV to obtain the bandwidth of the estimated error density via a two-step procedure in which the ROT and CV methods are applied twice separately. In contrast, our Bayesian algorithm simultaneously estimates bandwidths of regression function and kernel-form error density. Performance of the resulting bandwidth estimation for the kernel-form error density is summarized in Figure 2. When the errors are generated from the Gaussian density and mixture Gaussian density, the distribution of ISE values derived through the Bayesian method is generally smaller than those derived from either the ROT or CV method.

3.4 Comparison between the local constant and local linear estimators

We compare the proposed Bayesian bandwidth estimation method with the local linear estimator to the Bayesian bandwidth estimation with the local constant estimator introduced by Zhang, King, and Shang (2014). Although the two methods are both Bayesian, they differ in terms of the regression mean estimator. In Figure 3, we display boxplots of 1000 ISE values estimated by the local constant and local linear estimators. The local linear estimator improves the estimation accuracy of the regression mean function while maintaining similar values of ISE for the estimation of error density.

Further, model selection in Bayesian inference is conducted through the marginal likelihood, which is the expectation of the likelihood with respect to the prior parameters. Here, we use the marginal likelihood to compare the two estimators and assess their statistical significance. The integral of the product of the likelihood and prior of parameters is rarely calculated analytically, but it is often computed numerically (see, e.g., Chib 1995; Geweke 1999).

Let $\theta = (\mathbf{h}, b)$ denote the parameter vector and \mathbf{y} the data. The marginal likelihood reflects a summary of evidence provided by the data supporting one estimator as opposed to its competing estimator. Chib (1995) shows that the marginal likelihood under an estimator \mathcal{A} is expressed as

$$P_{\mathcal{A}}(\mathbf{y}) = \frac{l_{\mathcal{A}}(\mathbf{y}|\theta)\pi_{\mathcal{A}}(\theta)}{\pi_{\mathcal{A}}(\theta|\mathbf{y})}, \tag{11}$$




Figure 2: Boxplots of 1000 ISE derived through each bandwidth estimation method over $n = 200, 1000$ and 5000 generated samples for estimating the error density, where errors were simulated from the Gaussian density $N(0, 0.5^2)$ and mixture Gaussian density $0.7N(0, 0.4^2) + 0.3N(0, 0.8^2)$ for both one and three regressors.




Figure 3: Boxplots of 1000 ISE derived through the local constant and local linear estimators using the Bayesian bandwidth estimation method over $n = 1000$ generated samples for estimating the regression function and error density, where errors were simulated from the Gaussian density $N(0, 0.5^2)$ and mixture Gaussian density $0.7N(0, 0.4^2) + 0.3N(0, 0.8^2)$ for both one and three regressors.

where $l_{\mathcal{A}}(\mathbf{y}|\theta), \pi_{\mathcal{A}}(\theta)$ and $\pi_{\mathcal{A}}(\theta|\mathbf{y})$ denote likelihood, prior and posterior under estimator \mathcal{A} . $P_{\mathcal{A}}(\mathbf{y})$ is computed at the posterior estimate of θ . The numerator of (11) has a closed form and can be computed analytically, the denominator of (11) is the posterior of θ , which is replaced by its kernel density estimator from the simulated chain of θ through a posterior sampler. The Bayes factor of estimator \mathcal{A} against estimator \mathcal{B} is defined as

Table 2: Mean and standard deviation (sd) of 1000 LML derived through the Bayesian bandwidth estimation method with the local constant and local linear estimators averaged over 1000 generated samples for estimating the error density, where errors were simulated from the Gaussian density $N(0, 0.5^2)$ and mixture Gaussian density $0.7N(0, 0.4^2) + 0.3N(0, 0.8^2)$ for both one and three regressors. The minimum values are highlighted in bold. Chib denotes the LML computed by Chib (1995), while Geweke denotes the LML computed by Geweke (1999).

	<i>d</i> = 1				<i>d</i> = 3			
	Gaussian		Mixture Gaussian		Gaussian		Mixture Gaussian	
Estimator	Chib	Geweke	Chib	Geweke	Chib	Geweke	Chib	Geweke
LML (local constant)	-763.27	-763.35	-837.75	-837.86	-859.23	-859.56	-932.46	-932.81
LML (local linear)	-754.99	-755.05	-829.86	-829.94	-806.91	-807.17	-883.64	-883.92
Bayes factor	3944	4024	2670	2752	5.28e+22	5.66e+22	1.59e+21	1.71e+21

$$BF = \frac{P_{\mathcal{A}}(\mathbf{y})}{P_{\mathcal{B}}(\mathbf{y})} = \exp\{\ln P_{\mathcal{A}}(\mathbf{y}) - \ln P_{\mathcal{B}}(\mathbf{y})\},$$

where $P_{\mathcal{A}}(\mathbf{y}) > P_{\mathcal{B}}(\mathbf{y})$. The Bayes factor is used to decide whether estimator \mathcal{A} is favored over estimator \mathcal{B} . According to Kass and Raftery (1995), a Bayes factor value between 1 and 3 indicates that the evidence supporting estimator \mathcal{A} against estimator \mathcal{B} is insignificant. When the Bayes factor is between 3 and 20, estimator \mathcal{A} is favored over \mathcal{B} with positive evidence; when the Bayes factor is between 20 and 150, \mathcal{A} is favored over \mathcal{B} with strong evidence; and when the Bayes factor is above 150, \mathcal{A} is favored over \mathcal{B} with very strong evidence. In Table 2, based on 1000 repetitions, we compare the log marginal likelihood (LML) and Bayes factor between the local constant and local linear estimator and find a strong preference for the local linear estimator.

4 An application to firm ownership concentration

Yafeh and Yosha (2003) present an empirical investigation on how expenditure on managerial activities inside a firm is affected by a firm’s ownership concentration. Their data were collected from $n = 185$ firms listed on the Japanese stock market. Using this data set, Kuruwita, Kulasekera, and Gallagher (2011) investigate the relationship between a measure of expenses on managerial private benefits (y) and three explanatory variables – firm age (x_1), ownership concentration (x_2) and firm profit (x_3) – to justify their proposed varying-coefficient regression model. Note that the ownership concentration is represented by the percentage of shares possessed by the top 10 shareholders. Kuruwita, Kulasekera, and Gallagher (2011) find that all three explanatory variables have a significant impact on the response.

To illustrate our proposed method for bandwidth estimation, we considered the nonparametric regression model given by

$$y_i = m(x_{i,1}, x_{i,2}, x_{i,3}) + \epsilon_i, \quad \text{for } i = 1, 2, \dots, 185,$$

where the error density is unknown and approximated by our proposed location-mixture density given by (6).

To apply our proposed sampling algorithm, we assume that the priors of b^2 and h_k^2 , for $k = 1, 2, 3$, are the inverse gamma densities. As a sensitivity analysis, we also consider prior densities as the exponential densities, given by

$$\pi(z) = \tau \exp(-\tau z), \quad z \geq 0,$$

where the hyperparameter $\tau = 1$. The estimated values of b , h_k and their associated statistics are presented in Table 3. The batch-mean standard deviation and SIF values show that the sampler has reached good mixing performance.

Table 3: Parameter estimates and statistics of Bayesian bandwidth estimation under mixture Gaussian error density for Japanese chemical industry data. The burn-in period contains the first 1000 draws, and the successive 10,000 draws were recorded.

Bandwidth estimate	95% Bayesian credible interval	Standard deviation	Batch-mean standard dev	SIF
Inverse gamma densities ($\alpha = 1, \beta = 0.05$)				
$b = 0.3026$	(0.1998, 0.4304)	0.0574	0.0010	5.7
$h_1 = 4.6976$	(4.0505, 5.4587)	0.3489	0.0077	9.7
$h_2 = 3.4029$	(2.4674, 4.1056)	0.4087	0.0098	11.6
$h_3 = 0.9467$	(0.2345, 1.9318)	0.4485	0.0175	30.3
Exponential prior densities ($\tau = 1$)				
$b = 0.3049$	(0.1993, 0.4344)	0.0585	0.0010	5.6
$h_1 = 4.6856$	(3.9914, 5.3569)	0.3498	0.0079	10.1
$h_2 = 3.3958$	(2.3863, 4.1494)	0.4195	0.0103	12.0
$h_3 = 0.9589$	(0.2818, 1.9872)	0.4451	0.0173	30.0

4.1 In-sample fitting

For comparison purposes, we compare the accuracy of bandwidth estimation for the local linear estimator obtained by the ROT and CV methods. By using the *np* package (Hayfield and Racine 2008) in R (R Core Team 2020), the resulting bandwidth vector of the local linear estimator is $(0.6288 \times 10^8, 0.1315 \times 10^{12}, 0.3184 \times 10^{13})$, which is not a meaningful result. Consequently, the CV method fails to choose meaningful bandwidths. In contrast, the ROT and Bayesian methods both lead to reasonable bandwidth estimates, with bandwidths obtained by the ROT being $\mathbf{h} = (6.1720, 5.9416, 0.0233)^\top$.

Applying the ROT method to the kernel density estimator based on residuals, we obtained a bandwidth of $b = 3.68$, which differs from the bandwidth derived through the proposed Bayesian method. The resulting difference in kernel density estimators can be seen in Figure 4.

4.2 Out-of-sample prediction

To assess the out-of-sample point forecast accuracy of the nonparametric regression estimator, we split the original sample into two sub-samples. The first one is called the learning sample, which contains the first 148




Figure 4: A plot of the estimated densities of regression errors.

Table 4: Point forecast accuracy of the nonparametric regression with the local linear estimator, where the bandwidths are selected by the ROT and Bayesian methods.

Method	MSFE	MAFE	MAPE
ROT	229.6831	11.7569	1.0836%
Bayesian	197.8883	10.9798	1.0041%

units $\{y_i, x_{i,1}, x_{i,2}, x_{i,3}\}_{i=1, \dots, 148}$. The second one is called the testing sample, which contains the last 37 units $\{y_j, x_{j,1}, x_{j,2}, x_{j,3}\}_{j=149, \dots, 185}$. The learning sample allows us to build the regression model with the estimated bandwidths, while the testing sample allows us to evaluate prediction accuracy.

To measure prediction accuracy, we consider the mean square forecast error (MSFE), mean absolute forecast error (MAFE) and mean absolute percentage error (MAPE). These are expressed as:

$$\begin{aligned} \text{MSFE} &= \frac{1}{37} \sum_{\omega=1}^{37} (y_{\omega} - \hat{y}_{\omega})^2, \\ \text{MAFE} &= \frac{1}{37} \sum_{\omega=1}^{37} |y_{\omega} - \hat{y}_{\omega}|, \\ \text{MAPE} &= \frac{1}{37} \sum_{\omega=1}^{37} \left| \frac{y_{\omega} - \hat{y}_{\omega}}{y_{\omega}} \right| \times 100, \end{aligned}$$

where y_{ω} denotes holdout samples, while \hat{y}_{ω} denotes the point forecasts. In Table 4, the MSFE and MAFE of the nonparametric regression estimator with bandwidths selected by the ROT and Bayesian bandwidth estimation methods are presented.

With the Bayesian approach, we can also compute the prediction interval nonparametrically. To this end, we first compute the cumulative density function (CDF) of the error distribution, over a set of grid points with a range of residuals. Then, we take the inverse of the CDF and find two grid points that are closest to the $\alpha/2$ and $(1 - \alpha/2)$ quantile, where $\alpha = 0.05$ is the customarily level of significance. The 95% pointwise prediction intervals are obtained by adding the two corresponding grid points to the point forecasts. For instance, the holdout samples of response are shown in solid diamond-shaped dots; the point forecasts are shown as solid black circles, while the 95% pointwise prediction intervals are shown as vertical bars in Figure 5.

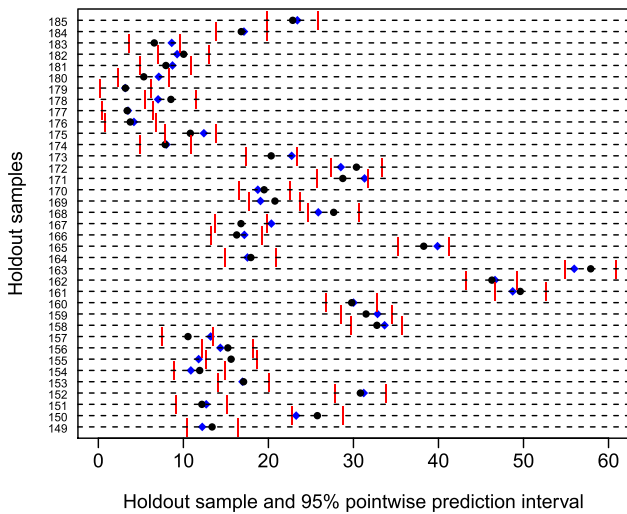


Figure 5: A plot of holdout samples and the 95% pointwise prediction intervals. The holdout data are shown as blue solid diamond-shaped dots. The point forecasts are shown as solid black circles, while the 95% prediction intervals are shown as vertical bars. The empirical coverage probability is 94.59%.

5 An application to SPD estimation

In the second application study, we consider a nonparametric regression model involving SPD estimation. Aït-Sahalia and Lo (1998) show that in a dynamic equilibrium model, the price of a security is expressed as

$$P_t = \exp(\gamma_{t,\lambda}) \mathbb{E}_t^* [Z(S_T)] = \exp(\gamma_{t,\lambda}) \int_{-\infty}^{\infty} Z(S_T) f_t^*(S_T) dS_T,$$

where $T = t + \lambda$, λ is the length of time to maturity, $\gamma_{t,\lambda}$ is a constant risk-free interest rate between t and T , \mathbb{E}_t^* is the expectation conditional on information available at date t , S_T is the price of the security at date T , $Z(S_T)$ is the payoff of the security at the maturity date T , and $f_t^*(S_T)$ is the date t SPD for the payoff of the security at date T .

When an option is of interest, the SPD is the second-order derivative of a call-option pricing formula with respect to the strike price calculated at S_T . According to Aït-Sahalia and Lo (1998), the date t price of a call option is a nonlinear function of $(S_t, X_t, \lambda, \gamma_{t,\lambda}, \delta_{t,\lambda})^\top$, which can be estimated through nonparametric regression, where X_t is the strike price on a stock with date t price S_t , and $\delta_{t,\lambda}$ is the dividend yield.

To reduce the number of regressors, Aït-Sahalia and Lo (1998) assume that the call-option pricing formula is given by the Black-Scholes (BS) formula, where the date t volatility denoted by σ_t , is estimated by the nonparametric regression of the implied volatility on $\tilde{z}_t = (F_t, X, \delta)$. Note that F_t is the futures price of the underlying asset, X is the strike price, and δ is the dividend yield. The kernel estimator of the regression function is

$$\hat{\sigma}_t(F_t, X, \lambda | \mathbf{h}) = \frac{\sum_{j=1}^n K_{\mathbf{h}}(\tilde{z}_t - \tilde{z}_j) \tilde{\sigma}_j}{\sum_{j=1}^n K_{\mathbf{h}}(\tilde{z}_t - \tilde{z}_j)},$$

where $\tilde{\sigma}_j$ is the volatility implied by the price of the call option, and $\mathbf{h} = (h_1, h_2, h_3)^\top$ is a vector of bandwidths. According to Aït-Sahalia and Lo (1998) and Huynh, Kervella, and Zheng (2002), the SPD is expressed as

$$f_{\text{BS},t}(S_T) = \frac{1}{S_T \sqrt{2\pi \hat{\sigma}_t^2 \lambda}} \exp \left\{ - \frac{[\ln(S_T/S_t) - (\gamma_{t,\lambda} - \delta_{t,\lambda} - \hat{\sigma}_t^2/2)\lambda]^2}{2\hat{\sigma}_t^2 \lambda} \right\}.$$

The sample contains S&P 500 options data from January 4 – December 31, 1993, and the sample size is $n = 14,431$. We apply the proposed Bayesian sampling algorithm to estimate bandwidths. Table 5 presents the estimates of the bandwidths and their associated statistics. As a result, the bandwidth vector for the local linear estimator is $(0.9340, 4.7531, 9.4147)^\top$, which clearly differs from $(2.5336, 8.1416, 17.8335)^\top$ and $(0.2976, 7.0521, 3.4184)^\top$ derived through the ROT and CV methods. Moreover, the bandwidth for the kernel-form error density estimated through Bayesian sampling is 0.6134.

Using the bandwidth vectors obtained from the ROT, CV and Bayesian methods, we plot the SPD graphs at maturities of two and 10 days respectively, in Figure 6. At these maturities, the SPDs derived through Bayesian

Table 5: Estimates of the bandwidths and statistics: S&P500 index options data. The burn-in period contains the first 100 draws, and the successive 500 draws were recorded.

Bandwidth estimate	95% Bayesian credible interval	Standard deviation	Batch-mean standard error	SIF
$b = 0.1132$	(0.0981, 0.1350)	0.0101	0.0014	9.5
$h_1 = 1.2887$	(1.2592, 1.3067)	0.0133	0.0040	44.5
$h_2 = 6.5582$	(6.3092, 6.8044)	0.1597	0.0525	54.1
$h_3 = 12.9901$	(12.2074, 13.2636)	0.3055	0.1003	53.8




Figure 6: Plots of the estimated SPD at maturities of 2 and 10 days based on S&P 500 index options data.

sampling are different from those derived through the ROT and CV methods, especially at the tail of the distribution.

6 Conclusion

This paper presents a Bayesian sampling method to simultaneously estimate bandwidths for the local linear estimator of the regression function and kernel-form error density estimator in a nonparametric regression model, where the unknown error density is approximated by the mixture of Gaussian densities centered at individual residuals, and scaled by a bandwidth. We conducted a simulation study and demonstrated that the proposed Bayesian method outperforms ROT and CV in estimating bandwidths under the criterion of ISE. The proposed Bayesian method is able to provide not only bandwidth estimates but Bayesian credible intervals, which allow for assessing uncertainties associated with bandwidth estimates. Our Bayesian sampling procedure represents a data-driven solution to the problem of simultaneously estimating bandwidths for the kernel estimator of a regression function and kernel-form error density. It is well known that local linear fitting has many advantages over local constant fitting (e.g., Fan and Gijbels 1996). As shown in Figure 3, this simple extension is necessary.

For a nonparametric regression model of firm ownership concentration, we estimated bandwidths for local linear fitting; notably, the CV method failed to provide reasonable bandwidths. We also implemented the proposed Bayesian sampling algorithm to estimate bandwidths for the nonparametric regression model involved in SPD estimation. The SPD with bandwidths estimated through the proposed Bayesian sampling method differed from those with bandwidths via the ROT and CV methods. In comparison with ROT, our proposed method relaxes the assumption of a strict Gaussian error density to a mixture Gaussian error density. In comparison with CV, our proposed method utilizes prior information to constrain the range of optimal bandwidths, and consequently, a meaningful bandwidth estimate is obtained.

Author contribution: All the authors have accepted responsibility for the entire content of this submitted manuscript and approved submission.

Research funding: None declared.

Conflict of interest statement: The authors declare no conflicts of interest regarding this article.

References

- Aït-Sahalia, Y., and A. W. Lo. 1998. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *The Journal of Finance* 53: 499–547.
- Bowman, A. W., and A. Azzalini. 1997. *Applied Smoothing Techniques for Data Analysis*. London: Oxford University Press.
- Cheng, F. 2004. "Weak and Strong Uniform Consistency of a Kernel Error Density Estimator in Nonparametric Regression." *Journal of Statistical Planning and Inference* 119: 95–107.

- Chernozhukov, V., and H. Hong. 2003. "An MCMC Approach to Classical Estimation." *Journal of Econometrics* 115: 293–346.
- Chib, S. 1995. "Marginal Likelihood from the Gibbs Output." *Journal of the American Statistical Association: Theory and Methods* 90: 1313–21.
- Efromovich, S. 2005. "Estimation of the Density of Regression Errors." *The Annals of Statistics* 33: 2194–227.
- Fan, J., and I. Gijbels. 1996. *Local Polynomial Modelling and its Applications*. Boca Raton: Chapman & Hall.
- Faraway, J. J., and M. Jhun. 1990. "Bootstrap Choice of Bandwidth for Density Estimation." *Journal of the American Statistical Association: Theory and Methods* 85: 1119–22.
- Garthwaite, P. H., Y. Fan, and S. A. Sisson. 2016. "Adaptive Optimal Scaling of Metropolis-Hastings Algorithms Using the Robbins-Monro Process." *Communications in Statistics - Theory and Methods* 45: 5098–111.
- Gelman, A., G. O. Roberts, and W. R. Gilks. 1996. "Efficient Metropolis Jumping Rules." *Bayesian Statistics* 5: 599–607.
- Geweke, J. F. 1999. "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication." *Econometric Reviews* 18: 1–73.
- Geweke, J. 2010. *Complete and Incomplete Econometric Models*. New Jersey: Princeton University Press.
- Gourieroux, C., A. Monfort, and A. Trognon. 1984. "Pseudo Maximum Likelihood Methods: Theory." *Econometrica* 52: 681–700.
- Hall, P., S. N. Lahiri, and J. Polzehl. 1995. "On Bandwidth Choice in Nonparametric Regression with Both Short-And Long-Range Dependent Errors." *The Annals of Statistics* 23: 1921–36.
- Härdle, W. 1990. *Applied Nonparametric Regression*. Cambridge: Cambridge University Press.
- Härdle, W. 1991. *Smoothing Techniques with Implementation in S*. New York: Springer-Verlag.
- Härdle, W., and E. Mammen. 1993. "Comparing Nonparametric Versus Parametric Regression Fits." *The Annals of Statistics* 21: 1926–47.
- Härdle, W., and J. S. Marron. 1985. "Optimal Bandwidth Selection in Nonparametric Regression Function Estimation." *The Annals of Statistics* 13: 1465–81.
- Härdle, W., and M. Müller. 2000. "Multivariate and Semiparametric Kernel Regression." In *Smoothing and Regression: Approaches, Computation, and Application*, edited by M. G. Schimek, 357–91. New York: John Wiley & Sons.
- Hayfield, T., and J. S. Racine. 2008. "Nonparametric Econometrics: The np Package." *Journal of Statistical Software* 27.
- Herrmann, E., J. Engel, M. P. Wand, and T. Gasser. 1995. "A Bandwidth Selector for Bivariate Kernel Regression." *Journal of the Royal Statistical Society: Series B* 57: 171–80.
- Huynh, K., P. Kervella, and J. Zheng. 2002. "Estimating State Price Densities with Nonparametric Regression." In *Applied Quantitative Finance*, edited by W. Härdle, T. Kleinow and G. Stahl, 171–96. Heidelberg: Springer-Verlag.
- Jaki, T., and R. W. West. 2008. "Maximum Kernel Likelihood Estimation." *Journal of Computational & Graphical Statistics* 17: 976–93.
- Jaki, T., and R. W. West. 2011. "Symmetric Maximum Kernel Likelihood Estimation." *Journal of Statistical Computation and Simulation* 81: 193–206.
- Kass, R. E., and A. E. Raftery. 1995. "Bayes Factors." *Journal of the American Statistical Association: Review Paper* 90: 773–95.
- Kim, S., N. Shephard, and S. Chib. 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models." *The Review of Economic Studies* 65: 361–93.
- Kuruwita, C. N., K. B. Kulasekera, and C. M. Gallagher. 2011. "Generalized Varying Coefficient Models with Unknown Link Function." *Biometrika* 98: 701–10.
- Lee, C., and P. J. Neal. 2018. "Optimal Scaling of the Independence Sampler: Theory and Practice." *Bernoulli* 24: 1636–52.
- Li, Q., and J. Racine. 2004. "Cross-validated Local Linear Nonparametric Regression." *Statistica Sinica* 14: 485–512.
- Li, Q., and J. Zhou. 2005. "The Uniqueness of Cross-validation Selected Smoothing Parameters in Kernel Estimation of Nonparametric Models." *Econometric Theory* 21: 1017–25.
- Linton, O., and Z. Xiao. 2007. "A Nonparametric Regression Estimator that Adapts to Error Distribution of Unknown Form." *Econometric Theory* 23: 371–413.
- Martins-Filho, C., F. Yao, and M. Torero. 2018. "Nonparametric Estimation of Conditional Value-at-risk and Expected Shortfall Based on Extreme Value Theory." *Econometric Theory* 34: 23–67.
- Müller, U. 2013. "Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix." *Econometrica* 81: 1805–49.
- Nadaraya, E. A. 1964. "On Estimating Regression." *Theory of Probability and Its Applications* 9: 141–2.
- Padgett, W. J., and L. A. Thombs. 1986. "Smooth Nonparametric Quantile Estimation under Censoring: Simulations and Bootstrap Methods." *Communications in Statistics – Simulation and Computation* 15: 1003–25.
- R Core Team. 2020. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing, <https://www.R-project.org/>.
- Roberts, G. O. 1996. "Markov Chain Concepts Related to Sampling Algorithms." In *Markov Chain Monte Carlo in Practice*, edited by W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, 45–57. London: Chapman & Hall.
- Roberts, G. O., and J. S. Rosenthal. 2001. "Optimal Scaling for Various Metropolis-Hastings Algorithms." *Statistical Science* 16: 351–67.
- Roberts, G. O., and J. S. Rosenthal. 2009. "Examples of Adaptive MCMC." *Journal of Computational & Graphical Statistics* 18: 349–67.

- Roberts, G. O., A. Gelman, and W. R. Gilks. 1997. "Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithm." *The Annals of Applied Probability* 7: 110–20.
- Romano, J. P. 1988. "On Weak Convergence and Optimality of Kernel Density Estimate of the Mode." *The Annals of Statistics* 16: 629–47.
- Sain, S. R., K. A. Baggerly, and D. W. Scott. 1994. "Cross-validation of Multivariate Densities." *Journal of the American Statistical Association: Theory and Methods* 89: 807–17.
- Samb, R. 2010. "Nonparametric Kernel Estimation of the Probability Density Function of Regression Errors Using Estimated Residuals," Working paper, Universite Pierre et Marie Curie, LSTA, <https://arxiv.org/abs/1010.0439>.
- Samb, R.. 2011. "Nonparametric Estimation of the Density of Regression Errors." *Comptes Rendus Mathematique* 349: 1281–5.
- Scott, D. W. 1992. *Multivariate Density Estimation: Theory, Practice, and Visualization*. New York: Wiley.
- Silverman, B. W. 1986. *Density Estimation for Statistics and Data Analysis*. New York: Chapman & Hall.
- Wahba, G., and S. Wold. 1975. "A Completely Automatic French Curve: Fitting Spline Functions by Cross Validation." *Communications in Statistics – Theory and Methods* 4: 1–17.
- Watson, G. S. 1964. "Smooth Regression Analysis," *Sankhya: The Indian Journal of Statistics, Series A* 26: 359–72.
- Yafeh, Y., and O. Yosha. 2003. "Large Shareholders and Banks: Who Monitors and How?" *The Economic Journal* 113: 128–46.
- Yuan, A., and J. G. de Gooijer. 2007. "Semiparametric Regression with Kernel Error Model." *Scandinavian Journal of Statistics* 34: 841–69.
- Zhang, X., R. D. Brooks, and M. L. King. 2009. "A Bayesian Approach to Bandwidth Selection for Multivariate Kernel Regression with an Application to State-price Density Estimation." *Journal of Econometrics* 153: 21–32.
- Zhang, X., M. L. King, and R. J. Hyndman. 2006. "A Bayesian Approach to Bandwidth Selection for Multivariate Kernel Density Estimation." *Computational Statistics & Data Analysis* 50: 3009–31.
- Zhang, X., M. L. King, and H. L. Shang. 2014. "A Sampling Algorithm for Bandwidth Estimation in a Nonparametric Regression Model with a Flexible Error Density." *Computational Statistics & Data Analysis* 78: 218–34.
- Zhang, X., M. L. King, and H. L. Shang. 2016. "Bayesian Bandwidth Selection for a Nonparametric Regression Model with Mixed Types of Regressors." *Econometrics* 4, Article 24.

Supplementary Material: The online version of this article offers supplementary material (<https://doi.org/10.1515/snde-2018-0050>).