

Linking frequentist and Bayesian change-point methods

D. Ardia, A. Dufays, C. Ordás Criado

G–2023–70

December 2023

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

CITATION ORIGINALE / ORIGINAL CITATION

D. Ardia, A. Dufays, C. Ordás Criado (2023). "Linking Frequentist and Bayesian Change-Point Methods", *Journal of Business and Economic Statistics*, <https://doi.org/10.1080/07350015.2023.2293166>.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2023
– Bibliothèque et Archives Canada, 2023

Legal deposit – Bibliothèque et Archives nationales du Québec, 2023
– Library and Archives Canada, 2023

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Télec. : 514 340-5665
info@gerad.ca
www.gerad.ca

Linking frequentist and Bayesian change-point methods

David Ardia ^a

Arnaud Dufays ^b

Carlos Ordás Criado ^c

^a *Department of Decision Sciences, HEC Montréal & GERAD, Montréal, (Qc), Canada, H3T 2A7*

^b *Faculty of Data Science, Economics & Finance, EDHEC Business School, 59100 Roubaix, France*

^c *Department of Economics, Laval University, Québec (Qc), Canada, G1V 0A6*

david.ardia@hec.ca

arnaud.dufays@edhec.edu

December 2023
Les Cahiers du GERAD
G–2023–70

Copyright © 2023 Ardia, Dufays, Ordás Criado

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : We show that the two-stage minimum description length (MDL) criterion widely used to estimate linear change-point (CP) models corresponds to the marginal likelihood of a Bayesian model with a specific class of prior distributions. This allows results from the frequentist and Bayesian paradigms to be bridged together. Thanks to this link, one can rely on the consistency of the number and locations of the estimated CPs and the computational efficiency of frequentist methods, and obtain a probability of observing a CP at a given time, compute model posterior probabilities, and select or combine CP methods via Bayesian posteriors. Furthermore, we adapt several CP methods to take advantage of the MDL probabilistic representation. Based on simulated data, we show that the adapted CP methods can improve structural break detection compared to state-of-the-art approaches. Finally, we empirically illustrate the usefulness of combining CP detection methods when dealing with long time series and forecasting.

Keywords : Change-point, minimum description length, model selection/combination, structural change

Acknowledgements: Preliminary versions of this article were presented at the 2019 annual R/Finance conference in Chicago, at the Econometric Society & Bocconi University Virtual World Congress 2020, the 35th European Economic Association Virtual Congress 2020, and the NBER-NSF Time Series Conference 2023. We benefited from helpful comments from Laurent Barras, Luc Bauwens, Keven Bluteau, Kris Boudt, Serge Darroles, Linda Mhalla, Jeroen Rombouts, and Yong Song. We are grateful to the associate editor and two anonymous reviewers for insightful comments, which helped to greatly improve the article. The usual disclaimer applies.

Data Availability Statement: Please contact the corresponding author to have access to the computer code and the datasets.

Disclosure Statement: The authors report there are no competing interests to declare.

Funding : We gratefully acknowledge the financial support of IVADO, NSERC (grant RGPIN-2022-03767), and the Fonds de recherche du Québec – Société et culture (grant #193277).

1 Introduction

Change-point (CP) models are used to account for abrupt changes in an ordered sequence of observations. When a series contains structural shifts, the data can be modeled as disjoint segments of different stochastic processes. CP models typically lead to a better understanding of critical events and consistent estimates of model parameters for each segment and can also improve predictions. Due to these appealing features, CP models have attracted considerable interest in economics and natural sciences. The number and locations of the breaks are generally unknown to the researcher and must be estimated. Many methods are now available to detect CPs both in frequentist and Bayesian frameworks (see, *e.g.*, Eckley, Fearnhead and Killick, 2011; Bauwens, Koop, Korobilis and Rombouts, 2015).

Frequentist and Bayesian CP approaches have mostly evolved independently in the literature, as they often explore different aspects of the CP problem. The frequentist approach has mainly focused on developing asymptotic frameworks for break detection (*e.g.*, Yao and Au, 1989; Liu, Wu and Zidek, 1997; Ciuperca, 2011; Davis, Yau et al., 2013), testing the number and location of breaks (*e.g.*, Bai and Perron, 1998), and providing efficient algorithms to explore possible segments (*e.g.*, Bai, 1997; Bai and Perron, 2003; Jackson et al., 2005; Killick et al., 2012; Fryzlewicz, 2014; Yau and Zhao, 2016). Depending on whether a CP detection method explores all possible partitions simultaneously or sequentially for a subset of segments, CP detection can be classified as global (or exact) or local, respectively. In contrast, the Bayesian CP literature has mainly focused on developing schemes for efficiently drawing from the posterior distribution of the CPs (*e.g.*, Carlin, Gelfand and Smith, 1992; Stephens, 1994; Chib, 1998; Giordani and Kohn, 2008; Maheu and Song, 2018) and on forecasting future breaks (*e.g.*, Pesaran, Pettenuzzo and Timmermann, 2006; Koop and Potter, 2007). Also, the criteria used to determine the number of CPs and their locations differ across paradigms. Frequentists mostly rely on (several variants of) the Schwarz information criterion (SIC; *e.g.* Yao and Au, 1989; Liu, Wu and Zidek, 1997; Ciuperca, 2011; Fryzlewicz, 2014), the sum of squared residuals (*e.g.* Bai, 1997; Bai and Perron, 1998), and the two-stage minimum description length (MDL) criterion (*e.g.* Davis, Lee and Rodriguez-Yam, 2006; Davis, Yau et al., 2013; Yau and Zhao, 2016), while Bayesians typically maximize the marginal likelihood (Chib, 1998; Bernardo, Bayarri, Berger, Dawid, Heckerman, Smith and West, 2007; Bauwens, Dufays and Rombouts, 2014b; Du, Kao and Kou, 2016). While most criteria used by frequentists consistently estimate the number and locations of the breaks, this is rarely investigated in the Bayesian framework. An exception is Du, Kao and Kou (2016), who establishes the consistency of the marginal likelihood to estimate the total number and location of the CPs for a piecewise-constant model.

CP model selection and combination have also been intensively studied in frequentist and Bayesian paradigms (*e.g.*, Chan, Yip and Zhang, 2014; Maheu and Gordon, 2008). Yet, links between the two statistical frameworks have been overlooked. Our article addresses this gap by showing that the MDL criterion corresponds to the marginal likelihood of a Bayesian model with a specific class of prior distributions. Thanks to this link, the marginal likelihood inherits the asymptotic properties of the MDL criterion, and can further be used to detect CPs, assign posterior probabilities to competing CP methods, and select or combine CP methods via their posterior probabilities. Moreover, popular state-of-the-art CP methods can be adapted to the proposed marginal likelihood to improve their performance.

Similar links have been proposed in the model selection literature, in particular for normal-linear models. However, none of them tackles the CP detection problem. For example, Smith and Spiegelhalter (1980) propose deterministic functions linking the (logarithm of the) Bayes factors (ratios of marginal likelihoods) of normal-linear regression models to the Schwarz or Akaike information criteria, but their focus is on how the prior specifications impact the Bayes factor. In the same vein, Kass and Raftery (1995) emphasize that the SIC can be viewed as a rough asymptotic approximation of the (logarithm of the) Bayes factor that avoids specifying priors. Hansen and Yu (2001) recall that the MDL criterion is a first-order approximation of the marginal likelihood. They also show that the marginal likelihood can be understood as an alternative description length criterion, referred to as

“mixture MDL.” However, these links are mainly asymptotic, the focus is on variable selection, and they ignore structural change. Our article extends this literature to CP detection for normal-linear regression models by (i) providing a class of proper priors belonging to the Normal Inverse-Gamma (\mathcal{NIG}) distributions that leads to an exact (and thus finite-sample) equivalence between the marginal likelihood of the CP model and the MDL criterion, and (ii) allowing the variance of linear models to vary over time.

The proposed calibration has major implications for frequentist and Bayesian CP approaches. Researchers typically obtain distinct breakpoints in applied settings when different (state-of-the-art) local or global CP detection methods are used. Providing a hyperparameter calibration that leads to an exact link between the marginal likelihood and a frequentist information criterion allows frequentists to rely on the Bayesian framework, make probabilistic statements about competing models, and combine or select candidate models via their posterior probabilities. In addition, frequentists can rely on Bayesian credible intervals for the break locations rather than confidence intervals. The former are straightforward to compute and have an advantage over the latter as they do not overlap and remain in-sample. Finally, the stopping rule of several sequential detection methods, such as binary segmentation, can be adapted to reflect a probabilistic statement between competing models. This feature strongly improves the performance of sequential methods. The proposed framework also benefits Bayesians. It provides hyperparameter values that motivate the prior specification of the Bayesian CP model as they allow the marginal likelihood to inherit asymptotic properties established in the frequentist literature. Asymptotic properties for the marginal likelihood have been developed in Du, Kao and Kou (2016) for piecewise-constant models. The proposed link allows us to extend this result to piecewise-linear Bayesian normal regression models. The consistency of the marginal likelihood directly follows from the theoretical works of Davis, Lee and Rodriguez-Yam (2006); Davis, Yau et al. (2013). Additionally, our framework accounts for future breaks in long time series context. It is worth noting that Bayesian methods that allow for future breaks, such as Pesaran, Pettenuzzo and Timmermann (2006), Koop and Potter (2007), and Maheu and Song (2018), cannot be applied to long time series due to their algorithmic complexity.

In practice, global detection methods cannot always be applied due to their computational cost. Local approaches, such as binary segmentation and its extensions, are feasible alternatives (*e.g.*, Bai, 1997; Fryzlewicz, 2014; Yau and Zhao, 2016; Korkas and Fryzlewicz, 2017). Our article broadens the set of existing CP detection methods in three ways: (i) by proposing three local algorithms that take advantage of the probabilistic representation of the MDL criterion, (ii) by adapting a fast and deterministic global method following Eckley, Fearnhead and Killick (2011), and (iii) by showing how to compute credible intervals for the estimated CPs and regression parameters, which account for both in-sample and out-of-sample breaks as well as model uncertainty. These novel approaches are contrasted with existing ones using simulated data. They detect and locate breaks with the same or better accuracy than existing methods. Two applications illustrate the relevance of combining local CP detection methods when dealing with long time series and forecasting. In the first application, we estimate a five-factor Fama-French CP model over a long time period by combining local CP methods. We explore the evolution of model parameters over time while accounting for both method and break uncertainty. The large set of detected breakpoints allows us to accurately infer the break process and to provide credible intervals for model parameters when future breakpoints occur. In the second application, we forecast the monthly U.S. inflation and demonstrate that accounting for breaks and method uncertainty leads to significant improvements in the predictive performance of the model in terms of root mean square forecast errors and mean absolute forecast errors. In particular, allowing for an out-of-sample break leads to better predictions over longer time horizons.

The article is organized as follows. In Section 2, we present the methodology. Section 3 introduces adapted local and global CP detection methods that build upon the probabilistic representation of the MDL criterion. We test how the adapted methods collectively perform in a simulation exercise in Section 4. Section 5 proposes a full Bayesian model providing posterior distributions that account for break uncertainty and model uncertainty. Section 6 develops the two applications, and Section 7 concludes.

2 Change-point framework

2.1 Model

Consider the following linear regression model with no structural breaks:

$$y_t = \beta_0 + \beta_1 x_{t,1} + \dots + \beta_{K-1} x_{t,K-1} + \sigma \epsilon_t = \mathbf{x}'_t \boldsymbol{\beta} + \sigma \epsilon_t, \quad t = 1, \dots, T, \quad (1)$$

where y_t is the observed dependent variable, $\mathbf{x}_t = (1, x_{t,1}, \dots, x_{t,K-1})'$ is a $(K \times 1)$ vector including a constant and $K - 1$ covariates, $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_{K-1})'$ and $\sigma > 0$ are unknown parameters, and $\epsilon_t \sim iid\mathcal{N}(0, 1)$.¹ In applied settings, it is often unreasonable to assume that $\boldsymbol{\beta}$ and σ are time-invariant, especially for long time series. A more flexible linear specification allowing for $m \geq 0$ breaks ($m + 1$ regimes) in the model parameters is:

$$y_t = \mathbf{x}'_t \boldsymbol{\beta}_i + \sigma_i \epsilon_t, \quad \text{for } \tau_{i-1} < t \leq \tau_i, \quad i = 1, \dots, m + 1, \quad (2)$$

where τ_i denotes the i -th CP location on the $[0, T]$ time segment, and $\boldsymbol{\beta}_i$ and σ_i are the model parameters holding for the period starting after time τ_{i-1} and ending at time τ_i , with $\tau_{i-1} < \tau_i$. The boundary conventions $\tau_0 = 0$ and $\tau_{m+1} = T$ are used. To simplify notation, the vectors of observations belonging to segment i are denoted by $\mathbf{y}_i = \mathbf{y}_{(\tau_{i-1}+1):\tau_i} = (y_{\tau_{i-1}+1}, \dots, y_{\tau_i})'$ for the dependent variable and $\mathbf{X}_i = [\mathbf{x}_{\tau_{i-1}+1} \ \mathbf{x}_{\tau_{i-1}+2} \ \dots \ \mathbf{x}_{\tau_i}]'$ for the covariates. Specification (2) encompasses stationary piecewise-linear autoregressive processes of order q_i (*i.e.*, CP-AR(q_i)). We collect the CPs located on the $]0, T[$ interval in vector $\boldsymbol{\tau} = (\tau_1, \dots, \tau_m)'$ and gather the segment parameters in $\Theta = \{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{m+1}\}$, where $\boldsymbol{\theta}_i = (\boldsymbol{\beta}_i, \sigma_i^2)$. The number of CPs, $\dim(\boldsymbol{\tau}) = m$, is typically restricted by a fixed upper bound $M \leq T$. We further assume that $\mathbf{X}'_i \mathbf{X}_i$ is nonsingular for all considered segments. The purpose of a CP method is to estimate the number of CPs m , break locations $\boldsymbol{\tau}$, and regression parameters Θ .

2.2 Marginal likelihood

Given the wide variety of CP methods available, model selection and model averaging are of particular interest in our setup. Model selection deals with model uncertainty by selecting one model from a set of P candidate models $\{\mathcal{M}_1, \dots, \mathcal{M}_P\}$ based on some selection criterion (*e.g.*, AIC, SIC, MDL). Model averaging deals with model uncertainty by combining several models from a set of candidate models based on a weighting scheme. In Bayesian statistics, dealing with model uncertainty is straightforward in theory: Model selection and model averaging can be achieved by using model posterior probabilities. To compute these probabilities, the marginal likelihood function under each model needs to be evaluated. In general, this step is not trivial. However, in some specific cases, analytical expressions exist. For instance, in the context of normal-linear regression with no structural change, the marginal likelihood exhibits a closed-form expression for \mathcal{NIG} priors (see, *e.g.*, Fernandez, Ley and Steel, 2001). This closed form extends to piecewise-linear models conditional on the CPs (see, *e.g.*, Du, Kao and Kou (2016) and Proposition 1 below). Without loss of generality and for the sake of simplicity, we drop the conditioning on \mathbf{X}_i in what follows. Consider first the general expression of the marginal likelihood for the set of $m + 1$ independent regimes associated with vector $\boldsymbol{\tau}$:

$$f(\mathbf{y}_{1:T} | \boldsymbol{\tau}) = \prod_{i=1}^{m+1} f(\mathbf{y}_i | \boldsymbol{\tau}), \quad (3)$$

where $f(\mathbf{y}_i | \boldsymbol{\tau})$ is the density of observing the data over the period $]\tau_{i-1}, \tau_i]$, and where the model parameters Θ have been integrated out.² Proposition 1 recalls a well-known result in Bayesian statistics

¹While restrictive, several recent theoretical and empirical CP frameworks rely on the normality assumption (see, *e.g.*, Rigaiil, Lebarbier and Robin (2012); Maheu and Song (2014, 2018); Smith, Bulkley and Leslie (2020) in the Bayesian paradigm and Safikhani and Shojaie (2020) in the frequentist paradigm).

²When the segments are not independent, (3) can be replaced with $f(\mathbf{y}_{1:T} | \boldsymbol{\tau}) = \prod_{i=1}^{m+1} f(\mathbf{y}_{(\tau_{i-1}+1):\tau_i} | \boldsymbol{\tau}, \mathbf{y}_{1:\tau_{i-1}})$, where $\mathbf{y}_{1:0} = \{\emptyset\}$.

and shows the prior distributions as well as the analytical expression of the marginal likelihood, which are used throughout the article.

Proposition 1. *Consider the piecewise-linear regression model and the following $\mathcal{N}\mathcal{I}\mathcal{G}$ priors:*

$$\begin{aligned}\beta_i | \sigma_i^2, \tau &\sim \mathcal{N} \left(\underline{\beta}_i, \sigma_i^2 g_i \underline{\mathbf{M}}_i^{-1} \right), \\ \sigma_i^2 | \tau &\sim \mathcal{IG} \left(\frac{\underline{\nu}_i}{2}, \frac{\underline{s}_i}{2} \right),\end{aligned}\tag{4}$$

for $i = 1, \dots, m+1$, where $\underline{\beta}_i \in \mathbb{R}^K$, $\underline{\mathbf{M}}_i \in \mathbb{R}^{K \times K}$ is a definite positive matrix, $\underline{\nu}_i > 0$ and $\underline{s}_i > 0$ are the hyperparameters holding in segment i . The marginal likelihood of the observations belonging to segment i is then given by:

$$f^{\beta_i}(\mathbf{y}_i | \tau) = (2\pi)^{-\frac{n_i}{2}} \left(\frac{|\underline{\mathbf{M}}_i^{-1}|}{|g_i \underline{\mathbf{M}}_i^{-1}|} \right)^{\frac{1}{2}} \frac{\Gamma(\frac{\underline{\nu}_i}{2}) (\frac{\underline{s}_i}{2})^{\frac{\underline{\nu}_i}{2}}}{\Gamma(\frac{\underline{\nu}_i}{2}) (\frac{\underline{s}_i}{2})^{\frac{\underline{\nu}_i}{2}}},\tag{5}$$

where $\bar{\nu}_i = n_i + \underline{\nu}_i$, $\bar{s}_i = \underline{s}_i + s_i + \underline{\beta}_i' g_i^{-1} \underline{\mathbf{M}}_i \underline{\beta}_i + \hat{\beta}_i' \mathbf{X}_i' \mathbf{X}_i \hat{\beta}_i - \bar{\beta}_i' \bar{\mathbf{M}}_i \bar{\beta}_i$, $\bar{\beta}_i = [g_i^{-1} \underline{\mathbf{M}}_i + \mathbf{X}_i' \mathbf{X}_i]^{-1} (g_i^{-1} \underline{\mathbf{M}}_i \underline{\beta}_i + \mathbf{X}_i' \mathbf{X}_i \hat{\beta}_i)$ and in which $\hat{\beta}_i$ and s_i denote the ordinary least-squares (OLS) estimates of β_i and the related sum of squared residuals computed for segment i , respectively.

Proof. See the online Appendix, Section I.A. □

Marginal likelihood $f^{\beta_i}(\mathbf{y}_i | \tau)$ in (5) depends on hyperparameters $\underline{\beta}_i$, $\underline{\mathbf{M}}_i$, $\underline{\nu}_i$, and \underline{s}_i , that we treat as constants (see, *e.g.*, Du, Kao and Kou, 2016). We further propose to set the constants so that $f^{\beta_i}(\mathbf{y}_i | \tau)$ reduces to the MDL criterion, which is commonly used in the frequentist CP literature to estimate piecewise-linear regression models (see Section 2.3). With well-justified presets at hand, $f^{\beta_i}(\mathbf{y}_i | \tau)$ can be used to compute model posterior probabilities and make probabilistic statements about competing models.

To illustrate the latter point, consider the set of P vectors $\{\tau^1, \dots, \tau^P\}$ and let prior $f(\tau^p)$ be uniform. Then, the posterior probability of τ^p ($p = 1, \dots, P$) is given by:

$$f(\tau^p | \mathbf{y}_{1:T}) = \frac{f(\mathbf{y}_{1:T} | \tau^p)}{\sum_{j=1}^P f(\mathbf{y}_{1:T} | \tau^j)}.\tag{6}$$

By using (3) and (5) in (6), the latter expression becomes feasible. Thus, we can use these model posterior probabilities for, say, selecting the candidate vector of CPs that exhibits the greatest posterior probability (*i.e.*, $\max_{p \in \{1, P\}} f(\tau^p | \mathbf{y}_{1:T})$) for the given data, or for combining CP methods and averaging posterior quantities of interest by weighting these posterior quantities by the corresponding model posterior probabilities (see Section 5). For instance, model uncertainty can be accounted for in the estimates of, say, parameter θ_t ($t = 1, \dots, T$), by mixing the parameter posterior densities as follows:

$$f(\theta_t | \mathbf{y}_{1:T}) = \sum_{p=1}^P f(\theta_t | \mathbf{y}_{1:T}, \tau^p) f(\tau^p | \mathbf{y}_{1:T}) = \sum_{p=1}^P f(\theta_t | \mathbf{y}_{(\tau_{k_p}^p - 1): \tau_{k_p}^p}^p) f(\tau^p | \mathbf{y}_{1:T}),\tag{7}$$

where $t \in]\tau_{k_p}^p - 1, \tau_{k_p}^p]$ and index k_p denotes the k^{th} segment of the p^{th} model ($k_p = 0, \dots, m_p + 1$). These model selection and model combination approaches are illustrated in Sections 4 and 6. In the context of forecasting, the h -step-ahead ($h \geq 1$) predictive density would similarly account for model uncertainty through the weighted average:

$$f(\mathbf{y}_{T+1:T+h} | \mathbf{y}_{1:T}) = \sum_{p=1}^P f(\mathbf{y}_{T+1:T+h} | \mathbf{y}_{(\tau_{k_p}^p - 1): \tau_{k_p}^p}^p) f(\tau^p | \mathbf{y}_{1:T}).\tag{8}$$

Section 6 explores forecasting performance accounting for model uncertainty.

2.3 MDL marginal likelihood

As emphasized earlier, the MDL criterion is widely used to (consistently) estimate piecewise-stationary autoregressive models (AR; *e.g.*, Davis, Lee and Rodriguez-Yam, 2006; Killick, Fearnhead and Eckley, 2012; Yau and Zhao, 2016; Ng, Pan and Yau, 2022). In this section, we build upon Proposition 1 and set the hyperparameters of the \mathcal{NIG} priors so that the marginal likelihood reduces to the MDL criterion.

For a piecewise-linear model with m breaks, the MDL criterion reads:

$$\text{MDL}(m, \boldsymbol{\tau}) = \ln f(\mathbf{y}_{1:T} | \boldsymbol{\tau}, \hat{\Theta}_{\text{MLE}}) - \ln^+(m) - (m+1) \ln T - \left(\frac{K+1}{2} \right) \sum_{i=1}^{m+1} \ln n_i, \quad (9)$$

where $f(\mathbf{y}_{1:T} | \boldsymbol{\tau}, \hat{\Theta}_{\text{MLE}})$ is the conditional likelihood function, $\hat{\Theta}_{\text{MLE}}$ is the set of model parameters maximizing the likelihood for the given vector $\boldsymbol{\tau}$, $\ln^+(x) = \max\{0, \ln(x)\}$ for $x \geq 0$, and $n_i = \dim(\mathbf{y}_i)$. Proposition 2 formally establishes the link between a class of Bayesian CP models and the MDL criterion.

Proposition 2. *Under the setting in Proposition 1, let the hyperparameters of the \mathcal{NIG} priors be given by:*

$$\underline{\boldsymbol{\beta}}_i = \hat{\boldsymbol{\beta}}_i, \quad \underline{\mathbf{M}}_i = \mathbf{X}_i' \mathbf{X}_i, \quad \underline{g}_i = \underline{f}_i n_i - 1, \quad \underline{f}_i = \left(\frac{((m^+)^{\frac{1}{m+1}}) n_i^{\frac{1}{4}} T}{(\frac{1}{\sqrt{n_i}} + 1)^{\frac{1}{2}}} \right)^{\frac{2}{K}} \exp\left(\frac{2}{K} \Delta R_{4,i}\right), \quad (10)$$

$$\underline{\nu}_i = \underline{k}_i n_i, \quad \underline{s}_i = \underline{k}_i s_i, \quad \underline{k}_i = \frac{1}{\sqrt{n_i}}, \quad (11)$$

for $i = 1, \dots, m+1$, where $m^+ = \max\{1, m\}$, $\Delta R_{4,i} = R_4(\frac{n_i + \underline{\nu}_i}{2}) - R_4(\frac{\underline{\nu}_i}{2})$ with $R_4(x) = \frac{1}{12x} - \frac{1}{360x^3} + \frac{1}{1260x^5}$ for $x > 0$. Then, the logarithm of the marginal likelihood conditional on $\boldsymbol{\tau}$ is given by:

$$\ln f^{\text{MDL}}(\mathbf{y}_{1:T} | \boldsymbol{\tau}) = \sum_{i=1}^{m+1} \ln f^{\hat{\boldsymbol{\beta}}_i}(\mathbf{y}_i | \boldsymbol{\tau}) = \text{MDL}(m, \boldsymbol{\tau}) + \mathcal{O}\left(\min_i(n_i)^{-\frac{7}{2}}\right), \quad (12)$$

where $\min_i(n_i)$ denotes the number of observations in the smallest regime.

Proof. See the online Appendix, Section I.B. □

Remark 1. The hyperparameters of the inverse-gamma distribution imply that the maximum likelihood estimator $\frac{s_i}{n_i}$ lies in a high-density region of the prior as it falls between the mode and the expectation; see the online Appendix, Section I.C.

Remark 2. The approximation order in (12) can be made arbitrarily small by selecting a high-order N in the remainder $\Delta R_{N,i}$; see the online Appendix, Section I.B.

In the rest of the article, we refer to the marginal likelihood obtained in Proposition 2 as the *MDL marginal likelihood*. Given that $\underline{\mathbf{M}}_i = \mathbf{X}_i' \mathbf{X}_i$ in (10), the priors suggested in Proposition 2 belong to the class of g-priors (see, *e.g.*, Zellner, 1986; Fernandez, Ley and Steel, 2001).

Besides being simple to use, the calibration proposed in Propositions 2 is of great interest to link the frequentist and Bayesian statistical paradigms. For example, the \mathcal{NIG} -prior distributions used in the proposition are also used in many Bayesian CP papers to improve the estimation efficiency (see, *e.g.*, Geweke and Jiang, 2011; Rigaiil, Lebarbier and Robin, 2012; Maheu and Song, 2014). Our calibrations further extend the consistency properties of the MDL criterion to these Bayesian CP setups. The proposed prior hyperparameters lead to an exact equivalence between the marginal likelihood and the MDL criterion. Davis, Yau et al. (2013) show that the MDL criterion is a consistent estimator of the number of CPs, their locations, the order of the linear model, and the parameter values in each segment for a broad class of piecewise-linear stationary autoregressive processes, which nest

canonical model (2). These asymptotic properties also apply to the MDL marginal likelihood under the regularity conditions stated in Davis, Yau et al. (2013) and under the presets given in Proposition 2. Moreover, by establishing the equivalence between the MDL criterion and the marginal likelihood for a specific class of priors, we extend the consistent maximum likelihood estimator of Du, Kao and Kou (2016) from a piecewise-constant to a piecewise-linear multiple regression setup. The consistency of the MDL marginal likelihood is shown numerically in Table 4 of the simulation section.

3 Revisiting CP methods

In this section, we revisit three popular CP methods (two local ones based on binary segmentation and the efficient and global method of Bai and Perron (2003)) by adapting them to the MDL marginal likelihood criterion. In addition, we propose a fourth method that “prunes” the adapted global approach and reduces its complexity to an order lower than $\mathcal{O}(T^2)$. The simulation results reported in Section 4 demonstrate that the pruning achieves a good compromise between accuracy in CP detection and computational cost, particularly for long time series.

3.1 Binary and wild binary segmentation

Binary segmentation, denoted by BS or BS(a, b) hereafter, is a generic procedure that sequentially detects multiple breaks in time series over a finite time segment $[a, b]$ with $a < b$ (see, *e.g.*, Gupta and Chen, 1996). BS relies on the “cumsum” statistic:

$$\tilde{y}_{a,b}(\tau) = \sqrt{\frac{b-\tau}{(b-a+1)(\tau-a+1)}} \sum_{t=a}^{\tau} y_t - \sqrt{\frac{\tau-a+1}{(b-a+1)(b-\tau)}} \sum_{t=\tau+1}^b y_t, \quad (13)$$

where $a < \tau < b$. When searching over the entire time span $[1, T]$, BS(a, b) starts by setting $a = 1$ and $b = T$ and then operates as follows:

- for $\tau \in (a, b)$, compute $\bar{\tau} = \operatorname{argmax}_{\tau} \tilde{y}_{a,b}(\tau)$;
- if $\tilde{y}_{a,b}(\bar{\tau}) > \delta_T$, add $\bar{\tau}$ to the set of CPs. Then, run BS($a, \bar{\tau}$) and BS($\bar{\tau}, b$).

If none of the CP candidates yield a statistic $\tilde{y}_{1,T}(\tau)$ strictly above δ_T , the algorithm stops, and the series does not exhibit structural change. Threshold δ_T is chosen based on theoretical and empirical considerations. As noted by Fryzlewicz (2014), if the series exhibits similarities across regimes (such as short regime duration or offsetting jumps at adjacent segments), BS can lead to a test statistic that is below the threshold. To overcome this weakness, this author proposes an alternative procedure called “wild binary segmentation” and denoted as WBS or WBS(a, b) hereafter. When applied to the observations belonging to the entire interval $[1, T]$, the algorithm starts by setting $a = 1$ and $b = T$ and proceeds as follows:

- generate a set of N random intervals $\{a_n, b_n\}_{n=1}^N$ within $[a, b]$.
- for each interval, apply BS(a_n, b_n) and save the maximum value of the statistic $\tilde{y}_{a,b}(\bar{\tau}) = \max_{n \in [1, N]} \tilde{y}_{a_n, b_n}(\bar{\tau}_n)$, where $\bar{\tau}_n = \operatorname{argmax}_{\tau} \tilde{y}_{a_n, b_n}(\tau)$.
- if $\tilde{y}_{a,b}(\bar{\tau}) > \delta_T$, set $\bar{\tau}$ as a new CP. Then, run WBS($a, \bar{\tau}$) and WBS($\bar{\tau}, b$).

Both BS and WBS rely on fitting criterion (13). By replacing this criterion with:

$$\hat{y}_{a,b}^{\text{MDL}}(\tau) = \ln \frac{f^{\text{MDL}}(\mathbf{y}_{a:b} | \boldsymbol{\tau} = \{a-1, \tau, b\})}{f^{\text{MDL}}(\mathbf{y}_{a:b} | \boldsymbol{\tau} = \{a-1, b\})}, \quad (14)$$

threshold δ_T can be given a probabilistic justification. That is, when $\hat{y}_{a,b}^{\text{MDL}}(\tau) > \delta_T = 0$, the posterior probability of a model exhibiting a break at location τ on interval $[a, b]$ is greater than that of its no-break counterpart. When $\delta_T > 3$, the probability in favor of a break is greater than 95%. By setting

$\delta_T = 3$, we can readily adapt BS and WBS to the Bayesian criterion (14). We call these two modified CP methods based on the MDL marginal likelihood “BSMDL” and “WBSMDL,” respectively, and test their performance in the simulation section.³

3.2 Bai and Perron (2003)’s approach

The efficient global algorithm proposed by Bai and Perron (2003) is one of the most popular algorithms for detecting multiple CPs in linear regression models. For a bounded set of breaks $m = 1, \dots, M$, this algorithm relies on dynamic programming to determine the CP locations; it minimizes the sum of squared residuals for each number of breakpoints in the set. The procedure requires at most least-squares operations of $\mathcal{O}(T^2)$, which is far fewer than the brute force approach of $\mathcal{O}(T^m)$ for any $m > 2$. It further handles minimum regime duration and can readily be used with a specific class of information criteria.

A CP information criterion can be optimized with this algorithm if it can be decomposed into the following additive components:

$$\text{IC}(m, \boldsymbol{\tau}, \mathbf{y}_{1:T}) = u(m, T) + \sum_{i=1}^{m+1} v(\mathbf{y}_i | \tau_{i-1}, \tau_i), \quad (15)$$

where $u(m, T)$ is a function that only depends on the number of CPs and on the sample size, and where $v(\mathbf{y}_i | \tau_{i-1}, \tau_i)$ is a function that depends on the starting and terminal dates τ_{i-1} and τ_i . For instance, Rigaiil, Lebarbier and Robin (2012), Maheu and Song (2014) and Du, Kao and Kou (2016) use different marginal likelihood criteria based on different \mathcal{NIG} priors that comply with (15).

We adapt this global method by replacing the sum of squared residuals with the MDL marginal likelihood. Given m CPs, MDL criterion (9) verifies (15). Indeed, as Eckley, Fearnhead and Killick (2011) point out, no heuristic algorithms such as the genetic ones proposed in Davis, Lee and Rodriguez-Yam (2006) or Li and Lund (2012) are necessary for finding the segments that maximize the MDL criterion. We call this global CP method “GMDL” and explore its performance in the simulation section. When multiple m values are explored, one can readily use posterior probabilities (6) to assess the most likely number of breaks.

Global segmentation based on efficient algorithms of order $\mathcal{O}(T^2)$ can be computationally demanding for long time series. This has motivated researchers to develop local alternatives that exhibit less complexity than $\mathcal{O}(T^2)$. Below, we propose a pruned global method of complexity $\mathcal{O}(T)$. As shown in Bai and Perron (2003), the total number of possible segments in their approach is at most $\frac{T(T+1)}{2}$. Figure 1 illustrates this statement for a sample of size $T = 10$ with a minimal duration of one time period before a potential break occurs.

To reduce the complexity of this global method, we only need to decrease the number of considered segments in Figure 1 by a function proportional to T . In the spirit of WBS, we can evaluate αT randomly-selected segments, where α is a fixed constant.

Alternatively, following the LRSM method of Yau and Zhao (2016), we first estimate a set of potential CPs, $\mathcal{J}^* = \{\tau_1^*, \dots, \tau_{m^*}^*\}$, with a local method and then find the best subset of CPs via a global method, such that:

$$(m^{**}, \mathcal{J}^{**}) = \operatorname{argmax}_{m \leq m^*, \mathcal{J} \subseteq \mathcal{J}^*} \text{MDL}(m, \mathcal{J}). \quad (16)$$

This two-step approach is now standard in the literature (see Chan, Yip and Zhang, 2014; Fryzlewicz, 2014). We replace the MDL criterion of the scanning window proposed in the original LRSM approach

³Note that this adaptation is proposed for the BS algorithm in the context of CP Poisson processes by Young Yang and Kuo (2001). In that case, however, the threshold δ_T is set to 0, and the prior for τ is a noninformative (uniform) distribution.

		Terminal date									
		1	2	3	4	5	6	7	8	9	10
Starting date	1	O	O	X	X	X	X	X	X	X	X
	2		O	O	X	X	X	X	X	X	X
	3			O	O	X	X	X	X	X	X
	4				O	O	X	X	X	X	X
	5					O	O	X	X	X	X
	6						O	O	X	X	X
	7							O	O	X	X
	8								O	O	X
	9									O	O
	10										O

Figure 1: Global method—number of possible segments. The vertical number indicates the initial date of a segment and the horizontal number indicates the ending date. We set the sample size $T = 10$ with a minimum regime duration of one observation (since the start) and an unspecified number of breaks m . For instance, in the first row, possible segments are 1 to 3, 1 to 4, ..., 1 to 10, where the initial and in-between-break dates of the possible segments are denoted by O, and the ending dates are denoted by Xs. In the fourth row, possible segments are 4 to 6, 4 to 7, ..., 4 to 10. Setting a fixed value for m would imply further restrictions (or O terms); see Bai and Perron (2003, Figure 1).

with criterion (14) and compute \mathcal{J}^* as follows:

$$\mathcal{J}^* = \left\{ l \in \{h, h + 1, \dots, T - h\} : \hat{y}_{1,T}^{\text{MDL}}(l) = \max_{t \in (l-h, l+h)} \hat{y}_{1,T}^{\text{MDL}}(t) \right\}, \tag{17}$$

where h is the closest integer to $\ln T$ and $t \in (l - h, l + h)$ is the scanning window centered on l and browsing the entire $[1, T]$ segment. Since radius h is a function of T , the complexity of the method is at most $\mathcal{O}\left(\left(\frac{T}{\ln T}\right)^2\right)$. While this improvement with respect to $\mathcal{O}(T^2)$ seems small, the number of segments is drastically reduced in applied settings. For instance, for each simulated series of the six DGPs considered in the simulation section, the number of potential breaks never exceeds 8% of the sample size. Figure 2 illustrates the pruned version of the global method for a sample size of ten observations and potential locally estimated break at periods $\mathcal{J}^* = \{4, 8\}$. We call this pruned version of the GMDL method “PGMDL” and test its performance in the simulation section.

		Terminal date									
		1	2	3	4	5	6	7	8	9	10
Starting date	1	O	O	O	X	O	O	O	X	O	X
	2		O	O	X	O	O	O	X	O	X
	3			O	O	O	O	O	X	O	X
	4				O	O	O	O	X	O	X
	5					O	O	O	X	O	X
	6						O	O	X	O	X
	7							O	O	O	X
	8								O	O	X
	9									O	O
	10										O

Figure 2: Pruned global method—number of possible segments. The vertical number indicates the initial date of a segment and the horizontal number indicates the ending date. We set the sample size to $T = 10$ with a possible set of locally estimated breaks at periods $\mathcal{J}^* = \{4, 8\}$, and a maximal number of breaks set to $M = m^* = 2$. For instance, in the first row, possible segments are 1 to 4, 1 to 8 and 1 to 10, where the initial and in-between-break dates of the possible segments are denoted by O while the terminal dates are denoted by Xs. In the fourth row, possible segments are 4 to 8 and 4 to 10.

4 Simulation study

We assess the performance of the CP methods proposed in Section 3 in a simulation study. The goal is to: (i) compare the performance of the adapted CP methods proposed in Section 3 with existing ones, (ii) illustrate how we can capture the uncertainty related to the choice of a CP method and the

number of breaks for inference, and (iii) show the computational gain of combining local methods for detecting breakpoints in long time series. The following CP methods are considered:

1. BS: Binary segmentation, a local procedure implemented in the R package **wbsts** (Korkas and Fryzlewicz, 2017; Korkas and Fryzlewicz, 2020). We use function `wbs.lsw(...,M=1)` for detecting the breakpoints, keeping the default values for all other function arguments.
2. WBS: Wild binary segmentation, a local procedure presented in Korkas and Fryzlewicz (2017) and also implemented in the R package **wbsts**. Again, function `wbs.lsw(...,M=0)` is used, keeping the default values for all other function arguments.
3. LRSM: The likelihood ratio scan method, a local procedure proposed by Yau and Zhao (2016) and available at <http://wileyonlinelibrary.com/journal/rss-datasets>. We set the window size to $h = \max\{50, 2\log(T)^2\}$, the maximum AR order to 10, the minimum distance between two CPs to 50, and the minimum distance between the relative position of two CPs to 1%.
4. BSMDL: The binary segmentation procedure based on break detection criterion (14). Threshold δ_T is set to 3.
5. WBSMDL: The wild binary segmentation procedure based on break detection criterion (14). Again, threshold δ_T is set to 3.
6. PGMDL: The pruned global MDL method based on (17). We use Bai and Perron (2003)'s algorithm to optimize the MDL criterion given the set of potential breakpoints obtained with scan statistics (17). This is also a local method. The maximum number of breaks and the minimum regime duration are set to 50 and $10K$, respectively.
7. GMDL: The global MDL method as described under (15). This corresponds to Bai and Perron (2003)'s approach applied to the MDL criterion for a given number m of breaks. The upper bound of m and the minimum regime duration are set to 50 and $10K$, respectively.
8. SEL: This selection scheme chooses the method among local methods 1–6 that delivers the highest posterior probability according to (6).

The simulated data considered are the first six stationary AR linear processes of Yau and Zhao (2016, sec. 4), which are reproduced in Table 1. All DGPs are of autoregressive order up to two. DGP A is a stationary AR(1) model with no structural change. DGPs B and C are two standard piecewise-stationary AR processes with two breaks in the conditional mean occurring far from the start and end periods, and with constant variance. DGP D is a single-CP model with an early break (after 50 observations) in mean and constant variance. The last two DGPs, E and F, are strongly persistent stationary piecewise-linear AR processes with changing means and variances, for which structural change is typically more difficult to detect: the near-unit-root AR parameter of DGP E is constant across segments, while the autocorrelation function of DGP F does not change much across segments.

Table 1: Data generating processes.

DGP: $y_t = \beta_0 + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \epsilon_t$ where $\epsilon_t \sim iid\mathcal{N}(0, \sigma^2)$							
DGP	# breaks	Break location	β_0	β_1	β_2	σ^2	Description
A	0	{-}	{0}	{-0.7}	{0}	{1}	Stationary AR(1), no break
B	2	{514,768}	{0,0,0}	{0.9,1.69,1.32}	{0,-0.81,-0.81}	{1,1,1}	Piecewise-stationary AR(2), two breaks
C	2	{400,612}	{0,0,0}	{0.4,-0.6,0.5}	{0,0,0}	{1,1,1}	Piecewise-stationary AR(1), two breaks
D	1	{50}	{0,0}	{0.75,-0.5}	{0,0}	{1,1}	Piecewise-stationary AR(1), one early break
E	2	{400,750}	{0,0,0}	{0.999,0.999,0.999}	{0,0,0}	{1,2.25,1}	Piecewise-near-unit-root AR(1), two breaks
F	2	{400,750}	{0,0,0}	{1.399,0.999,0.699}	{-0.4,0,0.3}	{1,2.25,1}	Piecewise-near-unit-root AR(2), two breaks

NOTE: This table reports the DGPs investigated in Yau and Zhao (2016, sec. 4) and used in the simulation study.

We perform 1000 replications of the above DGPs and consider two metrics to evaluate the performance of the CP methods: (i) the proportion of replications for which a method identifies the true

number of breaks, and (ii) the “exact frequency” defined as the proportion of replications for which a method achieves both the correctly estimated number of breaks and the absolute distance between each pair of true and estimated break locations within 50 (see Yau and Zhao, 2016, sec. 4.2).

The performance of the eight CP methods for the six DGPs and samples of size $T = 1024$ is reported in Table 2. Overall, the methods are reliable, both in detecting the true number of breaks (in bold) and their locations: 87.5% of the reported experiments (42 out of 48) detect the true number of breaks more than 70% of the time, and 79.2% of them (38 out of 48) achieve that same level of accuracy in terms of exact frequency. BS and WBS clearly exhibit the worst overall performance. Note, however, that all local methods except BS and WBS detect the CPs assuming the true lag orders of the autoregressive processes.⁴ Given the latter, the results for BS and WBS are reasonably good and rather remarkable in the context of the near-unit-root processes E and F. In particular, BS and WBS beat the local methods LRSM and PGMDL in detecting and locating the breaks of process E.⁵ As well-known, structural breaks happening at the beginning or end of the sample are typically more difficult to detect with standard CP methods (see, *e.g.*, Horváth, Miller and Rice, 2020). In our simulation results, the adapted global and local methods accurately detect and locate DGP D’s early break. In particular, the performance of the SEL approach stresses the potential benefits of combining the local methods to detect early or late breakpoints. Note also that the local methods based on the MDL marginal likelihood compare favorably with other existing state-of-the-art methods (see, *e.g.*, Yau and Zhao, 2016, Table 3). Notice that WBSMDL outperforms all methods investigated in Table 2 (including the global method GMDL) for both performance metrics in all investigated DGPs. As expected, the global GMDL method is very reliable. For this particular Monte Carlo exercise, the SEL method performs equally well as the GMDL method.

Table 3 shows the average posterior probabilities of each CP method for 1000 replications of the DGPs. Considering the local CP methods first (left panel), we notice that the percentage of candidate CP methods with at least two posterior probabilities above the threshold of 10% —mixture metric “# mix”— is higher than 90% for each DGP. This indicates that no single local method strongly dominates the others in terms of posterior probabilities. Each local method captures relevant features of the simulated data. For instance, focusing on DGP A, BSMDL, WBDMDL, LRSM, and PGMDL are almost all equally valid methods in terms of posterior probabilities, and BS and WBS also display relatively large values for that metric. While model averaging over the potential number of breaks has already been investigated (see, *e.g.*, Maheu and Gordon, 2008), we will advocate for averaging across local methods to reduce the computational cost, improve CP detection, and account for model uncertainty. Recall from Table 2 that the SEL method appears to be an estimation strategy as accurate as Bay and Perron’s global method. Table 4 will further show that the computational gains with respect to the global method can be significant for long time series.

Regarding the global method, the right panel of Table 3 reports the average posterior probabilities for a set of possible number of breaks. As expected, the highest average posterior probability corresponds to GMDL running with the true number of breaks (in bold). The average posterior probabilities are much more favorable to the true model when multiple values for m are explored with a single global method than when multiple local methods are used. DGPs E and F display somewhat more diffuse posterior probabilities.

To underscore the usefulness of combining local CP methods in long time series, we simulate 100 series with DGP B, varying in size from $T = 2^{10} = 1024$ to $T = 2^{14} = 16,384$ by assuming breaks at dates $\tau_1 = 0.5T$ and $\tau_2 = 0.75T$. Table 4 summarizes the simulation results.

⁴Recall that all methods based on the MDL marginal likelihood require OLS estimates —the $\hat{\beta}_s$ s defined in (10) and in Proposition 2— for all relevant segments as an input. For example, criterion (14) used to compute BSMDL and WBSMDL makes use of the true lag orders of the DGPs in these OLS estimates. In comparison, the CPs obtained via BS or WBS rely on criterion (13), which does not require any particular DGP.

⁵Note that our LRSM simulation results are more favorable for DGP E and F than those reported in Yau and Zhao (2016, sec. 4). These authors rely on Yule-Walker instead of maximum likelihood estimators and omit the term $\ln^+(m)$ so that the optimal partitioning algorithm of Jackson et al. (2005) can be used.

Table 2: Results of the simulation study—performance of the CP methods.

DGP	Exact	# of detected breaks				Exact	# of detected breaks			
		$m=0$	$m=1$	$m=2$	$m \geq 3$		$m=0$	$m=1$	$m=2$	$m \geq 3$
BS method						WBS method				
A	48.3	48.3	20.8	19.8	11.1	78.8	78.8	16.8	3.9	0.5
B	33	0	14.1	51.6	34.3	27.5	0	29.1	53.6	17.3
C	83.8	0	0.1	88.6	11.3	84.7	0.4	0	92.3	7.3
D	58.8	12.9	66	13.6	7.5	69.7	21.2	73.7	4.9	0.2
E	66.9	0.3	3.9	90.6	5.2	73.6	5	1.1	89.9	4
F	61.8	0.2	10.4	82.2	7.2	73.3	0	5.8	89.3	4.9
BSMDL method						WBSMDL method				
A	100	100	0	0	0	100	100	0	0	0
B	98.5	0	0	100	0	99.7	0	0	100	0
C	99.7	0	0	100	0	100	0	0	100	0
D	71.2	2.9	97.1	0	0	100	0	100	0	0
E	69.6	7.3	4.2	86.5	2	84.6	3.4	1.8	93.4	1.4
F	79.2	0	3.6	94.9	1.5	92.5	0	1.3	97.1	1.6
LRSM method						PGMDL method				
A	100	100	0	0	0	95.3	95.3	3.9	0.8	0
B	99.4	0	0	99.9	0.1	95.1	0	0	95.5	4.5
C	96	0.4	0	96.4	3.2	95.4	0	0	95.5	4.5
D	99.8	0	99.8	0.2	0	95.2	0	95.2	3.8	1
E	55.4	32.3	3.4	62.8	1.5	42.5	0.1	0.6	48.2	51.1
F	91.6	0	2.5	96.2	1.3	70.3	0	0.3	73.7	26
GMDL method						SEL method				
A	100	100	0	0	0	100	100	0	0	0
B	99.7	0	0	100	0	99.7	0	0	100	0
C	100	0	0	100	0	99.8	0	0	100	0
D	99.9	0	99.9	0.1	0	99.9	0	99.9	0.1	0
E	81.7	2.5	1.1	89.8	6.6	82.7	2.9	1.8	92.1	3.2
F	90.1	0	0.7	94.4	4.9	92.9	0	1.1	96.9	2

NOTE: This table reports the results of the simulation study for the eight CP methods (six local, one global, and one selection method) for the DGPs of Table 1. All simulated series are of size $T = 1024$. Each series has been generated 1,000 times. For each run, the SEL method picks the local CP method with the highest posterior probability of the six local methods. The “ $m = \cdot$ ” columns report the proportion of total replications for which a method estimates m breaks. Bold values highlight the true number of breaks in the corresponding DGP. The “Exact” column reports the proportion of total replications for which a method infers the true number of breaks and achieves an absolute distance between each pair of true and estimated break within 50.

Table 3: Results of the simulation study—model posterior probabilities.

DGP	# mix	Local methods						GMDL method				
		BS	WBS	BSMDL	WBSMDL	LRSM	PGMDL	# mix	$m=0$	$m=1$	$m=2$	$m \geq 3$
A	100	8.2	14.4	19.6	19.6	19.6	18.5	0.2	99.8	0.2	0.0	0.0
B	98.7	0.2	0.3	16.1	7.7	38.2	37.5	1.0	0.0	0.0	99.6	0.4
C	96.1	2.1	2.4	17.9	7.8	27.7	42.1	1.7	0.0	0.0	99.2	0.8
D	98.6	0.7	0.3	0.0	13.3	44.6	41.1	1.1	0.0	99.5	0.5	0.0
E	93.5	3.8	6.1	14.6	23.8	30.1	21.6	29.2	2.8	2.0	84.1	11.1
F	94.8	1.6	1.8	8.7	13.4	43.9	30.6	13.3	0.0	1.0	91.6	7.5

NOTE: This table reports two metrics related to the model posterior probabilities of six local methods and the global one for the 1000 replications of the DGPs from Table 1. The “BS”, ..., “PGMDL”, and “ $m = \cdot$ ” columns show the average posterior probabilities for each DGP. Bold values for the GMDL method highlight the true number of breaks in the corresponding DGP. “# mix” designates the percentage of candidate CP methods with at least two posterior probabilities above the threshold of 10% for each DGP.

Based on the local methods’ computation time, BS and WBS are the fastest algorithms, followed by BSMDL, LRSM, WBSMDL, and PGMDL. It takes about 12 min on a standard computer (Intel Xeon Gold 6252 CPU at 2.1 GHz) for the slowest local algorithm to run for the largest sample size.

Table 4: Results of the simulation study—long time series.

T	Time	$m=m_0$	$ \hat{\tau}_1 - \tau_1 $	$ \hat{\tau}_2 - \tau_2 $	MDL	Time	$m=m_0$	$ \hat{\tau}_1 - \tau_1 $	$ \hat{\tau}_2 - \tau_2 $	MDL
BS method						WBS method				
2^{10}	0.00	57	11.37	39.40	-1534.43	0.00	63	11.48	70.38	-1584.66
2^{11}	0.00	34	10.00	51.29	-2982.35	0.00	57	10.23	72.21	-2978.50
2^{12}	0.01	17	9.82	43.76	-6014.65	0.01	60	11.67	59.07	-6013.11
2^{13}	0.01	19	14.37	48.42	-11682.57	0.01	57	9.26	44.65	-11664.00
2^{14}	0.03	11	19.73	26.00	-23395.69	0.03	54	13.22	40.91	-23321.18
BSMDL method						WBSMDL method				
2^{10}	0.00	99	6.52	3.33	-1500.85	0.19	95	6.64	3.35	-1500.85
2^{11}	0.01	100	4.63	2.93	-2967.02	0.49	96	4.50	2.89	-2967.02
2^{12}	0.02	100	5.57	3.33	-5979.61	1.31	96	5.44	3.41	-5979.61
2^{13}	0.07	100	5.64	3.45	-11636.49	3.99	98	5.78	3.51	-11633.85
2^{14}	0.22	100	4.74	4.14	-23308.41	12.17	99	4.74	4.15	-23308.41
LRSM method						PGMDL method				
2^{10}	0.01	100	8.26	5.41	-1502.28	0.04	99	5.68	5.20	-1502.94
2^{11}	0.05	100	9.43	7.79	-2967.51	0.12	100	4.62	3.74	-2968.27
2^{12}	0.17	100	9.33	7.01	-5986.78	0.49	100	4.92	3.81	-5985.29
2^{13}	0.64	100	10.08	8.17	-11642.16	2.22	100	5.00	4.46	-11634.74
2^{14}	2.45	100	10.42	10.89	-23310.94	11.58	100	4.70	5.31	-23309.88
GMDL method						SEL method				
2^{10}	0.45	99	5.34	3.66	-1500.85		99	5.97	3.36	-1500.85
2^{11}	2.12	100	3.96	2.90	-2966.46		100	4.39	3.00	-2967.02
2^{12}	11.85	100	4.44	3.32	-5978.67		100	5.06	3.27	-5979.61
2^{13}		Computationally too demanding					100	5.21	3.39	-11633.85
2^{14}		Computationally too demanding					100	4.16	4.14	-23308.41

NOTE: This table reports results of simulations performed with DGP B for a long time series by setting the relative locations at periods $\tau_1 = 0.5T$ and $\tau_2 = 0.75T$. The results are based on 100 replications. “Time” is the average estimation time (in minutes). “ $m=m_0$ ” is the percentage of replications in which the true number of regimes has been detected. “ $|\hat{\tau}_i - \tau_i|$ ” ($i = 1, 2$) is the average of the absolute value of the difference between the estimated CPs and the true one. “MDL” is the average estimated MDL value. The computational burden prevented the computation of the GMDL method on series with more than $T = 2^{12} = 4096$ observations.

Turning to the global method, GMDL’s algorithm takes more than eleven minutes to run for a sample of size $2^{12} = 4096$. For the largest sample sizes considered, GMDL is feasible but too demanding for the 100 replications required for our simulation study. Regarding the detection rates of the true number of CPs (Column “ $m=m_0$ ”), all approaches except BS and WBS deliver excellent results, as expected, given the consistency of the MDL criterion. The adapted methods proposed in Section 3 perform very well: The true breaks are detected in at least 95% of the runs. BS and WBS have the fastest computation time but also the poorest performance.

When assessing whether the estimation of the breaks is accurate in Columns “ $|\hat{\tau}_i - \tau_i|$ ” of Table 4, BSMDL, WBSMDL, and PGMDL are as reliable as the global method. Indeed, these local methods deliver the largest average MDL values, very close to the averaged exact MDL scores computed with the GMDL method (when feasible) and larger MDL scores as compared to the LRSM method (also based on the MDL criterion). These results naturally extend to the SEL method. Overall, for this particular DGP, BSMDL is the best method as it produces fast and very accurate results. Of course, BSMDL did not perform as well for DGP E in Table 2. That is why selecting the local method with the highest posterior probability is preferable in practice, as the true DGP is unknown.

5 Bayesian CP estimation and forecasting

To quantify the uncertainty of the estimated CPs, we propose a full Bayesian framework that accounts for information regarding the number and location of the CPs provided by an existing CP method.

Our approach differs in three ways from the Bayesian frameworks of Rigaiil, Lebarbier and Robin (2012) and Maheu and Song (2014), which also address break uncertainty and use \mathcal{NIG} priors. First, we condition on a given number of breaks while they do not. Second, we can handle a broader class of priors for the CPs, such as the geometric regime duration (as in Chib, 1998) or more complex distributions such as Poisson or negative binomial (as in Koop and Potter, 2007; Bauwens, Dufays and De Backer, 2014a). Third, we can use the estimated CPs to set an informative prior on the breakpoints. To illustrate the latter two features, we focus on a binomial distribution that is more flexible than a Poisson distribution due to its two hyperparameters and more informative than a negative binomial distribution since its variance is smaller than its expectation.

Given a vector $\hat{\tau}$ obtained from a CP method, we assume the following prior distribution:

$$\tilde{\tau}_i \sim \mathcal{TB}(r_i, e_i), \quad (18)$$

for $i = 1, \dots, m$, where $\mathcal{TB}(r_i, e_i)$ denotes a binomial distribution with parameters r_i and e_i , truncated on support $\left(\left\lceil \frac{\tilde{\tau}_{i-1} + \tilde{\tau}_i}{2} \right\rceil, \left\lfloor \frac{\tilde{\tau}_i + \tilde{\tau}_{i+1}}{2} \right\rfloor \right)$, where $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ are the ceiling and floor functions, respectively. The two parameters of the binomial distribution are calibrated such that $r_i = \left\lfloor \frac{\tilde{\tau}_i + \tilde{\tau}_{i+1}}{2} \right\rfloor$ and $E[\tilde{\tau}_i] = \hat{\tau}_i$. Defining $\tilde{\tau} = (\tilde{\tau}_1, \dots, \tilde{\tau}_m)'$, the Bayesian model is completed as follows:

$$\begin{aligned} y_t | \tilde{\tau}, \beta_i, \sigma_i^2 &\sim \mathcal{N}(\mathbf{x}'_t \beta_i, \sigma_i^2), \quad \text{for } \tilde{\tau}_{i-1} < t \leq \tilde{\tau}_i, \\ \beta_i | \tilde{\tau}, \sigma_i^2 &\sim \mathcal{N}(\hat{\beta}_i, \sigma_i^2 g_i (\mathbf{X}'_i \mathbf{X}_i)^{-1}), \\ \sigma_i^2 | \tilde{\tau} &\sim \mathcal{IG}\left(\frac{\nu_i}{2}, \frac{s_i}{2}\right), \end{aligned}$$

for $i = 1, \dots, m+1$, where $\hat{\beta}_i$ is the OLS estimates over segment $]\tilde{\tau}_{i-1}, \tilde{\tau}_i]$ with $\tilde{\tau}_0 = 0$, $\tilde{\tau}_{m+1} = T$, $\nu_i = \sqrt{n_i}$, $s_i = \frac{s_i}{\sqrt{n_i}}$, and $g_i = \left(\frac{((m+1)^{\frac{1}{m+1}}) n_i^{\frac{1}{4}} T}{(\frac{1}{\sqrt{n_i}} + 1)^{\frac{1}{2}}} \right)^{\frac{2}{K}} \exp\left(\frac{2}{K} \Delta R_{4,i}\right) n_i - 1$ with $n_i = \tilde{\tau}_i - \tilde{\tau}_{i-1}$.

As shown in Proposition 2, β_i and σ_i^2 can be integrated out. We can thus directly sample from the posterior distribution of the CPs conditioning on the data. Doing so significantly improves the efficiency of the sampler, as it rules out dependencies between the CPs and other model parameters. The posterior distribution of $\tilde{\tau}$ is given by:

$$f(\tilde{\tau} | \mathbf{y}_{1:T}) \propto f^{\text{MDL}}(\mathbf{y}_{1:T} | \tilde{\tau}) \prod_{i=1}^m f_B(\tilde{\tau}_i | r_i, e_i) \mathbf{1} \left\{ \left\lfloor \frac{\tilde{\tau}_{i-1} + \tilde{\tau}_i}{2} \right\rfloor \leq \tilde{\tau}_i \leq \left\lfloor \frac{\tilde{\tau}_i + \tilde{\tau}_{i+1}}{2} \right\rfloor \right\}, \quad (19)$$

where $f_B(x | r_i, e_i)$ denotes the density function of the binomial distribution with parameters r_i and e_i , and $\mathbf{1}\{\bullet\}$ is the indicator function. To simulate from this posterior, we use the D-DREAM algorithm (Bauwens, Dufays and De Backer, 2014a). Once the CPs have been sampled, direct sampling from the posterior distributions of β_i and σ_i^2 is straightforward (see the online Appendix, Section II).

One advantage of the Bayesian paradigm is the flexibility of the prior structure. In a CP framework, hierarchical priors have been introduced in Pesaran, Pettenuzzo and Timmermann (2006) to learn the in-sample break process and generate informed future breaks and parameters. We show below how to adapt our framework to account for one future break.⁶ We use the following prior distributions:

$$\begin{aligned} \tilde{\tau}_{m+1} | \tilde{\tau} &\sim \Lambda \equiv T + \mathcal{G}(r_{m+1}), \\ \beta_{m+2} | \hat{\beta}_{1:m+1} &\sim \mathcal{N}\left(\frac{1}{m+1} \sum_{i=1}^{m+1} \hat{\beta}_i, \Sigma_{\beta_{m+2}}\right), \\ \sigma_{m+2}^2 | \sigma_{1:m+1}^2 &\sim \mathcal{IG}\left(\frac{T}{2}, \frac{1}{m+1} \sum_{i=1}^{m+1} s_i\right), \end{aligned} \quad (20)$$

⁶The framework is easily extended to any number of future breaks, but it makes notations cumbersome.

where $\mathcal{G}(r_{m+1})$ denotes a geometric distribution with break probability $r_{m+1} = \frac{m}{T}$ and $\Sigma_{\beta_{m+2}}$ stands for a diagonal matrix with the empirical variance of the OLS estimates as diagonal elements.⁷ Note that the deterministic transformation Λ implies that the breakpoint occurs beyond the sample size. The out-of-sample prior specification in (20) is similar to the one proposed in Pesaran, Pettenuzzo and Timmermann (2006). Since the information set includes all available data, this prior can be easily adapted to account for specific features such as negative correlations between parameters over consecutive segments (see, *e.g.*, Pástor and Stambaugh, 2001) or persistence found in the OLS estimates.

6 Empirical applications

In this section, we propose two applications to show the practical relevance of our CP framework. The first application illustrates the computational advantage of using local CP methods when dealing with a very long time series. It stresses the ability of the MDL marginal likelihood to estimate and combine multiple CP models, accounting for break and parameter uncertainty. The second application is a forecasting exercise on the U.S. Consumer Price Index (CPI). It illustrates how combining local or global methods and incorporating an out-of-sample break can improve forecasting.

6.1 Combination of local CP methods in Fama-French factor models

We revisit the application in Horváth, Miller and Rice (2020) that seeks to detect structural changes in the parameters of the Fama-French five-factor model (Fama and French, 2015) applied to a portfolio of U.S. banking sector stocks. Our application considers a longer time horizon, namely daily returns over a period ranging from July 1st, 1963 to December 30th, 2022. The sample size of 14,979 observations is substantial. The five-factor model takes the following form :

$$R_t - R_{F,t} = a + b_{\text{MKT}}(R_{\text{MKT},t} - R_{F,t}) + b_{\text{SMB}}R_{\text{SMB},t} + b_{\text{HML}}R_{\text{HML},t} + b_{\text{RMW}}R_{\text{RMW},t} + b_{\text{CMA}}R_{\text{CMA},t} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2), \quad (21)$$

where R_t is the return of the bank portfolio at time t , $R_{F,t}$ is the risk-free rate, $R_{\text{MKT},t}$ is the market return, $R_{\text{SMB},t}$ is the return on a diversified portfolio of small stocks minus the return on a diversified portfolio of big stocks; $R_{\text{HML},t}$ is the return of a portfolio of stocks with a high book-to-market (B/M) ratio minus the return of a portfolio of stocks with a low B/M ratio; $R_{\text{RMW},t}$ is the returns of a portfolio of stocks with robust profitability minus a portfolio of stocks with weak profitability; and finally $R_{\text{CMA},t}$ is the return of a portfolio of stocks with conservative investment minus the return of a portfolio of stocks with aggressive investment.

Five-factor model (21) is expected to exhibit evolving risk exposures over time, and some factors may not be relevant throughout the entire period. Therefore, to detect the most likely factors that impact the bank portfolio, we consider the $2^5 = 32$ possible combinations of factors and apply our local CP methods to each model. This methodology is similar to that used in Meligkotsidou and Vrontos (2008) and Dufays, Houndetoungan and Coën (2022), which also exploit different sets of explanatory variables in a CP framework to infer hedge fund investment strategies.

For each of the 32 models, we apply 11 local CP methods. Specifically, we use the BSMDL and WBSMDL methods (both with five thresholds given by 0, 1, ..., 4) and the PGMDL method. The (wild) binary segmentation methods do not require specifying a maximum number of CPs, while we set this maximum to 99 for the PGMDL method. The estimation time of the local CP methods is at most 3 min (for the PGMDL) on a standard computer. Once the $32 \times 11 = 352$ models are estimated, we compare their MDL marginal likelihood and select the most probable model(s). The five most probable models are reported in Table 5 along with the number of detected breaks, the included risk

⁷We assume implicitly that the model has more than two breakpoints since we cannot learn about the break process with less information.

factors, the MDL marginal log-likelihoods, and the posterior probabilities. We can see that the bank portfolio is exposed to all five risk factors since the five best models select this set of factors (FF5) and their posterior probabilities sum to 100%. The method with the largest posterior probability (91.1%) identifies 38 CPs, the second-best method captures 39 ones, and so on.

Table 5: Best local CP models.

	Ranking				
	Best	2nd	3rd	4th	5th
Number of CPs	38	39	33	32	43
Set of Factors	FF5	FF5	FF5	FF5	FF5
MDL MLL	-10,510.90	-10,513.23	-10,518.30	-10,520.32	-10,527.02
Posterior Probability	91.12%	8.81%	0.06%	0.01%	0.00%

NOTE: This table reports the five best local CP models with respect to the MDL marginal log-likelihood. The top five models (number of CPs and set of factors) are obtained with the local CP methods over a total of 352 models (32 sets of factors times 11 local methods). FF5 stands for the five-factor model of Fama and French (2015).

Focusing on the model with the five factors, we apply the Bayesian setup of Section 5 to each of the CP methods, yielding a mixture posterior that accounts for both break and method uncertainty in the model parameters. Figure 3 displays the posterior distribution for the model parameters over the entire time period. Combining the local CP methods with many CPs results in many parameter shifts. The banking sector is shown to be positively exposed to the market factor (MKT) over the whole sample, with a level of exposure ranging from 0.5 to 2 and a clear increasing trend in the market exposure until the mid-90s. The exposure to the other factors is null, positive or negative depending on the regime. Notice that the 2007–2008 financial crisis that sparked the Great Recession impacted all risk factors quite heavily, while the COVID-19 pandemic occurring at the end of the sample seems also to be detected. According to the asset pricing theory (see, *e.g.*, Cochrane, 2005, chap. 12), the intercept (the so-called “alpha”) of a linear factor model—the five-factor model here—should be zero. Without CPs, this is what we find at the 1% significance level. Accounting for CPs, we notice many regimes in the top-left panel of Figure 3 in which this is also the case as the posterior distribution overlaps the horizontal dashed line at zero. However, we can also identify some periods for which the alpha is away from zero, for instance, at the beginning of the sample. Explaining occasional persistent deviations from zero for the intercept is beyond the scope of this illustrative application. Additional robustness checks showing that the detected CPs are not driven by the volatility dynamics only and that our choice of priors has no noticeable impact on the posterior of the model parameters can be found in the online Appendix, Section III.A.

Given the significant number of CPs identified, we can gain insights into the break process and make inference about model parameters accounting for a possible future break. For example, we can examine the 38 CPs detected by the most probable model and use prior (20) to infer the bank’s market exposure following a future break. Doing so yields a 95% credible interval of [0.69; 1.63] for b_{MKT} . Thus, the banking sector is likely to remain highly exposed to the market in the upcoming period even in the case in which a structural shift occurs.

To verify the validity of our local CP results, we also estimate the five-factor model with the global GMDL method with a maximum of 80 regimes. In Table 6, we report the five best models according to their MDL marginal log-likelihood. The estimation procedure takes approximately two hours to run and identifies a number of CPs ranging from 33 to 40, with a mode at 36. Focusing on the best GMDL model and its 36 detected breaks, we find that one out of 36 breaks is at a distance of 230 days, five breaks are within a distance of 32 days, while all remaining breaks are within a range of nine days or less to the CPs identified by the local methods. Figure 4 highlights the very goodmatch between the mixture of local CPs and the GMDL in terms of the posterior probability of CPs over time.




Figure 3: Mixture of marginal posterior distributions. The plots show the mixture of marginal posteriors of the Fama-French five-factor model parameters. These posteriors are computed with the full Bayesian setup of Section 5. The mixing weights are the posterior probabilities of the eleven CP local methods.

Table 6: Best GMDL models.

	Ranking				
	Best	2nd	3rd	4th	5th
Number of CPs	36	37	35	38	34
MDL MLL	-10,443.83	-10,444.22	-10,444.74	-10,445.47	-10,446.61
Posterior Probability	41.53%	28.19%	16.73%	8.12%	2.58%

NOTE: This table reports the five best GMDL models with respect to the MDL marginal log-likelihood for the five-factor model of Fama and French (2015). In the estimation, we set a maximum number of 80 regimes.

6.2 Forecasting the U.S. consumer price index

For the second illustration, we use our CP methods to forecast the monthly U.S. CPI. We consider the series $y_t = 1200 \times \frac{CPI_t - CPI_{t-1}}{CPI_t}$, where CPI_t denotes the CPI in month t . The series ranges from January 1947 to August 2022 and includes a total of 908 monthly observations. Following Groen, Paap and Ravazzolo (2013) and Check and Piger (2021) who find that short-lag CP-AR specifications are most appropriate to capture the dynamics of U.S. inflation in- and out-of-sample, we explore autoregressive specifications of order one and two with and without breaks and apply local and global CP detection methods. The interested reader can find the in-sample parameter dynamics of the CPI series for the CP-AR(2) estimation in the online Appendix, Section III.B. In this section, we focus on the out-of-




Figure 4: Change-point probability over time. The figure displays the CP probability over time obtained with the GMDL (in black) and the mixture of local CP methods (in gray).

sample forecasting exercise. We use an expanding window scheme for the model estimation, starting with 10% of the observations. Overall, we have 817 out-of-sample observations. We use the same 11 local CP methods used in the previous section. In addition, we apply the global GMDL method with a maximum of ten regimes. All models are then estimated using the Bayesian framework of Section 5 to generate forecasts at monthly horizons $h \in \{1, 3, 6, 12\}$.⁸

The root mean square forecast error (RMSFE) and mean absolute forecast error (MAFE) are reported in Table 7 for the various models and forecasting horizons. In gray, we highlight the specifications that belong to the superior set of models at the 95% confidence level for a given loss function and forecasting horizon (Hansen, Lunde and Nason, 2011). Overall, for both metrics, we notice improvements when we use local or global combination methods, in particular when we incorporate a future break in the modeling and consider forecasting over longer time horizons. Performance comparisons between higher-order AR specifications and the lower ones are reported in the online Appendix, Section III.B. Overall, the former ones do not beat the shorter AR specifications in terms of RMSFE and MAFE across forecasting horizons.

7 Conclusion

Many methods are available for detecting structural breaks and estimating the corresponding model parameters in each regime in linear regression models where the number and locations of the breaks are unknown. These methods mainly differ in the statistical framework under which they are developed (*i.e.*, Bayesian or frequentist), the fitting criterion used to adjust the data, and the algorithm's complexity used to optimize the fitting criterion. Even the most reliable global detection methods can

⁸To do so, for each period and each method, we simulate from the posterior predictive distribution and average over the draws to predict the next h observations. We generate 5000 MCMC draws (with a burn-in of 2000 draws). Afterward, we combine the predictions from the local or global methods using the posterior model probabilities computed from the MDL marginal log-likelihood.

Table 7: Forecasting results.

Model	RMSFE for horizon $h =$				MAFE for horizon $h =$			
	1	3	6	12	1	3	6	12
AR(1)	3.00	3.60	3.77	3.83	2.16	2.56	2.67	2.72
CP-AR(1) - Local	2.96	3.28	3.41	3.59	2.08	2.32	2.42	2.59
CP-AR(1) - Local - Future Break	2.95	3.28	3.40	3.57	2.08	2.32	2.41	2.57
CP-AR(1) - Global	2.95	3.24	3.40	3.57	2.07	2.30	2.40	2.54
CP-AR(1) - Global - Future Break	2.95	3.24	3.39	3.54	2.07	2.29	2.39	2.52
AR(2)	2.98	3.45	3.65	3.81	2.12	2.45	2.57	2.70
CP-AR(2) - Local	2.91	3.29	3.41	3.58	2.05	2.32	2.41	2.58
CP-AR(2) - Local - Future Break	2.91	3.28	3.40	3.56	2.06	2.32	2.41	2.56
CP-AR(2) - Global	2.94	3.29	3.40	3.58	2.07	2.31	2.40	2.57
CP-AR(2) - Global - Future Break	2.94	3.28	3.38	3.53	2.07	2.31	2.38	2.53

NOTE: This table reports the root mean square forecast error (RMSFE) and mean absolute forecasting error (MAFE) for the various model specifications and forecasting horizons. We highlight in gray the specifications belonging to the set of superior models at the 95% confidence level for a given loss function and forecasting horizon (Hansen, Lunde and Nason, 2011). We use the function `mcsTest` implemented in the R package `rugarch` (Ghalanos, 2022). The forecasting exercise is conducted over 817 out-of-sample observations.

deliver quite different numbers and locations of CPs when applied to real-world data. Moreover, these methods are computationally demanding for large samples.

This article addresses both issues by focusing on normal-linear regression models. We propose a Bayesian framework that allows researchers to account for model and break uncertainty and obtain accurate and reliable estimates for large datasets at a low computational cost. Our methodology relies on a marginal likelihood that exactly corresponds to the MDL criterion widely used in the frequentist literature for CP detection in linear regression. We establish that a class of hyperparameters for the \mathcal{NIG} prior distribution achieves such an equivalence. We further propose one global and three local CP methods that build on the probabilistic interpretation of the MDL criterion and on recent CP algorithms proposed in the literature. We also explore a full Bayesian model that delivers credible intervals for the CPs and other quantities of interest. Extensive experiments based on simulated data demonstrate the excellent performance of the proposed approaches to detect CPs relative to other state-of-the-art methods. Finally, we provide empirical illustrations that show the practical relevance of our approach when dealing with long time series and forecasting.

Future research could consider developing this framework for normal multivariate piecewise-linear regressions. For instance, Smith (2022) uses a Normal Inverse-Wishart prior distribution to obtain a closed-form expression of the marginal likelihood given the breakpoints. We could extend his framework by calibrating those prior hyperparameters to match a consistent information criterion and using the efficient breakpoint sampler proposed in this article.

Supplementary materials

The online Appendix contains the proofs of the propositions and remarks, the posteriors of the Bayesian estimation in Section 5, and additional analyses on the two empirical illustrations.

References

- Bai, J., 1997. Estimating multiple breaks one at a time. *Econometric theory* 13, 315–352.
- Bai, J., Perron, P., 1998. Estimating and testing linear models with multiple structure changes. *Econometrica* 66, 47–78.
- Bai, J., Perron, P., 2003. Computation and analysis of multiple structural change models. *Journal of Applied Econometrics* 18, 1–22.

- Bauwens, L., Dufays, A., De Backer, B., 2014a. Estimating and forecasting structural breaks in financial time series. *Journal of Empirical Finance* 29, 207–229.
- Bauwens, L., Dufays, A., Rombouts, J., 2014b. Marginal likelihood for Markov-switching and change-point GARCH models. *Journal of Econometrics* 178, 508–522.
- Bauwens, L., Koop, G., Korobilis, D., Rombouts, J., 2015. The contribution of structural break models to forecasting macroeconomic series. *Journal of Applied Econometrics* 30, 596–620.
- Bernardo, J., Bayarri, M., Berger, J., Dawid, A., Heckerman, D., Smith, A., West, M., 2007. Objective Bayesian analysis of multiple changepoints for linear models, in: *Bayesian statistics 8: proceedings of the eighth Valencia International Meeting, June 2-6, 2006*, Oxford University Press, USA. pp. 1–27.
- Carlin, B.P., Gelfand, A.E., Smith, A.F.M., 1992. Hierarchical Bayesian analysis of changepoint problems. *Journal of the Royal Statistical Society C* 41, 389–405.
- Chan, N., Yip, C., Zhang, R., 2014. Group LASSO for structural break time series. *Journal of the American Statistical Association* 109, 590–599.
- Check, A., Piger, J., 2021. Structural breaks in u.s. macroeconomic time series: A bayesian model averaging approach. *Journal of Money, Credit and Banking* 53, 1999–2036.
- Chib, S., 1998. Estimation and comparison of multiple change-point models. *Journal of Econometrics* 86, 221–241.
- Ciuperca, G., 2011. A general criterion to determine the number of change-points. *Statistics & Probability Letters* 81, 1267–1275.
- Cochrane, J.H., 2005. *Asset Pricing*. Princeton University Press.
- Davis, R.A., Lee, T.C.M., Rodriguez-Yam, G.A., 2006. Structural break estimation for nonstationary time series models. *Journal of the American Statistical Association* 101, 223–239.
- Davis, R.A., Yau, C.Y., et al., 2013. Consistency of minimum description length model selection for piecewise stationary time series models. *Electronic Journal of Statistics* 7, 381–411.
- Du, C., Kao, C.M., Kou, S.C., 2016. Stepwise signal extraction via marginal likelihood. *Journal of the American Statistical Association* 111, 314–330.
- Dufays, A., Houndetoungan, E.A., Coën, A., 2022. Selective linear segmentation for detecting relevant parameter changes. *Journal of Financial Econometrics* 116, 762–805.
- Eckley, I.A., Fearnhead, P., Killick, R., 2011. *Analysis of changepoint models*. Cambridge University Press, Cambridge. chapter 10. pp. 205–224.
- Fama, E., French, K., 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116, 1–22.
- Fernandez, C., Ley, E., Steel, M.F., 2001. Benchmark priors for Bayesian model averaging. *Journal of Econometrics* 100, 381–427.
- Fryzlewicz, P., 2014. Wild binary segmentation for multiple change-point detection. *Annals of Statistics* 42, 2243–2281.
- Geweke, J., Jiang, Y., 2011. Inference and prediction in a multiple-structural-break model. *Journal of Econometrics* 163, 172–185.
- Ghalanos, A., 2022. *rugarch: Univariate GARCH models*. R package version 1.4-9.
- Giordani, P., Kohn, R., 2008. Efficient Bayesian inference for multiple change-point and mixture innovation models. *Journal of Business & Economic Statistics* 26, 66–77.
- Groen, J.J.J., Paap, R., Ravazzolo, F., 2013. Real-time inflation forecasting in a changing world. *Journal of Business & Economic Statistics* 31, 29–44.
- Gupta, A., Chen, J., 1996. Detecting changes of mean in multidimensional normal sequences with applications to literature and geology. *Computational Statistics* 11, 211–221.
- Hansen, M.H., Yu, B., 2001. Model selection and the principle of minimum description length. *Journal of the American Statistical Association* 96, 746–774.
- Hansen, P.R., Lunde, A., Nason, J.M., 2011. The model confidence set. *Econometrica* 79, 453–497.
- Horváth, L., Miller, C., Rice, G., 2020. A new class of change point test statistics of rényi type. *Journal of Business & Economic Statistics* 38, 570–579.
- Jackson, B., Scargle, J.D., Barnes, D., Arabhi, S., Alt, A., Gioumoussis, P., Gwin, E., Sangtrakulcharoen, P., Tan, L., Tsai, T.T., 2005. An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters* 12, 105–108.
- Kass, R., Raftery, A., 1995. Bayes factors. *Journal of the American Statistical Association* 90, 773–795.

- Killick, R., Fearnhead, P., Eckley, I.A., 2012. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association* 107, 1590–1598.
- Koop, G., Potter, S., 2007. Estimation and forecasting with multiple breaks. *Review of Economic Studies* 74, 763–789.
- Korkas, K., Fryzlewicz, P., 2020. `wbsts`: Multiple change-point detection for nonstationary time series. R package version 2.1.
- Korkas, K.K., Fryzlewicz, P., 2017. Multiple change-point detection for non-stationary time series using wild binary segmentation. *Statistica Sinica* 27, 287–311.
- Li, S., Lund, R., 2012. Multiple changepoint detection via genetic algorithms. *Journal of Climate* 25, 674–686.
- Liu, J., Wu, S., Zidek, J.V., 1997. On segmented multivariate regressions. *Statistica Sinica* 7, 497–525.
- Maheu, J., Song, Y., 2014. A new structural break model, with an application to Canadian inflation forecasting. *International Journal of Forecasting* 30, 144–160.
- Maheu, J.M., Gordon, S., 2008. Learning, forecasting and structural breaks. *Journal of Applied Econometrics* 23, 553–583.
- Maheu, J.M., Song, Y., 2018. An efficient Bayesian approach to multiple structural change in multivariate time series. *Journal of Applied Econometrics* 33, 251–270.
- Meligkotsidou, L., Vrontos, I.D., 2008. Detecting structural breaks and identifying risk factors in hedge fund returns: A Bayesian approach. *Journal of Banking & Finance* 32, 2471–2481.
- Nemes, G., 2015. Error bounds and exponential improvements for the asymptotic expansions of the gamma function and its reciprocal. *Proceedings of the Royal Society of Edinburgh Section A: Mathematics* 145, 571–596.
- Ng, W.L., Pan, S., Yau, C.Y., 2022. Bootstrap inference for multiple change-points in time series. *Econometric Theory* 38, 752–792.
- Pástor, L., Stambaugh, R.F., 2001. The equity premium and structural breaks. *Journal of Finance* 56, 1207–1239.
- Pesaran, M.H., Pettenuzzo, D., Timmermann, A., 2006. Forecasting time series subject to multiple structural breaks. *Review of Economic Studies* 73, 1057–1084.
- Rigaill, G., Lebarbier, E., Robin, S., 2012. Exact posterior distributions and model selection criteria for multiple change-point detection problems. *Statistics and Computing* 22, 917–929.
- Safikhani, A., Shojaie, A., 2020. Joint structural break detection and parameter estimation in high-dimensional non-stationary VAR models. *Journal of the American Statistical Association* 117, 1–26.
- Smith, A.F., Spiegelhalter, D.J., 1980. Bayes factors and choice criteria for linear models. *Journal of the Royal Statistical Society B* 42, 213–220.
- Smith, S.C., 2022. Structural breaks in grouped heterogeneity. Forthcoming in the *Journal of Business & Economic Statistics*.
- Smith, S.C., Bulkley, G., Leslie, D.S., 2020. Equity premium forecasts with an unknown number of structural breaks. *Journal of Financial Econometrics* 18, 59–94.
- Stephens, D.A., 1994. Bayesian retrospective multiple-changepoint identification. *Applied Statistics* 1, 159–178.
- Yao, Y.C., Au, S.T., 1989. Least-squares estimation of a step function. *Sankhya: The Indian Journal of Statistics* 51, 370–381.
- Yau, C.Y., Zhao, Z., 2016. Inference for multiple change points in time series via likelihood ratio scan statistics. *Journal of the Royal Statistical Society B* 78, 895–916.
- Young Yang, T., Kuo, L., 2001. Bayesian binary segmentation procedure for a Poisson process with multiple changepoints. *Journal of Computational and Graphical Statistics* 10, 772–785.
- Zellner, A., 1986. On assessing prior distributions and Bayesian regression analysis with g-prior distributions, in: Goel, P., Zellner, A. (Eds.), *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*. Elsevier Science, New York, pp. 233–243.