axioms

MDPI

Article

Boundary-Regularized Bayesian Autoregressive Changepoint Detection with Applications to Natural Gas Markets

Jibin Yang [1,*], Maozai Tian [2] and Fuguo Liu [1]

[1] School of Mathematics and Data Science, Changji University, Changji 831100, China; lfg53880@cjc.edu.cn
[2] Center for Applied Statistics, School of Statistics, Renmin University of China, Beijing 100872, China; mztian@ruc.edu.cn
* Correspondence: 18124yjb@cjc.edu.cn or yjb24959@163.com

Abstract

Standard Bayesian autoregressive changepoint models can become unstable near sample boundaries. As a candidate changepoint approaches either edge of the series, the local residual degrees of freedom shrink, producing a Gamma-function singularity in the marginal likelihood that can strongly bias the posterior toward spurious edge detections. To address this issue, we introduce a regularization framework driven by local degrees of freedom. By incorporating a centripetal prior of the form $\pi(k) \propto (\nu_1\nu_2)^{\lambda}$—where $\nu_1 = k - 2p - 1$ and $\nu_2 = n - k - p - 1$—the proposed method is designed to counteract this boundary effect. Theoretical analysis shows that a regularization intensity of $\lambda \geq 1$ is sufficient to offset this boundary effect asymptotically. Simulation results confirm that this approach substantially mitigates the U-shaped error profile typical of unregularized estimators, yielding a more favorable accuracy–robustness trade-off relative to the standard frequentist baselines considered in our study. Finally, empirical applications to several 2022 natural gas benchmarks, including TTF, SHPGX LNG, JKM, NBP, and NYMEX Henry Hub, demonstrate the framework's ability to distinguish persistent structural transitions from transient market turbulence. These results suggest that degree-of-freedom-based centripetal prior regularization can improve the stability of Bayesian changepoint inference in nonstationary time series.

Keywords: bayesian changepoint detection; boundary singularity; autoregressive models; centripetal prior regularization; degrees of freedom; natural gas markets

MSC: 62M10; 62F15; 91B84

Check for updates

Academic Editors: Jong-Min Kim and Xiaonan Zhu

Received: 19 April 2026
Revised: 17 May 2026
Accepted: 19 May 2026
Published: 21 May 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Time series data arise in a wide range of fields, including environmental monitoring [1], epidemiological surveillance [2] and financial econometrics [3]. Broader statistical applications of changepoint methods have also been examined in recent studies [4]. Detecting these structural breaks is important for understanding regime shifts, improving forecasts, and informing policy decisions [5,6]. Among the many models used for this purpose, autoregressive (AR) changepoint models are particularly important because they explicitly capture temporal dependence [7,8]. However, reliable detection remains difficult in highly volatile or irregular settings, where heavy-tailed behavior and unstable local dynamics can affect autoregressive inference [9]. In particular, AR changepoint inference can become unstable when candidate changepoints are located near sequence boundaries [10].

Traditionally, changepoint detection methods can be broadly divided into frequentist cost-minimization approaches and Bayesian probabilistic inference approaches [11,12]. Many frequentist algorithms are computationally efficient in practice and are available in standard software implementations [13,14]. Representative examples include dynamic programming, which finds the optimal segmentation by evaluating admissible partitions, and binary segmentation, which recursively splits the sequence at the most prominent changepoint [14]. Despite their practical efficiency, these methods are often sensitive to boundary effects and do not naturally quantify uncertainty. To reduce failures near the sequence ends, they commonly rely on heuristic constraints such as a minimum segment length, introduced to avoid extremely short segments [15]. By contrast, Bayesian frameworks provide a natural mechanism for uncertainty quantification and allow prior information about changepoint locations to be incorporated into the inference [16–18].

Despite these advantages, Bayesian changepoint inference in AR models can suffer from boundary-related pathologies, including an explosion of the marginal likelihood near sequence boundaries [19–21]. Near the boundaries, the local sample size decreases and the residual degrees of freedom can approach zero. Under uninformative priors, such as the uniform prior, the marginalization step then develops a Gamma-function singularity. As a result, the posterior can become spuriously concentrated near edge locations, favoring artificial boundary changepoints over genuine structural changes in the interior of the sequence.

To address this issue, this paper introduces a centripetal prior regularization framework for Bayesian AR changepoint inference. Rather than relying on heuristic remedies such as domain truncation, the proposed approach directly targets the mathematical source of the boundary singularity. Specifically, we construct a data-adaptive prior based on the local residual degrees of freedom, which induces a centripetal penalty against degenerate boundary configurations. Through asymptotic analysis and Monte Carlo experiments, we show that this prior, with regularization level $\lambda \geq 1$, effectively neutralizes the boundary singularity while remaining essentially inert under clean data conditions. The resulting framework improves robustness against spurious changepoint signals and achieves a more favorable accuracy–robustness trade-off in highly stochastic environments.

Beyond controlled simulations, the proposed framework is also motivated by real-world time series with extreme jumps and time-varying local structures, as commonly observed in energy commodity markets. The 2022 European natural gas crisis and the associated global gas-market turbulence highlighted the need for robust changepoint detection under severe market turbulence [22,23]. During this period, the Dutch Title Transfer Facility (TTF), Europe's primary natural gas trading hub, experienced sustained price volatility driven by geopolitical shocks. Similar instability was also observed in other international natural gas benchmarks, including SHPGX LNG, JKM, NBP, and NYMEX Henry Hub. These disruptions generated substantial local noise and irregular boundary behavior, which can confound conventional linear and threshold-based estimators [24,25]. In this setting, the proposed Bayesian framework provides a suitable tool for distinguishing genuine regime shifts from spurious edge-driven fluctuations.

In summary, the main contributions of this paper are threefold:

1. **Theoretical characterization**: We identify the mathematical source of the boundary singularity in unregularized Bayesian AR changepoint models and provide an asymptotic analysis of its divergence mechanism.
2. **Methodological contribution**: We propose a degree-of-freedom-based centripetal prior that neutralizes boundary-induced posterior explosion without relying on heuristic masking rules.

3. **Empirical validation and real-world application**: Using Monte Carlo experiments and several 2022 natural gas benchmarks, including TTF, SHPGX, JKM, NBP, and NYMEX Henry Hub, we show that the proposed regularization framework improves robustness in volatile markets with asymmetric local shocks.

The remainder of this paper is organized as follows: Section 2 presents the methodological framework, including the boundary singularity analysis and the proposed centripetal prior. Section 3 reports the simulation studies and benchmark comparisons. Section 4 presents the empirical application to several 2022 natural gas benchmarks. Section 5 concludes the paper and discusses directions for future research.

2. Methodology

In this section, we present the theoretical framework underlying the proposed method. We begin by reformulating the Bayesian autoregressive (AR) changepoint model, making the residual degrees of freedom explicit. This formulation allows us to analyze the asymptotic behavior of the marginal likelihood near the sequence boundaries and to identify the mathematical source of the boundary explosion. We then introduce a degree-of-freedom-based centripetal prior and show how it regularizes the boundary behavior of the posterior distribution.

2.1. Standard Bayesian AR Changepoint Formulation

Consider a univariate time series $y = \{y_1, y_2, \ldots, y_n\}$, which is assumed to follow an autoregressive process of order p. Assume a structural break occurs at an unknown discrete timepoint k. The data are partitioned into two regimes. To account for the lag structure of the AR model, the effective response vector in the first regime is denoted by $Y_1 \in \mathbb{R}^{k_1}$ where $k_1 = k - p$, and that in the second regime is denoted by $Y_2 \in \mathbb{R}^{k_2}$ where $k_2 = n - k$.

The resulting piecewise linear regression model is given by

$$Y_i = X_i \beta_i + \epsilon_i, \quad i \in \{1, 2\} \tag{1}$$

where X_i is the corresponding design matrix of dimensions $k_i \times d$ ($d = p + 1$ including the intercept), $\beta_i \in \mathbb{R}^d$ is the coefficient vector, and $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2 I_{k_i})$.

To ensure model identifiability and the invertibility of $X_i^T X_i$, the number of effective observations must strictly exceed the number of parameters. We define the residual degrees of freedom for each regime as $\nu_i = k_i - d$. Thus, $\nu_1 = k - 2p - 1$ and $\nu_2 = n - k - p - 1$. To guarantee $\nu_i \geq 1$ for both regimes, the changepoint is constrained to the search domain $k \in [k_{min}, k_{max}]$ where $k_{min} = 2p + 2$ and $k_{max} = n - p - 2$.

Under the noninformative Jeffreys prior $\pi(\beta_i, \sigma_i^2) \propto \sigma_i^{-2}$, the closed-form log marginal likelihood $\mathcal{L}_i$ for regime i is evaluated as:

$$\mathcal{L}_i(\nu_i) = -\frac{1}{2} \log \left| X_i^T X_i \right| - \frac{\nu_i}{2} \log \left(\frac{SSR_i}{2} \right) + \log \Gamma \left(\frac{\nu_i}{2} \right) + C \tag{2}$$

where SSR_i denotes the sum of squared residuals. Under a uniform prior $\pi(k) \propto 1$, the overall objective function is obtained by summing the two regime-wise contributions: $\mathcal{L}(k) = \mathcal{L}_1(\nu_1) + \mathcal{L}_2(\nu_2)$. This requires a minimum sequence length of $n \geq 3p + 4$ to ensure $k_{\min} \leq k_{\max}$ and all $\nu_i \geq 1$.

2.2. Asymptotic Analysis of Boundary Singularity

Although a uniform prior on the changepoint location imposes no structural preference on k, it gives rise to a pronounced instability in the marginal likelihood near the sequence boundaries. To characterize this phenomenon, consider the asymptotic regime in which

the candidate changepoint approaches the left boundary, so that the residual degrees of freedom in the first regime satisfy $v_1 \to 0^+$. We now examine the limiting behavior of the terms in Equation (2):

1. Behavior of the SSR term: As $\nu_1 \to 0^+$, the projection matrix approaches identity, causing $SSR_1 \to 0^+$. However, this term is not the source of the divergence. Indeed, since $\lim_{x\to 0} x\log(x) = 0$, it follows that $-\frac{v_1}{2}\log(SSR_1) \to 0$. Therefore, the contribution of the residual term vanishes asymptotically and cannot explain the boundary explosion.
2. Singular contribution of the Gamma term: The divergence instead originates from the Gamma term in the marginal likelihood, which arises from integrating out the nuisance parameters. Using the classical asymptotic expansion $\Gamma(z) \sim \frac{1}{z}$ as $z \to 0^+$, we obtain:

$$\log\Gamma\left(\frac{\nu_1}{2}\right) \sim \log\left(\frac{2}{\nu_1}\right) = \log(2) - \log(\nu_1) \tag{3}$$

Hence, as $\nu_1 \to 0^+$, $\log\Gamma(\nu_1/2) \to +\infty$. Coupled with the potential ill-conditioning of the minimal covariance matrix $\left(-\frac{1}{2}\log\left|X_1^T X_1\right| \to +\infty\right)$, the total marginal likelihood exhibits a divergence of the form $\mathcal{L}_1 \sim -\log(\nu_1) \to +\infty$. Therefore, the boundary explosion is fundamentally driven by a normalization singularity induced by vanishing residual degrees of freedom, rather than by an actual improvement in local model fit. Under a standard uniform prior on k, this singularity is translated into excessive posterior mass near the boundaries, which in turn favors spurious edge changepoints.

2.3. Centripetal Prior as an Asymptotic Neutralizer

Motivated by the asymptotic analysis above, we regularize the prior on the changepoint location using the joint residual degrees of freedom of the two regimes. Rather than imposing heuristic truncation rules, this construction directly targets the boundary singularity identified in Section 2.2. Specifically, we define the normalized centripetal prior as:

$$\pi(k) = \frac{1}{Z}[\nu_1\nu_2]^\lambda = \frac{1}{Z}[(k-2p-1)(n-k-p-1)]^\lambda \tag{4}$$

where $Z = \sum_{k=k_{min}}^{k_{max}}[\nu_1\nu_2]^\lambda$ is the normalizing constant over the discrete search domain, and $\lambda > 0$ denotes the regularization strength. The corresponding regularized log-posterior is given by:

$$\log P(k|y) \propto \mathcal{L}(k) + \lambda\log(\nu_1) + \lambda\log(\nu_2) \tag{5}$$

Property 1: Boundary neutralization

Consider the same asymptotic boundary regime as in Section 2.2, namely $\nu_1 \to 0^+$. Under the centripetal prior, the additional term $\lambda\log(\nu_1)$ contributes a counterbalancing divergence. Since the singular part of the log-marginal likelihood is asymptotically of order $-\log(v_1)$, the combined boundary behavior satisfies:

$$\lim_{\nu_1\to 0^+}\Big(\underbrace{-\log(\nu_1)}_{\text{Gamma Singularity}} + \underbrace{\lambda\log(\nu_1)}_{\text{Prior Penalty}}\Big) = \lim_{\nu_1\to 0^+}(\lambda-1)\log(\nu_1) \tag{6}$$

Therefore, when $\lambda = 1$, the singularity is exactly canceled and the limit remains bounded; when $\lambda > 1$, the prior term dominates and drives the regularized log-posterior to $-\infty$. Hence, any choice $\lambda \geq 1$ eliminates the spurious boundary preference induced by the unregularized model.

Property 2: Flatness at the sequence center

A useful regularization should suppress artificial boundary attraction without distorting well-supported interior signals. To examine this behavior, consider a changepoint near the center of the sequence where $\nu_1 \approx \nu_2$. Treating k as a continuous variable for the purpose of local asymptotic interpretation and differentiating the log-prior with respect to k gives:

$$\frac{\partial}{\partial k} \log \pi(k) = \lambda \left(\frac{1}{\nu_1} - \frac{1}{\nu_2} \right) = \lambda \left(\frac{1}{k - 2p - 1} - \frac{1}{n - k - p - 1} \right) \tag{7}$$

At the midpoint, where $\nu_1 = \nu_2$, this derivative is zero. Thus, the centripetal prior is locally flat around the center of the sequence and does not introduce directional bias in regions where the information from the two regimes is balanced.

2.4. Implementation and MAP Estimation

Based on the regularized formulation developed above, changepoint estimation is carried out by combining the two regime-wise log-marginal likelihoods with the centripetal prior term. The resulting regularized log-posterior score for a changepoint located at position k is defined as

$$S(k, \lambda) = \mathcal{L}_1(\nu_1) + \mathcal{L}_2(\nu_2) + \lambda \log(\nu_1) + \lambda \log(\nu_2) + \text{const} \tag{8}$$

where $\mathcal{L}_i(\nu_i)$ is given by Equation (2), and the constant term includes the normalization factor $-\log Z$, which is independent of k and therefore irrelevant for optimization.

Given $\lambda \geq 1$ and the admissible search domain $k \in [k_{\min}, k_{\max}]$, the maximum a posteriori (MAP) estimator of the changepoint is defined by

$$\hat{k}_{\text{MAP}} = \arg \max_{k \in [k_{\min}, k_{\max}]} S(k, \lambda). \tag{9}$$

The implementation proceeds as follows:

1. Search range specification: Fix the AR order p and the regularization parameter $\lambda \geq 1$. Determine the admissible search range using $k_{\min} = 2p + 2$ and $k_{\max} = n - p - 2$.
2. Regime-wise likelihood evaluation: For each candidate changepoint $k \in [k_{\min}, k_{\max}]$, construct the corresponding design matrices X_1 and X_2, compute the residual sums of squares SSR_1 and SSR_2, and evaluate the log-marginal likelihoods $\mathcal{L}_1$ and $\mathcal{L}_2$ using the associated residual degrees of freedom ν_1 and ν_2.
3. Prior regularization: Add the centripetal prior contribution $\lambda\{\log(\nu_1) + \log(\nu_2)\}$ to obtain the regularized score $S(k, \lambda)$.
4. MAP optimization: Perform an exhaustive search over the discrete candidate set and select the value of k that maximizes $S(k, \lambda)$.

By incorporating the centripetal prior, the resulting MAP estimator avoids the artificial boundary preference that arises under a uniform prior on k. In contrast to the unregularized case, where the boundary singularity may dominate the posterior score, the proposed formulation yields changepoint estimates that are driven primarily by interior data evidence rather than by boundary-induced artifacts.

2.5. Extension to Multiple Changepoints

The theoretical development in this section is centered on the singular boundary phenomenon in the single-changepoint setting. Nevertheless, the proposed MAP estimator, denoted by $\hat{k}_{MAP}$, can be directly incorporated into standard recursive segmentation procedures—such as binary segmentation (BS) [26,27]—for multiple changepoint detection, without requiring any structural modification to the proposed prior. In this sense, the

centripetal prior serves as a modular regularization component that remains compatible with sequential partitioning strategies.

Although the present work focuses empirically on the single-changepoint regime, the structure of the proposed regularizer suggests two relevant implications when applied recursively:

- Implicit Suppression of Degenerate Micro-segments
  As recursive partitioning proceeds, the resulting child segments become progressively shorter, and the residual degrees of freedom (ν_1, ν_2) correspondingly decrease. As a consequence, the logarithmic penalty term $\lambda(\log \nu_1 + \log \nu_2)$ becomes increasingly influential relative to the bounded marginal likelihood component. This mechanism suggests that the proposed prior can act as an intrinsic regularizer against over-segmentation in weakly supported sub-intervals, reducing the tendency to generate degenerate micro-segments without the need for ad hoc stopping thresholds.
- Robustness Against Error Propagation
  A common limitation of recursive segmentation is that small localization errors in an early split may propagate to subsequent child segments. From the perspective of Property 2, the centripetal prior introduces no artificial directional preference when local sample support is balanced. Consequently, moderate deviations in a parent split are less likely to systematically distort the objective landscape of downstream searches. In contrast to uninformative priors, which may amplify artificial edge attraction in recursively shortened intervals, the proposed prior is expected to provide a stabilizing effect by discouraging spurious boundary solutions.

Overall, these observations indicate that the proposed regularization framework is not confined to the single-changepoint formulation, but can also be naturally extended to recursive multiple-changepoint settings. A dedicated empirical evaluation of the multiple-changepoint setting is deferred to future work.

3. Simulation and Benchmark Comparison

To assess the empirical implications of the theoretical analysis in Section 2, we conduct a controlled simulation study. The aim is to compare the proposed centripetal prior with both unregularized Bayesian inference and standard frequentist baselines in a setting that contains localized anomalies near the sequence boundary.

3.1. Simulation Setup

We simulate a univariate autoregressive time series of length $n = 150$. The sequence is designed to contain both a genuine structural break and a localized boundary perturbation to examine how different methods behave when the data include both an informative signal and boundary-induced instability.

- True Structural Break: A parameter shift is introduced at the midpoint of the series, $k_{\text{true}} = 75$, which serves as the target changepoint.
- Localized Boundary Perturbation: To create a finite-sample analog of the boundary-attraction mechanism analyzed in Section 2, we replace the data points in the interval $t \in [13, 21]$ with a highly oscillatory sequence. This construction produces a localized boundary anomaly that can generate spurious likelihood peaks near the sample edge.

3.2. Algorithmic Benchmarks

We compare the proposed method with two classes of changepoint detection approaches:

1. Frequentist baseline methods: We consider three standard frequentist algorithms based on global cost minimization, all implemented using the ruptures library in

Python 3.12: dynamic programming (Dynp), binary segmentation (BinSeg), and Bottom-Up [14]. All three methods minimize the same autoregressive cost function.

2. Bayesian inference under different priors: We evaluate the posterior probability $P(k \mid y)$ under three prior distributions corresponding to different values of the regularization hyperparameter λ in the centripetal prior $\pi(k) \propto [\nu_1 \nu_2]^{\lambda}$:
   - Uniform Prior ($\lambda = 0$): The standard baseline prior $\pi(k) \propto 1$.
   - Weak Centripetal Prior ($\lambda = 0.5$): An intermediate regularization setting included to assess the effect of partial boundary regularization.
   - Centripetal Prior ($\lambda = 1$): The theoretically motivated regularization level suggested by the asymptotic analysis in Section 2.

3.3. Simulation Results and Discussion

The estimation results for the illustrative example are summarized in Table 1, while Figure 1 provides a visual comparison of the estimated changepoints and posterior distributions. Together, these results illustrate how the choice of prior affects changepoint inference when the sequence contains a localized boundary anomaly.

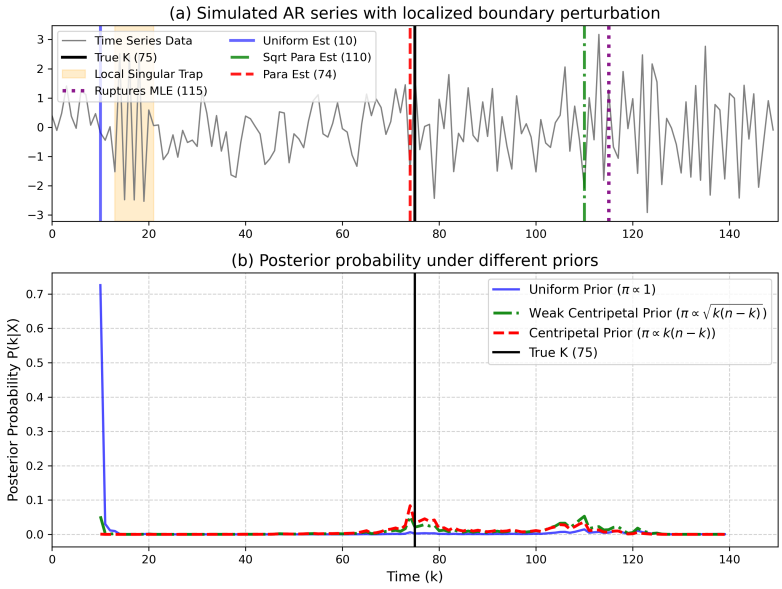


Figure 1. Performance of changepoint detection methods under a localized boundary perturbation. (**a**) Simulated AR series with a perturbed boundary segment (shaded yellow) and the estimated changepoints obtained by different methods. (**b**) Posterior probability distributions $P(k \mid y)$ under the three Bayesian priors considered in this study. The uniform prior places substantial posterior mass near the boundary; the weak centripetal prior reduces this concentration but retains a secondary interior mode, and the centripetal prior with $\lambda = 1$ concentrates near the true changepoint.

1. Frequentist baselines under localized boundary noise

Table 1 and the top panel of Figure 1 show that all three frequentist baseline methods selected $\hat{k} = 115$ rather than the true changepoint at $k_{\text{true}} = 75$, despite the use of minimal-size constraints. In this example, the global cost criterion is attracted to a spurious interior minimum generated by the perturbed segment, indicating sensitivity to local irregularities in the data.

2. Uniform prior under boundary singularity

Under the uniform prior, the posterior mode shifts to $\hat{k} = 10$, close to the left boundary. As shown in Figure 1b, the posterior mass becomes strongly concentrated in the boundary region. This behavior is consistent with the singular mechanism analyzed in Section 2,

where the marginal likelihood can become artificially inflated as the residual degrees of freedom approach zero.

3. Weak centripetal prior with partial regularization

Using $\lambda = 0.5$ reduces the extreme boundary attraction observed under the uniform prior, but does not fully eliminate the influence of the perturbation. The posterior distribution remains split between the true changepoint region and a secondary mode near $\hat{k} = 110$, resulting in a substantial estimation error. This intermediate outcome is consistent with the asymptotic analysis, which suggests that partial regularization may be insufficient to fully neutralize the boundary effect.

4. Centripetal prior at $\lambda = 1$

When $\lambda = 1$, the estimated changepoint is $\hat{k} = 74$, with an absolute error of 1. Compared with the other priors, the corresponding posterior distribution is more concentrated around the true changepoint and no longer exhibits the strong boundary preference seen under the uniform prior. In this example, the degree-of-freedom-based prior regularization provides the most accurate estimate among the methods considered.

Table 1. Boundary-perturbation example: changepoint detection performance ($k_{\text{true}} = 75$).

Framework	Algorithm/Prior	Constraint	Est. $\hat{K}$	Absolute Error
Ruptures	Dynp (AR Cost)	$min_size = 5$	115 *	40
	BinSeg (AR Cost)	$min_size = 5$	115 *	40
	BottomUp (AR Cost)	$min_size = 5$	115 *	40
Bayesian	Uniform ($\lambda = 0$)	None	10 †	65
	Weak Centripetal ($\lambda = 0.5$)	None	110	35
	Centripetal ($\lambda = 1$)	None	74	1

* Selected a spurious local optimum near the perturbed region. † Selected a changepoint near the left boundary, consistent with the singular behavior analyzed in Section 2.

3.4. Monte Carlo Analysis

To assess the performance of the proposed prior more systematically, we conduct a Monte Carlo study under two scenarios: a baseline setting with clean simulated data (Scenario A) and a stress-test setting with localized boundary noise (Scenario B). In both scenarios, the true changepoint location varies from the left side of the sequence ($k_{\text{true}} = 15$) to the right side ($k_{\text{true}} = 135$). For each location, we generate 50 independent AR(2) series of length $n = 150$. The resulting mean absolute errors (MAEs) are summarized in Figure 2. The Monte Carlo analysis focuses on the Bayesian specifications in order to isolate the effect of the prior regularization parameter λ.

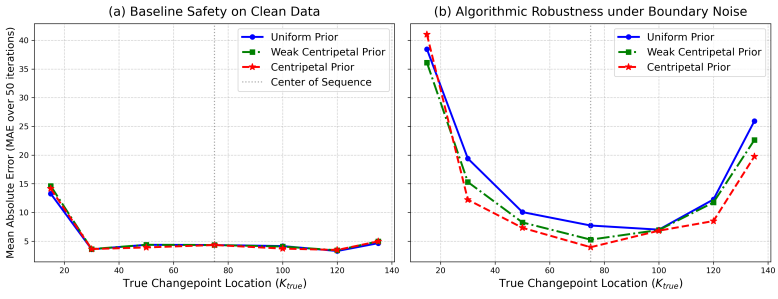


Figure 2. Monte Carlo comparison of changepoint estimation error across true changepoint locations (50 replications per location). (**a**) Mean absolute error under Scenario A (clean data). (**b**) Mean absolute error under Scenario B (localized boundary noise).

3.4.1. Performance on Clean Data

A natural concern when introducing prior regularization is the possibility of over-regularization, namely that the prior may distort changepoint inference even when the data are well behaved.

We examine this issue in Scenario A (Table 2), where the observations follow the underlying AR(2) model without additional anomalies. As shown in Figure 2a, the uniform prior ($\lambda = 0$), the weak centripetal prior ($\lambda = 0.5$), and the centripetal prior ($\lambda = 1$) produce very similar MAE curves. For example, at the sequence center ($k_{\text{true}} = 75$), all three methods attain an MAE of approximately 4.32. The close agreement across priors in Scenario A suggests that the proposed regularization does not introduce noticeable structural bias when the signal is clear.

Table 2. Mean absolute error under Scenario A: clean data.

True Location (k_{true})	Uniform Prior	Weak Centripetal Prior	Centripetal Prior
15	13.28	14.62	14.14
30	3.58	3.62	3.62
50	4.36	4.36	3.90
75 (Center)	4.32	4.32	4.32
100	4.14	4.06	3.70
120	3.28	3.40	3.46
135	4.64	4.96	4.96

3.4.2. Robustness to Boundary Noise

The clean data setting in Scenario A does not reflect situations in which localized irregularities appear near the sequence boundaries. To study such cases, Scenario B introduces boundary perturbations intended to mimic short anomalous segments near the beginning or end of the series. This setting is designed to test the sensitivity of changepoint inference to spurious likelihood spikes caused by limited boundary samples.

Figure 2b and Table 3 show that the three priors behave quite differently under boundary noise. The uniform prior exhibits substantially larger errors near the sequence edges, producing the characteristic U-shaped pattern in the MAE curve. For example, at $k_{\text{true}} = 30$, its MAE increases to 19.40, indicating strong sensitivity to spurious boundary effects.

Table 3. Mean absolute error under Scenario B: boundary-corrupted data.

True Location (k_{true})	Uniform Prior	Weak Centripetal Prior	Centripetal Prior
15	38.46	36.12	41.02
30	19.40	15.30	12.22
50	10.06	8.26	7.32
75 (Center)	7.72	5.26	3.94
100	7.00	6.94	6.82
120	12.26	11.74	8.50
135	25.92	22.62	19.76

By contrast, the centripetal prior reduces the MAE by 35% at $k_{\text{true}} = 30$ and by 49% at $k_{\text{true}} = 75$ relative to the uniform prior. Over most of the interior range ($k_{\text{true}} \in [30, 135]$), it yields the smallest MAE among the three methods, including 12.22 at $k_{\text{true}} = 30$ and 3.94 at $k_{\text{true}} = 75$. These results indicate that down-weighting boundary configurations can improve estimation stability in the presence of localized edge perturbations.

At the extreme boundary case $k_{\text{true}} = 15$, however, the centripetal prior yields a slightly larger MAE than the uniform prior (41.02 versus 38.46). This suggests a trade-off between boundary robustness and accuracy at the most extreme edge of the search domain. Nevertheless, across the majority of locations considered in Scenario B, the centripetal prior provides the most stable overall performance.

Overall, our simulation design directly targets the theoretical mechanism developed in Section 2. The boundary-perturbation example illustrates how localized edge anomalies can draw unregularized posterior mass toward weakly supported boundary configurations. The Monte Carlo experiments then test two key aspects of our hypothesis: robustness under boundary-corrupted data and the absence of significant bias under clean data. Importantly, these simulations are not intended as a forecasting exercise, but as a controlled test of the boundary-regularization mechanism introduced by our centripetal prior.

4. Case Studies in Natural Gas Markets

To examine how the proposed centripetal prior behaves in real data, we apply the Bayesian changepoint framework to five natural gas benchmarks: the Dutch Title Transfer Facility (TTF), the Shanghai Petroleum and Natural Gas Exchange LNG index (SHPGX), the Japan Korea Marker (JKM), the UK National Balancing Point (NBP), and the US NYMEX Henry Hub market. These markets differ in regional exposure, institutional setting, and dominant shock transmission channel, providing a useful setting for evaluating whether the proposed prior mitigates boundary-driven posterior concentration beyond a single benchmark.

To ensure comparability across markets, we use the same changepoint specification in each case and compare the same set of estimators: the uniform prior, the weak centripetal prior, the proposed centripetal prior, and the frequentist dynamic programming benchmark implemented in ruptures. Notably, the reference dates are not treated as true changepoints in a strict statistical sense; rather, they serve as external market-calendar benchmarks against which the localization behavior of different estimators can be compared.

4.1. Dutch TTF During the 2022 European Gas Crisis

Financial and commodity time series are often highly volatile and subject to short-lived extreme shocks. We consider the daily log returns of the Dutch TTF natural gas benchmark over the period from 1 May to 31 December 2022. This sample window shows pronounced volatility near the boundaries, reflecting residual market turbulence following the geopolitical disruptions earlier in 2022, as well as a major event in late August: the indefinite halt of the Nord Stream 1 pipeline on or around 31 August 2022. The combination of boundary volatility and a plausible mid-sample structural transition makes this period a useful test case for evaluating whether an estimation method is overly attracted to the sample boundary or instead identifies a changepoint aligned with the major August event.

Given the pronounced boundary volatility in the TTF series, the proposed centripetal prior is implemented with $\lambda = 1.5$ in the empirical application. We analyze the TTF log returns using an AR(2)-based changepoint model under four configurations: (1) a uniform prior as the baseline; (2) the proposed centripetal prior; (3) a weak centripetal prior as an ablation setting; and (4) a frequentist dynamic programming approach implemented through ruptures.

The estimated changepoints and posterior distributions are shown in Figure 3, and the absolute deviations from the reference event date are summarized in Table 4.

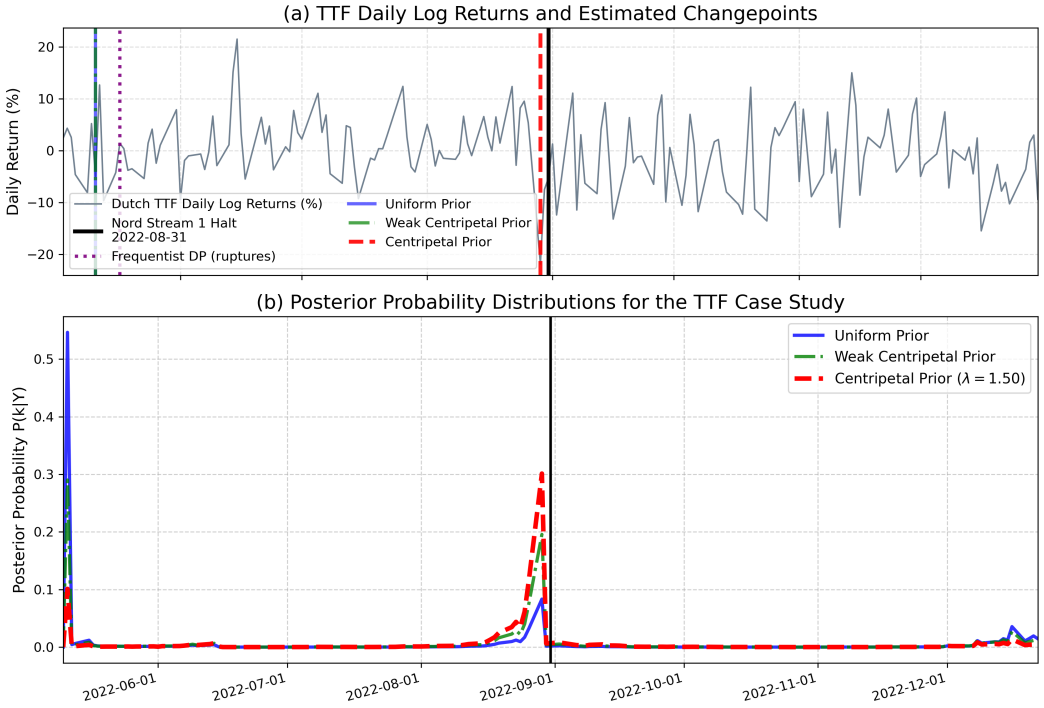


Figure 3. Case study based on the 2022 Dutch TTF log returns. (**a**) Daily log returns together with the changepoint estimates produced by the methods under comparison. (**b**) Posterior probability distributions under the Bayesian priors considered in this study. The uniform and weak centripetal priors place substantial posterior mass near the left boundary, whereas the stronger centripetal prior concentrates near the late-August event.

Table 4. Changepoint estimates for the TTF market.

Method	Prior Distribution	Est. Date	Abs. Error (Days)
Reference Event	–	31 Aug 2022	**0**
Uniform Prior (Baseline)	$\pi(k) \propto 1$	11 May 2022	112
Weak Centripetal Prior	$\pi(k) \propto [v_1 v_2]^{1/2}$	11 May 2022	112
Frequentist DP (Ruptures)	–	17 May 2022	106
Centripetal Prior (Proposed)	$\pi(k) \propto [v_1 v_2]^{3/2}$	29 Aug 2022	2

As shown in Figure 3 and Table 4, the uniform prior and the frequentist dynamic programming method both place the estimated changepoint near the left boundary of the sample (mid-May), rather than near the late-August reference event. The resulting deviations exceed 100 days. This pattern is consistent with the theoretical mechanism discussed earlier: without sufficient centripetal regularization, the likelihood can be overly influenced by localized edge volatility.

An additional observation concerns the weak centripetal prior. Although it favors interior locations relative to the uniform prior, its estimate still falls within the mid-May boundary region. This behavior can be interpreted in terms of penalty strength. In log-space, the weak centripetal prior contributes a penalty term proportional to $0.5 \log(v_1 v_2)$. In the present application, this level of regularization appears too weak to offset the concentration of likelihood near the sample boundary.

Given the pronounced empirical volatility of the TTF series, a stronger regularization level was used in this application. Under this specification, the estimated changepoint shifts from the left boundary to the late-August transition period, with an absolute deviation of two days from the reference event. Taken together, these results suggest that

merely adopting a mildly non-uniform prior may be insufficient in highly volatile real-world settings, whereas stronger centripetal regularization can improve robustness against boundary concentration.

The late-August changepoint identified by the centripetal prior ($\hat{k}$ = 29 August 2022) is economically interpretable and may be informative for understanding the timing of regime change in the European natural gas market during 2022. Throughout 2022, the European natural gas market was highly fragmented due to a series of geopolitical announcements [28,29]. In such an environment, early-sample volatility may be difficult to distinguish from more persistent structural change. As suggested by the mid-May estimates of the baseline methods (the "boundary trap"), procedures that are insufficiently regularized can place excessive weight on transient market disturbances near the beginning of the sample [30,31].

However, recent empirical studies in energy economics often associate the major market transition in 2022 with the indefinite shutdown of the Nord Stream 1 pipeline at the end of August, when the pricing mechanism of TTF moved further away from pipeline-dominated supply constraints and toward a stronger dependence on liquefied natural gas (LNG) integration.

The changepoint estimate obtained under the centripetal prior is broadly consistent with this interpretation. From this perspective, the May fluctuations reflect temporary volatility driven by uncertainty. In contrast, the August 29 estimate is plausibly associated with a more persistent market transition. This suggests that probabilistic changepoint models with appropriate regularization may be useful for distinguishing short-lived turbulence from more persistent regime transitions in energy-market risk analysis during periods of systemic stress [32].

4.2. External Validation: SHPGX LNG

We next examine the SHPGX China LNG price index as an external validation case. Unlike the TTF benchmark, SHPGX reflects an LNG market outside Europe, with different institutional features and demand-side seasonality [24]. It therefore provides a useful setting for assessing whether the proposed prior remains informative outside the European pipeline gas market.

We use daily log returns from 1 April 2022 to 28 February 2023. The sample shows a pronounced early fluctuation in June 2022 and a later transition period around late October, when the Chinese LNG market entered the winter heating and supply security season. We use 31 October 2022 as an economically motivated reference date. As in the TTF case, this date is treated as an external market-calendar benchmark rather than a known statistical changepoint.

The same AR(2)-based changepoint specification is applied to the SHPGX log-return series. As in the TTF case study, we compare the uniform prior, the weak centripetal prior, the proposed centripetal prior, and the frequentist dynamic programming benchmark implemented through `ruptures`. The empirical results are reported in Figure 4 and Table 5.

As shown in Table 5, both the uniform prior and the weak centripetal prior select 13 June 2022, whereas the frequentist dynamic programming benchmark selects 23 June 2022. These estimates are located close to the early-sample volatility episode and are more than 100 days away from the late-October reference date. The posterior distributions in Figure 4 show the same pattern: under the uniform and weak centripetal priors, posterior mass is concentrated around the June fluctuation.

By contrast, the proposed centripetal prior estimates the changepoint to 27 October 2022, only four days before the reference event. The corresponding posterior distribution shifts away from the early-sample likelihood peak and concentrates around the late-October

transition region. This behavior is consistent with the theoretical mechanism discussed in Section 2: when a localized fluctuation occurs near the boundary, the unregularized marginal likelihood can be inflated by limited local degrees of freedom, whereas the centripetal prior penalizes such weakly supported boundary configurations through the factor $[\nu_1\nu_2]^{\lambda}$.

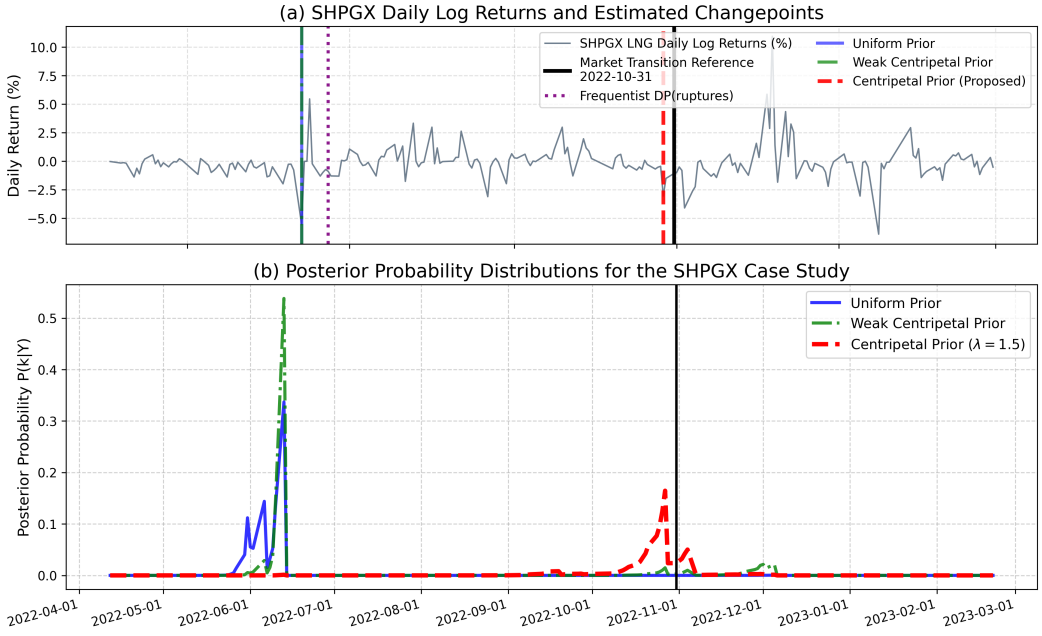


Figure 4. Changepoint estimates and posterior distributions for the SHPGX LNG price index. Panel (**a**) shows daily log returns and estimated changepoints. Panel (**b**) reports posterior probabilities under the Bayesian priors. The uniform and weak centripetal priors concentrate near the June fluctuation, whereas the proposed centripetal prior shifts the posterior mode toward the late-October transition period.

Table 5. Changepoint estimates for the SHPGX LNG price index.

Method	Prior Distribution	Est. Date	Abs. Error (Days)
Reference Event	–	31 Oct 2022	0
Uniform Prior (Baseline)	$\pi(k) \propto 1$	13 Jun 2022	140
Weak Centripetal Prior	$\pi(k) \propto [\nu_1\nu_2]^{1/2}$	13 Jun 2022	140
Frequentist DP (ruptures)	–	23 Jun 2022	130
Centripetal Prior (Proposed)	$\pi(k) \propto [\nu_1\nu_2]^{3/2}$	27 Oct 2022	4

In this sense, the SHPGX case provides a useful external check on the TTF results. It is consistent with the view that degree-of-freedom-based regularization improves the stability of changepoint localization in volatile energy-market series.

4.3. Cross-Market Robustness: JKM, NBP, and NYMEX

4.3.1. Market Selection and Reference Events

To further assess whether the empirical behavior observed in the TTF and SHPGX cases is specific to those two benchmarks, we consider three additional natural gas markets: the Japan Korea Marker (JKM), the UK National Balancing Point (NBP), and the NYMEX Henry Hub natural gas benchmark [31,33]. JKM reflects the Northeast Asian LNG market, NBP is closely connected to the European gas system, and NYMEX is more directly linked to the US domestic supply and LNG export conditions. They therefore provide a useful

setting for examining whether the proposed prior continues to reduce boundary-driven posterior concentration across heterogeneous market environments.

For comparability, all three series are analyzed over the common window from 1 March to 31 December 2022. Daily log returns are constructed from the corresponding price series, and the same AR(2)-based changepoint specification is applied throughout. The reference dates are chosen to reflect market-calendar events relevant to each benchmark: 26 August 2022 for JKM, corresponding to the late-summer LNG price surge; 31 August 2022 for NBP, corresponding to the late-August escalation in the European gas crisis; and 8 June 2022 for NYMEX, corresponding to the Freeport LNG facility explosion. As in the preceding cases, these dates are used as external benchmarks rather than as known statistical changepoints. The proposed centripetal prior is implemented with $\lambda = 1.5$.

4.3.2. Results and Cross-Market Interpretation

The numerical results are reported in Table 6, and the corresponding time series plots and posterior distributions are shown in Figure 5. Across the three markets, the uniform prior and the weak centripetal prior tend to place their posterior modes close to early-sample fluctuations. This pattern is most evident in JKM and NBP, where both priors select 10 March 2022. In the JKM case, this leads to an absolute error of 169 days relative to the August 26 reference date. The frequentist dynamic programming benchmark selects 19 May 2022, reducing the error to 99 days but remaining well before the late-summer transition period. By contrast, the proposed centripetal prior estimates the JKM changepoint as 28 September 2022, with an absolute error of 33 days. This estimate does not exactly coincide with the reference date, but it shifts the posterior mode away from the March boundary region and toward the post-peak adjustment period in the LNG market.

The NBP case shows a similar boundary-attraction pattern, although the proposed estimate is less tightly aligned with the reference event than in the TTF or SHPGX cases. The uniform prior and weak centripetal prior both select 10 March 2022, and the frequentist dynamic programming benchmark selects 16 March 2022. These estimates are close to the beginning of the sample and produce errors exceeding 160 days. The proposed centripetal prior shifts the estimated changepoint to 31 October 2022, reducing the absolute error to 61 days. This result should not be read as an exact dating of the late-August event. Rather, it suggests that, in a highly volatile European gas benchmark, the proposed prior moves posterior mass away from weakly supported boundary configurations and toward a later interior transition region.

The NYMEX case aligns more closely with the market reference date. The uniform prior and the weak centripetal prior both select 19 April 2022, which is 50 days before the Freeport LNG event. The frequentist dynamic programming benchmark selects 16 December 2022, producing a much larger error. In contrast, the proposed centripetal prior estimates the changepoint as 10 June 2022, only two days after the reference event. The posterior distribution in Figure 5 also shows that the proposed prior reallocates mass away from the early-sample peak and toward the economically relevant June region.

Taken together with the TTF and SHPGX evidence, these cross-market results are consistent with the interpretation developed in the preceding sections. The proposed prior does not impose a common changepoint date across markets: the estimated dates differ across JKM, NBP, and NYMEX, reflecting the different timings and transmission channels of regional gas-market shocks. At the same time, the posterior mass under the proposed prior is less concentrated near the sample boundary than under the uniform and weak centripetal priors. The cross-market evidence therefore provides an additional robustness check, suggesting that degree-of-freedom-based centripetal regularization can improve the

stability of changepoint localization in volatile energy-market series without forcing all markets into the same calendar interpretation.

Table 6. Changepoint estimates for the JKM, NBP, and NYMEX markets.

Market	Method	Prior Distribution	Est. Date	Abs. Error (Days)
JKM	Reference Event	–	26 Aug 2022	0
JKM	Uniform Prior	$\pi(k) \propto 1$	10 Mar 2022	169
JKM	Weak Centripetal Prior	$\pi(k) \propto [\nu_1 \nu_2]^{1/2}$	10 Mar 2022	169
JKM	Frequentist DP (ruptures)	–	19 May 2022	99
JKM	Centripetal Prior (Proposed)	$\pi(k) \propto [\nu_1 \nu_2]^{3/2}$	28 Sep 2022	33
NBP	Reference Event	–	31 Aug 2022	0
NBP	Uniform Prior	$\pi(k) \propto 1$	10 Mar 2022	174
NBP	Weak Centripetal Prior	$\pi(k) \propto [\nu_1 \nu_2]^{1/2}$	10 Mar 2022	174
NBP	Frequentist DP (ruptures)	–	16 Mar 2022	168
NBP	Centripetal Prior (Proposed)	$\pi(k) \propto [\nu_1 \nu_2]^{3/2}$	31 Oct 2022	61
NYMEX	Reference Event	–	08 Jun 2022	0
NYMEX	Uniform Prior	$\pi(k) \propto 1$	19 Apr 2022	50
NYMEX	Weak Centripetal Prior	$\pi(k) \propto [\nu_1 \nu_2]^{1/2}$	19 Apr 2022	50
NYMEX	Frequentist DP (`ruptures`)	–	16 Dec 2022	191
NYMEX	Centripetal Prior (Proposed)	$\pi(k) \propto [\nu_1 \nu_2]^{3/2}$	10 Jun 2022	2

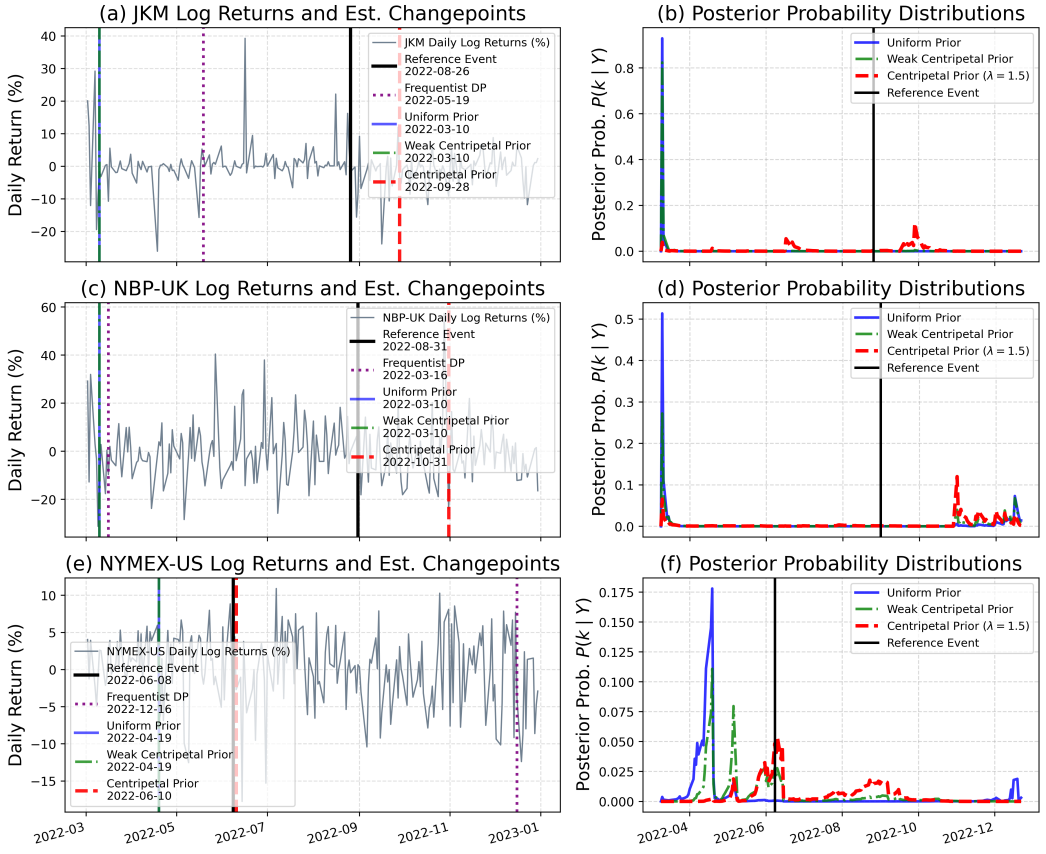


Figure 5. Cross-market robustness results for JKM, NBP, and NYMEX. Panels (**a**,**c**,**e**) show daily log returns and estimated changepoints. Panels (**b**,**d**,**f**) report posterior probability distributions under the Bayesian priors. The proposed centripetal prior generally reduces early-sample posterior concentration and shifts the posterior mode toward market-specific interior transition regions.

5. Conclusions

This paper investigated the boundary-trap problem in Bayesian autoregressive changepoint inference. We demonstrated that when a candidate changepoint approaches either

end of a time series, the shrinking residual degrees of freedom induce a singularity in the marginal likelihood. This mathematical artifact causes unregularized Bayesian procedures to exhibit excessive posterior concentration near sample boundaries, often leading to the misidentification of local noise or transient disturbances as genuine structural breaks.

To resolve this instability, we developed a Bayesian framework incorporating centripetal prior regularization. The proposed prior, $\pi(k) \propto (\nu_1\nu_2)^\lambda$, is constructed directly from the local degrees of freedom, allowing the regularization strength to adapt to the model's geometric structure. Theoretical analysis confirmed that setting $\lambda \geq 1$ is sufficient to neutralize the boundary likelihood explosion asymptotically. Unlike heuristic trimming methods, this approach provides a principled way to suppress spurious edge detections while preserving all available observations for inference.

Simulation results across various scenarios supported these theoretical findings. Compared to standard frequentist benchmarks—such as dynamic programming and binary segmentation—the centripetal prior improved localization accuracy in boundary-trap settings and reduced the characteristic U-shaped error profile near sample edges. Importantly, Monte Carlo experiments showed that the regularization remains effectively neutral on clean data, suggesting that no unnecessary structural bias is introduced when the underlying signal is well behaved.

The practical utility of the framework was further demonstrated through a series of empirical applications to 2022 natural gas markets. In the primary TTF case study, the proposed method helped distinguish persistent regime shifts from the extreme volatility characterizing the European energy crisis. While unregularized estimators were attracted to the sample boundary, the centripetal prior identified a changepoint closely aligned with the late-August Nord Stream 1 shutdown. To evaluate the empirical generality of these results, the analysis was extended to several other international gas benchmarks, including China's SHPGX LNG index, as well as the JKM, NBP, and NYMEX Henry Hub. Across these heterogeneous trading hubs, the centripetal prior generally reduced boundary attraction and yielded changepoint estimates that were consistent with major geopolitical and supply-side developments. These findings suggest that degree-of-freedom-based regularization improves the stability and reliability of Bayesian inference in highly volatile real-world environments.

More broadly, this study offers a theoretically grounded solution to the boundary singularity problem in Bayesian changepoint inference. The proposed framework helps distinguish short-lived anomalies from persistent regime transitions, particularly in time series characterized by heavy tails and elevated uncertainty. Although the empirical analysis focused on natural gas markets, the same principle may apply to other commodity and energy-market time series, as well as to a wider range of financial, economic, and engineering applications.

Several directions for future research remain open. First, extending the degree-of-freedom-based regularization principle to multivariate models, such as VAR and dynamic factor frameworks, may allow joint detection of synchronous and asynchronous structural changes across multiple variables. Second, incorporating centripetal prior regularization into conditional heteroskedastic models, including GARCH and stochastic volatility specifications, may further improve the method's ability to handle volatility clustering in high-frequency and extreme-volatility settings. Finally, evaluating the performance of the proposed prior in multiple-changepoint settings remains a valuable direction for future work.

Author Contributions: Conceptualization, J.Y.; Methodology, J.Y.; Software, J.Y. and M.T.; Validation, J.Y.; Formal analysis, J.Y.; Writing—original draft, J.Y.; Writing—review & editing, J.Y., M.T. and

F.L.; Visualization, J.Y.; Supervision, M.T. and F.L.; Project administration, M.T. and F.L.; Funding acquisition, J.Y. and F.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by three projects: the Platform Construction Project of the Key Laboratory of Financial Big Data, Changji University, grant number Changyuan Xiaofa [2024] No. 46; the Construction Project for the Proposed Master's Degree Program in Applied Statistics, Changji University, grant number Changyuan Xiaofa [2025] No. 10; and the 2024 University-Level Scientific Research Project of Changji University, grant number KY2024039.

Data Availability Statement: The simulated data and code used in this study are available from the corresponding author upon reasonable request. The empirical data consist of several natural gas market price series, including Dutch TTF, SHPGX LNG, JKM, NBP, and NYMEX Henry Hub, obtained from Investing.com and are subject to the terms and conditions of the data provider. These third-party data cannot be publicly redistributed by the authors, but can be accessed directly from the source website: https://cn.investing.com/, accessed on 18 March 2026.

Acknowledgments: Generative AI was used solely for translation, language polishing, and grammar correction to improve the clarity of the manuscript. No part of the scientific content, data analysis, or conclusions was generated by AI. The authors take full responsibility for the content of the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Reeves, J.; Chen, J.; Wang, X.L.; Lund, R.; Lu, Q.Q. A review and comparison of changepoint detection techniques for climate data. *J. Appl. Meteorol. Climatol.* **2007**, *46*, 900–915. [CrossRef]
2. Jegede, S.L.; Szajowski, K.J. Change-Point Detection in Homogeneous Segments of COVID-19 Daily Infection. *Axioms* **2022**, *11*, 213. [CrossRef]
3. Wang, J.; Zivot, E. A Bayesian time series model of multiple structural changes in level, trend, and variance. *J. Bus. Econ. Stat.* **2000**, *18*, 374–386. [CrossRef]
4. Tian, F.; Qi, Y.; Wang, Y.; Tian, B. Study of Change-Point Detection and Applications Based on Several Statistical Methods. *Symmetry* **2025**, *17*, 302. [CrossRef]
5. Truong, C.; Oudre, L.; Vayatis, N. Selective review of offline change point detection methods. *Signal Process.* **2020**, *167*, 107299. [CrossRef]
6. Aminikhanghahi, S.; Cook, D.J. A survey of methods for time series change point detection. *Knowl. Inf. Syst.* **2017**, *51*, 339–367. [CrossRef]
7. Davis, R.A.; Lee, T.C.M.; Rodriguez-Yam, G.A. Structural break estimation for nonstationary time series models. *J. Am. Stat. Assoc.* **2006**, *101*, 223–239. [CrossRef]
8. Sheikhrabori, R.; Aminnayeri, M. Maximum likelihood estimation of the change point in stationary state of auto regressive moving average (ARMA) models, using SVD-based smoothing. *Commun. Stat. Theory Methods* **2021**, *51*, 7801–7818. [CrossRef]
9. Cline, D.B.H.; Brockwell, P.J. Linear prediction of ARMA processes with infinite variance. *Stoch. Process. Their Appl.* **1985**, *19*, 281–296. [CrossRef]
10. Timmer, D.H.; Pignatiello, J.J., Jr. Change point estimates for the parameters of an AR (1) process. *Qual. Reliab. Eng. Int.* **2003**, *19*, 355–369. [CrossRef]
11. Bai, J.; Perron, P. Estimating and testing linear models with multiple structural changes. *Econometrica* **1998**, *66*, 47–78. [CrossRef]
12. Eckley, I.A.; Fearnhead, P.; Killick, R. Analysis of changepoint models. In *Bayesian Time-Series Models*; Cambridge University Press: Cambridge, UK, 2011; pp. 203–224. [CrossRef]
13. Killick, R.; Fearnhead, P.; Eckley, I.A. Optimal detection of changepoints with a linear computational cost. *J. Am. Stat. Assoc.* **2012**, *107*, 1590–1598. [CrossRef]
14. Truong, C.; Oudre, L.; Vayatis, N. ruptures: Change point detection in Python. *arXiv* **2018**, arXiv:1801.00826. [CrossRef]
15. Lavielle, M.; Teyssière, G. Detection of multiple change-points in multivariate time series. *Lith. Math. J.* **2006**, *46*, 287–306. [CrossRef]
16. Adams, R.P.; MacKay, D.J.C. Bayesian online changepoint detection. *arXiv* **2007**, arXiv:0710.3742. [CrossRef]
17. Fearnhead, P. Exact and efficient Bayesian inference for multiple changepoint problems. *Stat. Comput.* **2006**, *16*, 203–213. [CrossRef]

18. Zhu, H.; Jiao, H.; Gao, W.; Meng, X.B. Bayesian Change-Point Analysis Approach to Detecting Aberrant Test-Taking Behavior Using Response Times. *J. Educ. Behav. Stat.* **2023**, *48*, 490–520. [CrossRef]
19. Hall, P.; Simar, L. Estimating a changepoint, boundary, or frontier in the presence of observation error. *J. Am. Stat. Assoc.* **2002**, *97*, 523–534. [CrossRef]
20. Fearnhead, P.; Rigaill, G. Changepoint detection in the presence of outliers. *J. Am. Stat. Assoc.* **2019**, *114*, 169–183. [CrossRef]
21. Lu, S.; Wang, T. Bayesian Multiple Change Point Detection in the Presence of Outliers and Its Application to the Magnitude-Frequency Distributions. *Environmetrics* **2025**, *36*, e2877. [CrossRef]
22. Vo, D.H.; Tran, M.P.-B. Volatility spillovers between energy and agriculture markets during the ongoing food & energy crisis: Does uncertainty from the Russo-Ukrainian conflict matter? *Technol. Forecast. Soc. Change* **2024**, *208*, 123723. [CrossRef]
23. Roeger, W.; Welfens, P.J.J. Gas price caps and electricity production effects in the context of the Russo-Ukrainian War: Modeling and new policy reforms. *Int. Econ. Econ. Policy* **2022**, *19*, 645–673. [CrossRef]
24. Adekoya, O.O.; Adefemi, A.; Tula, O.A.; Umoh, A.A.; Gidiagba, J.O. A comprehensive review of Liquefied Natural Gas (LNG) market dynamics: Analyzing the current trends, challenges, and opportunities in the global LNG market. *World J. Adv. Res. Rev.* **2024**, *21*, 58–74. [CrossRef]
25. Hu, H.; Wei, W.; Chang, C.-P. The relationship between shale gas production and natural gas prices: An environmental investigation using structural breaks. *Sci. Total Environ.* **2020**, *713*, 136545. [CrossRef]
26. Kovács, S.; Bühlmann, P.; Li, H.; Munk, A. Seeded binary segmentation: A general methodology for fast and optimal changepoint detection. *Biometrika* **2023**, *110*, 249–256. [CrossRef]
27. Verzelen, N.; Fromont, M.; Lerasle, M.; Reynaud-Bouret, P. Optimal change-point detection and localization. *Ann. Stat.* **2023**, *51*, 1586–1610. [CrossRef]
28. Rogge, E. The European energy crisis, the Dutch TTF, and the market correction mechanism: A financial markets perspective. *J. World Energy Law Bus.* **2024**, *17*, 184–200. [CrossRef]
29. Priem, R. A market correction mechanism regulation as a consequence of the 2021–2023 energy crisis. *J. World Energy Law Bus.* **2024**, *17*, 407–421. [CrossRef]
30. Marhold, A.-A. Towards a 'security-centred' energy transition: Balancing the European Union's ambitions and geopolitical realities. *J. Int. Econ. Law* **2023**, *26*, 756–769. [CrossRef]
31. Zhang, Y. Tensions Between International Economic Law and the EU Energy Security Regulations During the Securitization of EU-Russian Gas Relations: Way Forward? *J. World Trade* **2023**, *57*, 433–456. [CrossRef]
32. Ciciretti, V.; Bucci, A. Building optimal regime-switching portfolios. *N. Am. J. Econ. Financ.* **2023**, *64*, 101837. [CrossRef]
33. Pindyck, R.S. Volatility in natural gas and oil markets. *J. Energy Dev.* **2004**, *30*, 1–19.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.