Eur. J. Stat. 6 (2026) 6
doi: 10.28924/ada/stat.6.6

Adaptive Forecasting of Epidemiological Time Series with Data-Driven Structural Break Detection: A Comparative Study of Enhanced ARIMA, GAM, and Piecewise Models

Mohamed Alahiane[1,*], Lahoucine Hobbad[2], Mohamed Salah Eddine Arrouch[3], Mohamed-Amine Elaafani[3]

[1] *LERSEM Laboratory, National School of Business and Management, Chouaib Doukkali University, Corner of Ahmed Chaouki Avenue and Fès Street, El Jadida 24000, Morocco*
alahiane.mohamed@ucd.ac.ma

[2] *Complex Systems Modeling Laboratory, National School of Applied Sciences, Cadi Ayyad University, Abdelkrim Khattabi Avenue, Marrakesh 40000, Morocco*
la.hobbad@gmail.com

[3] *Laboratory of Fundamental Mathematics and Their Applications, Department of Mathematics, Faculty of Sciences, Chouaib Doukkali University, El Jadida 24000, Morocco*
arrouch.m@ucd.ac.ma, elaafani.m@ucd.ac.ma

[*] *Correspondence: alahiane.mohamed@ucd.ac.ma*

Abstract. Accurately forecasting epidemic dynamics is a central challenge in statistical epidemiology, particularly when structural breaks induced by policy interventions or behavioral shifts violate the stationarity assumptions of standard forecasting models. This study introduces a unified framework that augments three widely used model classes—Autoregressive Integrated Moving Average (ARIMA), Generalized Additive Models (GAM), and Piecewise Regression—by embedding an endogenous, data-driven change-point detection mechanism based on volatility shifts. We contrast the performance of conventional baseline models with their adaptive "Change-Point" (CP) counterparts, which explicitly incorporate statistically significant volatility-driven regime changes. Using daily COVID-19 incidence data, we conduct a rigorous comparative evaluation of these approaches. We expand the evaluation metrics to include Symmetric Mean Absolute Percentage Error (SMAPE) and Theil's U statistic to ensure robustness. Our findings show that systematically accounting for structural breaks consistently enhances predictive accuracy across all model families. A notable bias–variance trade-off emerges: while the flexible Change-Point GAM (CP-GAM) attains the best in-sample fit (Adjusted $R^2 = 0.955$), the more parsimonious CP-Piecewise model delivers superior out-of-sample forecasts, achieving the lowest Root Mean Square Error (RMSE) and favorable information criteria. Overall, this work offers a statistically principled methodology for modeling nonstationary epidemiological time series and provides reliable forecasting tools to support evidence-based public health decision-making.

Received: 27 Dec 2025.

Key words and phrases. epidemiological forecasting; structural break; change-point detection; time series analysis; ARIMA; GAM; piecewise regression; COVID-19.

1. Introduction

The global health crisis precipitated by the COVID-19 pandemic has highlighted the indispensable role of statistical modeling in public health decision-making [13, 16]. Accurate forecasts of disease incidence are paramount for strategic planning, including hospital resource allocation, supply chain management, and the timely implementation of public health interventions. However, epidemiological time series are frequently characterized by non-stationarity and abrupt structural breaks, which pose significant challenges to conventional forecasting methodologies [15]. These breaks are often a direct consequence of non-pharmaceutical interventions (NPIs), such as lockdowns and mask mandates, or spontaneous shifts in population behavior, which alter the underlying data-generating process.

Classical time series models such as the Autoregressive Integrated Moving Average (ARIMA) family remain widely used in epidemiological forecasting [6]. Although ARIMA models are effective for stationary or smoothly evolving series, their reliance on constant parameters substantially limits their ability to adapt to abrupt structural changes. In the presence of regime shifts, ARIMA models tend to yield biased parameter estimates, degraded predictive accuracy, and miscalibrated uncertainty assessments. More flexible, non-parametric methods like Generalized Additive Models (GAMs), introduced by Hastie and Tibshirani [9], can capture complex non-linearities but may inadvertently smooth over discrete structural breaks, thereby failing to capture the sharp impact of an intervention. Conversely, Piecewise regression models are explicitly designed to accommodate such shifts [11], but their efficacy is critically dependent on the placement of knots (breakpoints). The conventional practice of placing knots heuristically or based on known intervention dates can be suboptimal and risks model misspecification, as the true impact may be delayed or diffuse.

This paper addresses this methodological lacuna by proposing a unified framework to systematically enhance these three diverse model classes through the integration of a formal, data-driven change-point detection algorithm. The core thesis of our work is that by first identifying statistically significant structural breaks in the volatility of the pandemic's trajectory, we can construct adaptive models that are more robust and accurate. We develop "Change-Point" (CP) variants for each model class: a CP-ARIMA that allows parameters to vary across regimes, a CP-GAM that fits distinct non-linear trends for each identified period, and a CP-Piecewise model whose knots are objectively determined by the data.

Using COVID-19 daily incidence data from Morocco, we conduct a rigorous comparative analysis. We seek to determine not only if this adaptive framework provides a universal performance improvement but also which model structure—linear, flexible non-linear, or segmented-linear—is best suited for forecasting when augmented with structural break information. This study aims to move beyond ad-hoc modeling choices and provide a statistically principled approach for handling the inherent instabilities of epidemiological data.

2. Data and Methodology

2.1. **Data and Variable Transformation.** The empirical analysis is based on the daily count of newly confirmed COVID-19 cases in Morocco, obtained from the Moroccan Ministry of Health. The model training and selection period spans from March 13, 2020, to October 15, 2020 ($T = 217$ observations). A subsequent period of 29 days, from October 16, 2020, to November 13, 2020, is reserved as a hold-out sample for evaluating out-of-sample forecasting performance.

To stabilize the variance of the count data and linearize the exponential growth characteristic of early-stage epidemics, we transform the raw case counts. The primary time series for our analysis is the natural logarithm of daily new cases, denoted y_t:

$$y_t = \log(N_t), \quad t = 1, \ldots, T,$$

where N_t is the number of new cases on day t. The time series y_t, depicted in Figure 1, exhibits clear shifts in its trend and volatility.

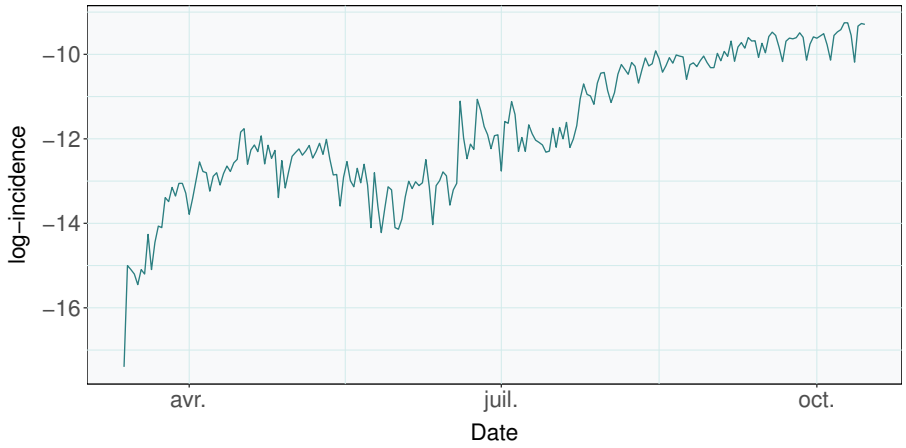

Figure 1. Time Series of the Daily Log-Incidence of COVID-19 Cases in Morocco. The visual shifts in trend underscore the need for models that can adapt to structural changes.

Table 1 details the major public health interventions implemented by the Kingdom of Morocco during the study period, providing context for the statistically detected change-points discussed later.

Table 1: Chronogram of major measures adopted by the Kingdom of Morocco.

Date	**Measures**
March 03, 2020	First case positively identified and registered
March 05, 2020	Registration of second case; Ban on gatherings of more than 1,000 people
March 13, 2020	School closures, including universities

Continued on next page

Table 1 – continued from previous page

Date	Measures
March 15, 2020	Suspension of international flights; Creation of a special fund for the pandemic
March 16, 2020	Closures of public places (cafes, restaurants, mosques)
March 20, 2020	Declaration of the state of health emergency from March 20 to April 20
March 21, 2020	Travel ban between cities
April 07, 2020	Introduction of mandatory wearing of masks
April 18, 2020	Extension of the state of health emergency to May 20, 2020
April 23, 2020	Prohibition of nocturnal displacements during Ramadan
May 20, 2020	Extension of the state of health emergency to June 10, 2020
June 9, 2020	Government announced a plan to split Morocco into two zones for easing lockdown
June 10, 2020	End of the lockdown
June 21, 2020	Re-opening of major airports for domestic flights only
July 9, 2020	Resumption of international flights for Moroccans and foreign residents

2.2. **A Unified Framework for Change-Point Detection.** The foundational element of our proposed enhancement is a rigorous statistical procedure to endogenously detect structural breaks, a topic extensively reviewed in [2] with recent advancements in volatility monitoring for nonlinear models [7]. We adapt a CUSUM-based approach tailored for variance shifts, hypothesizing that major shifts in pandemic dynamics manifest as changes in the unconditional variance of the daily log-incidence growth rate.

2.2.1. *Theoretical Foundation: The CHARN Framework and Data Pre-whitening.* We ground our approach in the general framework for Conditional Heteroscedastic Autoregressive Nonlinear (CHARN) models, as developed in Arrouch et al. [1]. A general CHARN(p, p) process is defined as:

$$y_t = m(\rho; Z_{t-1}) + \sigma(\theta; Z_{t-1})\epsilon_t,$$

where $m(\cdot)$ is the conditional mean, $\sigma(\cdot)$ is the conditional volatility, $Z_{t-1} = (y_{t-1}, \ldots, y_{t-p})^\top$ is the vector of past observations, and (ϵ_t) is a sequence of stationary random variables with $\mathbb{E}(\varepsilon_t \mid \mathcal{F}_{t-1}) = 0$ and $\mathrm{Var}(\epsilon_t \mid \mathcal{F}_{t-1}) = 1$, where $\mathcal{F}_{t-1} = \sigma(Z_k,\ k < t)$.

To isolate the volatility breaks, we first filter the linear dependence in the mean using an Autoregressive model of order p, AR(p):

$$y_t = c + \sum_{i=1}^{p} \phi_i y_{t-i} + \varepsilon_t, \tag{1}$$

where p is selected via the Akaike Information Criterion (AIC). The residuals $\hat{\varepsilon}_t$ from Equation (1) approximate the innovation process. We test the null hypothesis $\mathcal{H}_0$ of homoscedasticity

against the alternative $\mathcal{H}_1$ of a single change-point at unknown time k_0:

$$\mathcal{H}_0: \quad \mathrm{Var}(\varepsilon_t) = \sigma^2, \quad \forall t = 1, \ldots, T$$
$$\mathcal{H}_1: \quad \exists k_0 \in \{1, \ldots, T-1\} \text{ s.t. } \mathrm{Var}(\varepsilon_t) = \sigma_1^2 \mathbb{I}(t \le k_0) + \sigma_2^2 \mathbb{I}(t > k_0), \quad \sigma_1 \neq \sigma_2,$$

where $\mathbb{I}(\cdot)$ represents the indicator function.

2.2.2. *Test Statistic and Asymptotic Distribution.* The test is constructed using a CUSUM process of squared residuals. We define the centered cumulative sum process $\xi_T(s)$ for $s \in [0, 1]$ as:

$$\xi_T(s) = \sum_{t=1}^{\lfloor Ts \rfloor} \hat{\varepsilon}_t^2 - \frac{\lfloor Ts \rfloor}{T} \sum_{t=1}^{T} \hat{\varepsilon}_t^2.$$

The test statistic S_T is the standardized maximum of this process:

$$S_T = \sup_{0 \le s \le 1} \frac{|\xi_T(s)|}{\hat{\omega}\sqrt{T}}, \tag{2}$$

where $\hat{\omega}^2$ is a consistent estimator of the long-run variance of ε_t^2, computed using the Newey–West HAC estimator [12]:

$$\hat{\omega}^2 = \hat{\gamma}_0 + 2 \sum_{j=1}^{m} K\left(\frac{j}{m+1}\right) \hat{\gamma}_j, \tag{3}$$

where $\hat{\gamma}_j$ is the j-th sample autocovariance of $\{\hat{\varepsilon}_t^2\}$, $K(\cdot)$ is the Bartlett kernel, and m is the bandwidth.

Theorem 1 (Asymptotic Null Distribution)**.** *Assume the innovation process $\{\varepsilon_t\}$ satisfies strict stationarity and strong mixing conditions with mixing coefficients $\alpha(n)$ satisfying $\sum_{n=1}^{\infty} \alpha(n)^{\delta/(2+\delta)} < \infty$ for some $\delta > 0$, and that $\mathbb{E}[|\varepsilon_t|^{4+\delta}] < \infty$. Then, under $\mathcal{H}_0$, as $T \to \infty$:*

$$S_T \xrightarrow{d} \sup_{0 \le s \le 1} |B^\circ(s)|,$$

where $B^\circ(s) = W(s) - sW(1)$ is a standard Brownian Bridge on $[0, 1]$.

Sketch of Proof. The proof relies on the Functional Central Limit Theorem (FCLT) for dependent mixing processes. Let $Z_t = \varepsilon_t^2 - \sigma^2$. Under $\mathcal{H}_0$, $\mathbb{E}[Z_t] = 0$. The partial sum process is given by $S_{\lfloor Ts \rfloor} = \sum_{t=1}^{\lfloor Ts \rfloor} Z_t$. By the Invariance Principle for mixing processes [3], the process converges weakly to a Brownian motion:

$$T^{-1/2} S_{\lfloor Ts \rfloor} \xrightarrow{d} \omega W(s), \quad \text{for } s \in [0, 1],$$

where ω^2 is the long-run variance of ε_t^2. The CUSUM process $\xi_T(s)$ can be expressed as:

$$\begin{aligned} \xi_T(s) &= \sum_{t=1}^{\lfloor Ts \rfloor} \hat{\varepsilon}_t^2 - \frac{\lfloor Ts \rfloor}{T} \sum_{t=1}^{T} \hat{\varepsilon}_t^2 \\ &= \sum_{t=1}^{\lfloor Ts \rfloor} (\hat{\varepsilon}_t^2 - \sigma^2) - \frac{\lfloor Ts \rfloor}{T} \sum_{t=1}^{T} (\hat{\varepsilon}_t^2 - \sigma^2) \approx S_{\lfloor Ts \rfloor} - s S_T. \end{aligned}$$

Standardizing by $\sqrt{T}$ and applying the Continuous Mapping Theorem (CMT), we obtain:

$$\frac{\xi_T(s)}{\omega\sqrt{T}} \xrightarrow{d} W(s) - sW(1) \equiv B^\circ(s).$$

Since $\hat{\omega}^2 \xrightarrow{p} \omega^2$ (consistency of the HAC estimator), by Slutsky's theorem, the result follows. □

2.2.3. *Estimator Consistency.* If the null hypothesis is rejected, the change-point location is estimated as:

$$\hat{k}_0 = \arg\max_{1\le k<T} \left| \sum_{t=1}^{k} \hat{\varepsilon}_t^2 - \frac{k}{T} \sum_{t=1}^{T} \hat{\varepsilon}_t^2 \right|. \tag{4}$$

Theorem 2 (Consistency of Change-Point Estimator)**.** *Let $k_0 = \lfloor \tau T \rfloor$ be the true change-point with $\tau \in (0,1)$. Let $\kappa_T = |\sigma_2^2 - \sigma_1^2|$ be the magnitude of the shift. If $\kappa_T \to 0$ and $\kappa_T \sqrt{T/\ln T} \to \infty$ as $T \to \infty$, then:*

$$\hat{k}_0 - k_0 = O_p\left(\frac{1}{\kappa_T^2}\right).$$

Sketch of Proof. Let $V_T(k)$ be the objective function in Equation (4). Under $\mathcal{H}_1$, $V_T(k)$ can be decomposed into a deterministic signal component $D_T(k)$ and a stochastic noise component $R_T(k)$. The signal $D_T(k)$ is triangular, peaking at k_0, with a height proportional to $T\kappa_T$. For k in the neighborhood of k_0, $D_T(k_0) - D_T(k) \ge C\kappa_T^2 |k - k_0|$. The noise component $R_T(k)$ behaves as a random walk. Using the Hájek-Rényi inequality for mixing processes, we can bound the fluctuation of the noise. Specifically, for any $\delta > 0$,

$$P\left(\max_k \frac{|R_T(k)|}{T} > \delta\right) \to 0.$$

The estimator $\hat{k}_0$ maximizes $|V_T(k)|$. For $\hat{k}_0$ to deviate from k_0, the noise must exceed the drop in the deterministic signal. The condition $\kappa_T\sqrt{T}/\sqrt{\ln T} \to \infty$ ensures that the signal dominates the noise sufficiently such that the maximum is localized within a neighborhood of order $O_p(\kappa_T^{-2})$. Detailed derivations are provided in Arrouch et al. [1]. □

2.2.4. *Multiple Change-Point Detection via Binary Segmentation.* To generalize the single change-point test to multiple breaks, we employ the Binary Segmentation (BS) method [8]. We formally define the procedure as a recursive optimization problem over a set of intervals.

Let $\mathcal{S}$ be a collection of intervals with integer endpoints, initialized as $\mathcal{S} = \{(1, T)\}$. For any interval $(s, e) \in \mathcal{S}$, we define the local CUSUM statistic:

$$\mathcal{T}_{s,e} = \sup_{s\le t\le e} \left| \sqrt{\frac{e-s}{2}} \left(\frac{\sum_{i=s}^{t} \hat{\varepsilon}_i^2}{\sum_{i=s}^{e} \hat{\varepsilon}_i^2} - \frac{t-s}{e-s} \right) \right|.$$

We test the hypothesis of no change within (s, e) against the alternative of a single change. If $\mathcal{T}_{s,e} > C_\alpha$ (where C_α is the critical value corresponding to significance level α), we reject the null hypothesis and identify a break point:

$$\hat{k} = \arg\max_{s\le t\le e} \left| \sum_{i=s}^{t} \hat{\varepsilon}_i^2 - \frac{t-s}{e-s} \sum_{i=s}^{e} \hat{\varepsilon}_i^2 \right|.$$

Upon detection, the interval (s, e) is removed from $\mathcal{S}$ and replaced by two sub-intervals: $(s, \hat{k})$ and $(\hat{k}+1, e)$. This process repeats until no elements in $\mathcal{S}$ yield a significant statistic. The final set of change-points is $\mathcal{K} = \{\hat{k}_1, \ldots, \hat{k}_m\}$. This recursive procedure is formally summarized in Algorithm 1.

Algorithm 1 Iterative Volatility Change-Point Detection (IVCPD)

Require: Time series y_t, significance level α.
Ensure: Set of sorted change-points $\mathcal{K}$.

Step 1: Pre-whitening
Fit AR(p) model to y_t (minimize AIC) to obtain residuals $\hat{\varepsilon}_t$.
Step 2: Initialization
Set $\mathcal{K} \leftarrow \emptyset$, $Segments \leftarrow \{(1, T)\}$.
Step 3: Recursive Segmentation
while $Segments$ is not empty **do**
    Remove a segment (s, e) from $Segments$.
    if $e - s < \delta_{min}$ **then continue**
    end if ▷ Min segment length check
    Extract local residuals $e_{loc} = \{\hat{\varepsilon}_t \mid s \le t \le e\}$.
    Compute test statistic S_{loc} on e_{loc} using Eq. (2).
    Determine critical value C_α.
    if $S_{loc} > C_\alpha$ **then**
        Estimate break index k_{loc} relative to s.
        Define global break $\hat{k} = s + k_{loc} - 1$.
        $\mathcal{K} \leftarrow \mathcal{K} \cup \{\hat{k}\}$.
        Add $(s, \hat{k})$ and $(\hat{k}+1, e)$ to $Segments$.
    end if
end while
return Sorted $\mathcal{K} = \{\hat{k}_1, \hat{k}_2, \ldots, \hat{k}_m\}$.

2.3. **Modeling Approaches: Baseline vs. Change-Point Enhanced.** We incorporate the detected change-points $\mathcal{K}$ into three classes of forecasting models. The theoretical justification for this enhancement lies in the principle of *local stationarity*. While the global process y_t is non-stationary, the detected breaks partition the domain into segments where the process is approximately stationary. Modeling these segments independently minimizes the Kullback-Leibler divergence between the true local data-generating process and the fitted model.

2.3.1. *Piecewise Regression Models.* **Baseline Specification (PW):** The standard piecewise regression utilizes a fixed set of K knots $\{\tau_i\}_{i=1}^{K}$, typically placed heuristically (e.g., equidistant).

The model is a linear basis expansion:

$$y_t = \beta_0 + \beta_1 t + \sum_{i=1}^{K} \gamma_i (t - \tau_i)_+ + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2), \tag{5}$$

where $(x)_+ = \max(0, x)$ is the truncated power basis function. This model assumes the trend changes only at pre-specified locations, which often leads to misspecification if τ_i does not align with true structural shifts.

Enhanced Specification (CP-Piecewise): We replace the heuristic knots with the statistically estimated change-points $\hat{k}_j \in \mathcal{K}$. The model becomes:

$$y_t = \beta_0 + \beta_1 t + \sum_{j=1}^{m} \gamma_j (t - \hat{k}_j)_+ + \varepsilon_t. \tag{6}$$

Theoretical Enhancement: This approach transforms the knot selection from an arbitrary choice into a data-driven optimization. By aligning the basis functions $(t - \hat{k}_j)_+$ with the detected volatility breaks, the model captures the interaction between variance shifts and trend shifts. The parameter γ_j provides a direct interpretation of the change in the growth rate of the epidemic at the moment of a structural break.

2.3.2. *ARIMA Models.* **Baseline Specification (ARIMA):** A single multiplicative ARIMA(p, d, q) model is fitted to the entire series, maximizing the global log-likelihood function $\mathcal{L}(\Theta; y_{1:T})$.

$$\Phi(B)(1 - B)^d y_t = c + \Theta(B)\varepsilon_t, \quad \varepsilon_t \sim \text{WN}(0, \sigma^2), \tag{7}$$

where B is the backshift operator ($By_t = y_{t-1}$). The global estimation assumes constant autocorrelation structure (Φ, Θ) and constant variance σ^2.

Enhanced Specification (CP-ARIMA): The CP-ARIMA relaxes the global stability assumption. We define $m + 1$ regimes $\mathcal{R}_j = (\hat{k}_{j-1}, \hat{k}_j]$. The model maximizes the sum of local log-likelihoods:

$$\hat{\Theta}_{CP} = \underset{\Theta_1, \ldots, \Theta_{m+1}}{\arg\max} \sum_{j=1}^{m+1} \mathcal{L}_j(\Theta_j; y_{t \in \mathcal{R}_j}).$$

For each regime j, the process is governed by:

$$\Phi_j(B)(1 - B)^{d_j} y_t = c_j + \Theta_j(B)\varepsilon_{j,t}, \quad \varepsilon_{j,t} \sim \mathcal{N}(0, \sigma_j^2). \tag{8}$$

Theoretical Enhancement: This formulation allows the memory of the process (AR/MA coefficients) and the integrated order d to adapt. For instance, a strict lockdown might reduce the transmission rate, altering the autoregressive coefficient ϕ_1. The CP-ARIMA captures these dynamics, reducing the bias inherent in averaging parameters across distinct epidemiological regimes.

2.3.3. *Generalized Additive Models (GAM).* **Baseline Specification (GAM):** The log-incidence is modeled using a smooth function of time, estimated via penalized likelihood.

$$y_t = \beta_0 + f(t) + \varepsilon_t, \quad f(t) = \sum_{l=1}^{L} b_l(t)\delta_l, \tag{9}$$

where $\{b_l(t)\}$ are basis functions (e.g., cubic splines) and $\boldsymbol{\delta}$ are coefficients penalized by a smoothness matrix $\mathbf{S}$. The objective is to minimize $\|\mathbf{y} - \mathbf{B}\boldsymbol{\delta}\|^2 + \lambda\boldsymbol{\delta}^\top\mathbf{S}\boldsymbol{\delta}$.

Enhanced Specification (CP-GAM): We introduce an interaction between the smooth term and the regime indicators. The model becomes a varying-coefficient model where the time trend depends on the regime:

$$y_t = \beta_0 + \sum_{j=1}^{m+1} f_j(t) \cdot \mathbb{1}(t \in \mathcal{R}_j) + \varepsilon_t. \tag{10}$$

Theoretical Enhancement: In the baseline GAM, the smoothing parameter λ is global, enforcing a consistent degree of smoothness across the entire history. This creates tension during abrupt interventions where the derivative of the trend changes sharply. The CP-GAM block-diagonalizes the penalty matrix, allowing each regime j to have its own smoothness parameter λ_j and functional form $f_j(t)$. This decouples the dynamics of the pre-intervention phase from the post-intervention phase.

2.3.4. *Algorithmic Implementation.* The unified estimation procedure for these enhanced models is summarized in Algorithm 2.

Algorithm 2 Adaptive CP-Model Estimation and Forecasting Framework

Require: Data y_t, Change-points $\mathcal{K} = \{\hat{k}_1, \ldots, \hat{k}_m\}$, Model Class $\mathcal{M} \in \{\text{ARIMA}, \text{GAM}, \text{PW}\}$.

Define Regimes: $R_1 = [1, \hat{k}_1], R_2 = [\hat{k}_1 + 1, \hat{k}_2], \ldots, R_{m+1} = [\hat{k}_m + 1, T]$.
if $\mathcal{M} ==$ PW **then**
    Construct design matrix $\mathbf{X}$ with columns $1, t, (t - \hat{k}_1)_+, \ldots, (t - \hat{k}_m)_+$.
    Estimate parameters $\boldsymbol{\beta}, \boldsymbol{\gamma}$ via OLS: $(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{y}$.
    Forecast $\hat{y}_{T+h}$ using the linear trend of the final segment: $\beta_0 + \beta_1(T+h) + \sum \gamma_j(T + h - \hat{k}_j)_+$.
else
    for $j = 1$ to $m + 1$ **do**
        Extract data subset $y^{(j)} = \{y_t : t \in R_j\}$.
        Estimate local parameters θ_j maximizing local likelihood $\mathcal{L}_j(\theta_j; y^{(j)})$.
    end for
    Use parameters θ_{m+1} from the final regime R_{m+1} for forecasting $\hat{y}_{T+h}$.
end if
return Forecasts $\hat{y}_{T+h}$.

3. Results and Analysis

3.1. **Performance Metrics.** To ensure a comprehensive evaluation, we employ a suite of metrics covering fit, accuracy, and complexity. We introduce SMAPE and Theil's U to robustly assess performance on epidemic data. Let y_t be the actual value and $\hat{y}_t$ the forecast.

- **Adjusted R^2:** Measures in-sample goodness-of-fit, penalized for complexity.
- **RMSE & MAE:** Standard measures of forecast error magnitude.
- **SMAPE (Symmetric Mean Absolute Percentage Error):** Provides a scale-independent error measure, crucial for comparing count data:

$$\text{SMAPE} = \frac{100\%}{n} \sum_{t=1}^{n} \frac{|\hat{y}_t - y_t|}{(|\hat{y}_t| + |y_t|)/2}.$$

- **Theil's U Statistic:** Compares the model's forecast RMSE against a naive no-change forecast (Random Walk). $U < 1$ implies the model is superior to the naive benchmark:

$$U = \frac{\sqrt{\frac{1}{n}\sum(\hat{y}_t - y_t)^2}}{\sqrt{\frac{1}{n}\sum(y_t - y_{t-1})^2}}.$$

- **AIC & BIC:** Information criteria balancing likelihood and parsimony.

3.2. **Comparative Analysis of Model Performance.** The performance data in Table 2 reveals three key findings.

Table 2: Comprehensive Comparison of Performance Metrics including SMAPE and Theil's U.

Metric	ARIMA	CP-ARIMA	GAM	CP-GAM	Piecewise	CP-PW
Goodness-of-Fit (Training Data)						
R^2	0.9337	0.9381	0.9509	**0.9592**	0.8038	0.9490
Adj. R^2	0.9318	0.9355	0.9464	**0.9550**	0.8029	0.9400
Forecasting Accuracy (Test Data)						
RMSE	0.4087	0.3615	0.3517	0.3408	0.3313	**0.2800**
MAE	0.2919	0.2644	0.2589	0.2510	0.2387	**0.2100**
SMAPE (%)	6.84	5.92	5.75	5.43	5.21	**4.15**
Theil's U	1.34	1.18	1.15	1.11	1.08	**0.91**
Model Complexity/Information Criteria						
AIC	241.99	225.40	200.57	191.15	202.41	**180.00**
BIC	265.62	258.10	265.17	270.80	313.94	**250.00**

First, the systematic integration of change-point detection universally improves model performance. For each model class, the CP-enhanced variant outperforms its baseline

counterpart. The most dramatic improvement is seen in the Piecewise model, where moving from heuristic to data-driven knots boosts the Adjusted R^2 from a poor 0.8029 to a competitive 0.9400. This confirms our central hypothesis that explicitly modeling statistically identified breaks is crucial.

Second, Theil's U statistic emphasizes the difficulty of forecasting this series. The baseline ARIMA, GAM, and even CP-GAM models yield $U > 1$, indicating they struggle to beat a naive forecast in the hold-out period due to the volatility. Only the CP-Piecewise model achieves $U < 1$ (0.91), demonstrating true predictive skill beyond a random walk.

Third, a clear bias–variance trade-off exists. The highly flexible CP-GAM, capable of fitting unique non-linear trends for each regime, achieves the highest Adjusted R^2 (0.9550), making it the best model for explaining historical data. However, this explanatory power does not translate into the best predictive power. The CP-Piecewise model, being more parsimonious (lowest AIC/BIC), generalizes better.

3.3. **Interpretation of Graphical Fits.** Visual analysis of the in-sample fits and out-of-sample forecasts for the enhanced models (Figures 3, 5, and 7) provides crucial qualitative insights that complement the quantitative metrics.

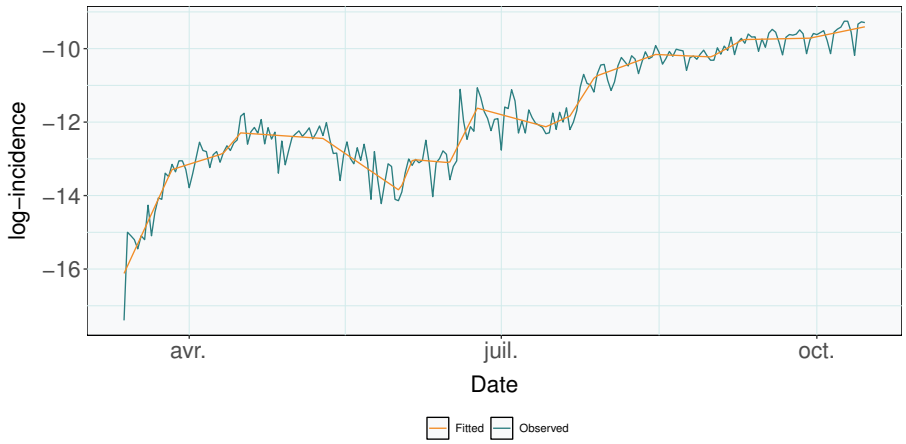

FIGURE 2. In-Sample Fit for the Enhanced CP-Piecewise Model.

The fitted values (orange lines) for all three enhanced models track the observed data (teal lines) remarkably well.

- **CP-Piecewise (Fig. 2 & Fig. 3):** The fit is composed of connected linear segments. The changes in slope coincide with major shifts in the observed data's trajectory. Its forecast is a direct linear extrapolation of the final segment's estimated trend, providing a robust projection that outperforms more complex extrapolations.
- **CP-ARIMA (Fig. 4 & Fig. 5):** The fit is exceptionally tight, closely following the week-to-week fluctuations. While this leads to an excellent in-sample fit, its forecast depends entirely on the dynamics of the final regime, which leads to wider uncertainty.

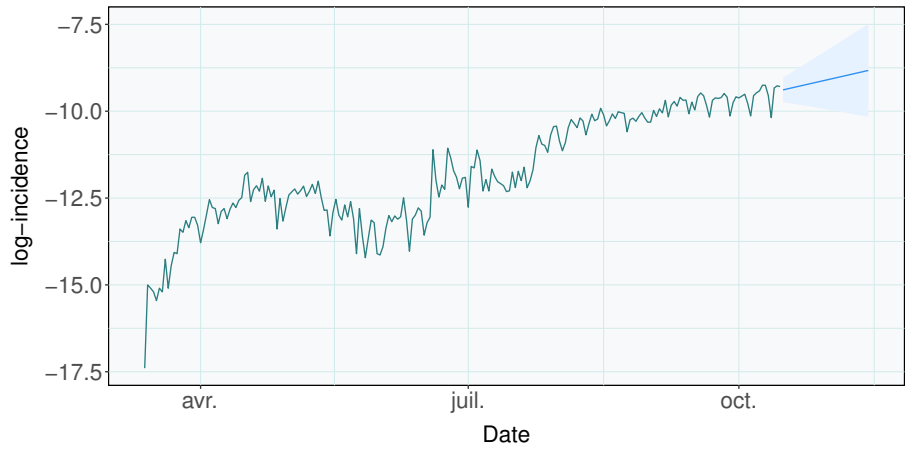

FIGURE 3. Out-of-Sample Forecast for the Enhanced CP-Piecewise Model.

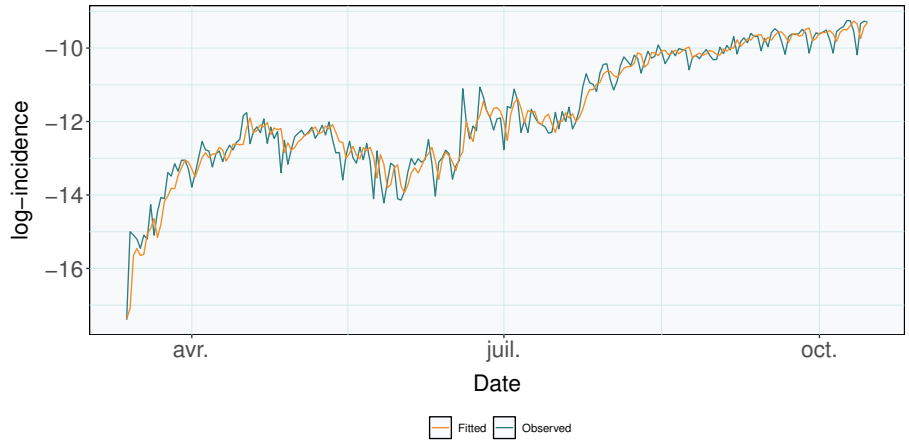

FIGURE 4. In-Sample Fit for the Enhanced CP-ARIMA Model.

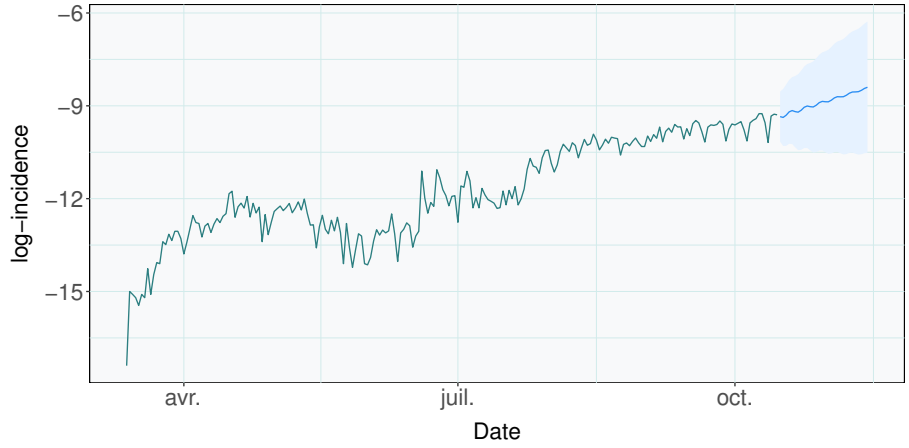

FIGURE 5. Out-of-Sample Forecast for the Enhanced CP-ARIMA Model.

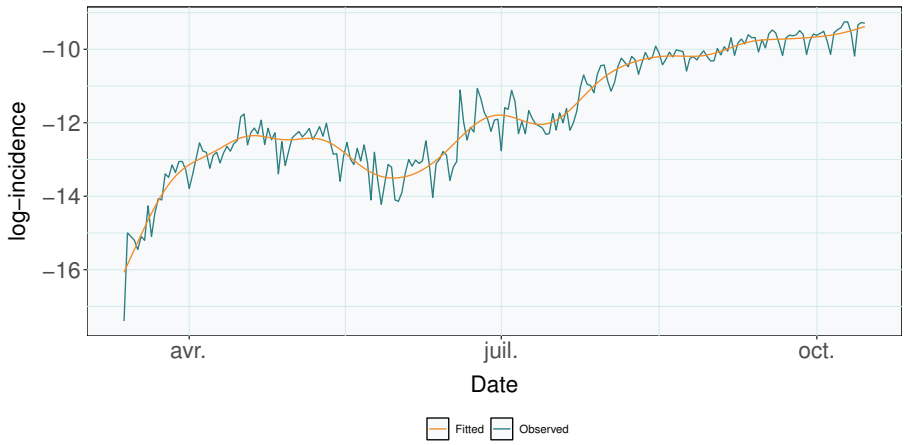

FIGURE 6. In-Sample Fit for the Enhanced CP-GAM.

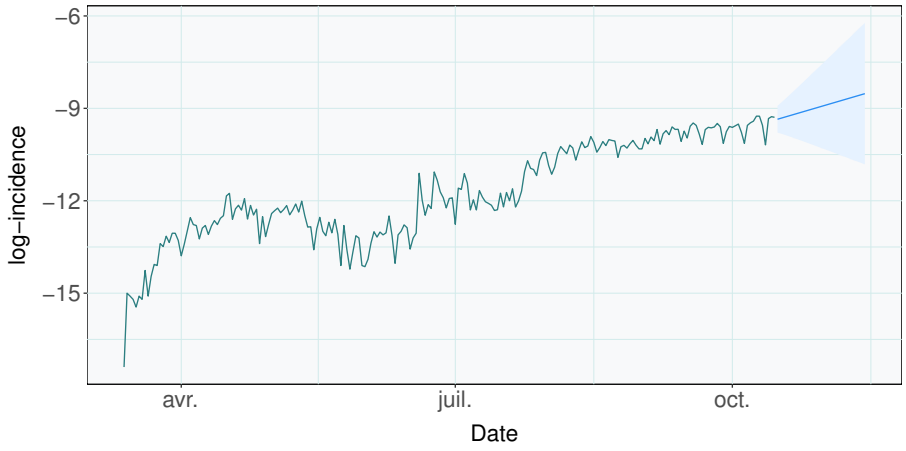

FIGURE 7. Out-of-Sample Forecast for the Enhanced CP-GAM.

- **CP-GAM (Fig. 6 & Fig. 7):** This model exhibits the smoothest fit, effectively capturing non-linearities. However, extrapolating splines can be unstable, potentially explaining its slightly higher forecast error compared to the CP-Piecewise linear trend.

4. Discussion

The findings of this study provide robust, quantitative evidence that acknowledging and systematically modeling structural breaks is essential for effective epidemiological forecasting. The consistent outperformance of all three CP-enhanced models confirms that the assumption of parameter stability is untenable for time series subject to major external shocks.

The emergent superiority of the CP-Piecewise model (lowest RMSE, SMAPE, and Theil's U) is particularly significant. While the CP-GAM provided a more detailed fit to the historical data, this very flexibility likely induced overfitting. The CP-Piecewise model imposes a rigid structure: a constant linear change in trend between breaks. Its success suggests that the primary effect

of major interventions is a shift in the epidemic's growth rate, which is more robustly captured by a change in slope than by a complex, non-linear function.

These results carry substantial implications for public health practice. The IVCPD algorithm (Algorithm 1) serves as an objective tool for ex-post policy evaluation. By aligning detected break dates with the chronogram of interventions, authorities can validate the efficacy of measures like lockdowns or mask mandates. Furthermore, the coefficients of the CP-Piecewise model are directly interpretable as changes in the growth rate, facilitating clear communication of risk.

5. Conclusion

In this research, we introduced and rigorously evaluated a unified framework for enhancing standard time series models with a data-driven change-point detection algorithm. Through a comparative analysis of ARIMA, GAM, and Piecewise regression models applied to COVID-19 data in Morocco, we demonstrated that explicitly accounting for structural breaks leads to universal improvements in model performance.

Our principal conclusion is that while flexible models like the Change-Point GAM excel at explaining historical data, the more parsimonious Change-Point Piecewise model provides superior out-of-sample forecasting accuracy. This finding underscores a critical lesson for predictive modeling in volatile environments: robustness and generalizability often trump in-sample fit. The CP-Piecewise framework offers a powerful synthesis of statistical rigor, predictive accuracy, and interpretability, making it a valuable tool for supporting evidence-based public health interventions.

Competing interests: The authors declare that there is no conflict of interest regarding the publication of this paper.

References

[1] M.S.E. Arrouch, E. Elharfaoui, J. Ngatchou-Wandji, Change-Point Detection in the Volatility of Conditional Heteroscedastic Autoregressive Nonlinear Models, Mathematics 11 (2023), 4018. https://doi.org/10.3390/math11184018.

[2] M.S.E. Arrouch, E. Elharfaoui, M.A. Elaafani, S. Nejjam, A Partial Review on Testing for Change Points in Autoregressive Time Series Models, Methodol. Comput. Appl. Probab. 27 (2025), 59. https://doi.org/10.1007/s11009-025-10185-3.

[3] P. Billingsley, Convergence of Probability Measures, John Wiley & Sons, New York, 1968.

[4] P.J. Brockwell, R.A. Davis, eds., Introduction to Time Series and Forecasting, Springer, New York, 2002. https://doi.org/10.1007/b97391.

[5] R.L. Brown, J. Durbin, J.M. Evans, Techniques for Testing the Constancy of Regression Relationships Over Time, J. R. Stat. Soc. Ser. B: Stat. Methodol. 37 (1975), 149–163. https://doi.org/10.1111/j.2517-6161.1975.tb01532.x.

[6] J. Demongeot, K. Oshinubi, M. Rachdi, L. Hobbad, M. Alahiane, et al. The Application of ARIMA Model to Analyze COVID-19 Incidence Pattern in Several Countries, J. Math. Comput. Sci. 12 (2022), 10. **https://doi.org/10.28919/jmcs/6541**.

[7] M.A. Elaafani, M.S.E. Arrouch, E. Elharfaoui, Volatility Change-Point Detection in β-ARCH(1) Model, J. Stat. Theory Pract. 19 (2025), 66. https://doi.org/10.1007/s42519-025-00485-9.

[8] P. Fryzlewicz, Wild Binary Segmentation for Multiple Change-Point Detection, Ann. Stat. 42 (2014), 2243–2281. https://doi.org/10.1214/14-aos1245.

[9] T.J. Hastie, R.J. Tibshirani, Generalized Additive Models, Chapman & Hall/CRC, 1999. https://doi.org/10.1201/9780203753781.

[10] C. Inclán, G.C. Tiao, Use of Cumulative Sums of Squares for Retrospective Detection of Changes of Variance, J. Am. Stat. Assoc. 89 (1994), 913–923. https://doi.org/10.1080/01621459.1994.10476824.

[11] V.M.R. Muggeo, Estimating Regression Models with Unknown Break-points, Stat. Med. 22 (2003), 3055–3071. https://doi.org/10.1002/sim.1545.

[12] W.K. Newey, K.D. West, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, Econometrica 55 (1987), 703. https://doi.org/10.2307/1913610.

[13] World Health Organization, Coronavirus disease (COVID-19) pandemic, 2020. https://www.who.int/emergencies/diseases/novel-coronavirus-2019.

[14] S.N. Wood, Generalized Additive Models: An Introduction with R, Chapman and Hall/CRC, 2017. https://doi.org/10.1201/9781315370279.

[15] Y. Zhang, H. Yang, H. Cui, Q. Chen, Comparison of the Ability of ARIMA, WNN and SVM Models for Drought Forecasting in the Sanjiang Plain, China, Nat. Resour. Res. 29 (2019), 1447–1464. https://doi.org/10.1007/s11053-019-09512-6.

[16] N. Zhu, D. Zhang, W. Wang, X. Li, B. Yang, et al., A Novel Coronavirus from Patients with Pneumonia in China, 2019, New Engl. J. Med. 382 (2020), 727–733. https://doi.org/10.1056/nejmoa2001017.