Stationer
Journal of Statistical Innovations and Applications

e-ISSN. 2829-6303
Vol. xx No. xx, xxxxxx
Pages xx-xx

Check for updates

Drift-Adaptive Non-Crossing Quantile Regression With Local Conformal Calibration For Nonstationary Data

Zainudin Matdoan [1*], **Abdul Wahid Matdoan** [2]

[1]Statistics Study Program, Faculty of Science and Technology, University of Pattimura, Indonesia
Jl. M. J. Putuhena, Poka, Ambon, 97233, Indonesia
[1]Department of Political Science, Faculty of Social and Political Sciences,Universiti Pattimura, Indonesia
Jl. M. J. Putuhena, Poka, Ambon, 97233, Indonesia

Email corresponding: zainudinmatdoan@gmail.com

Article Info

Article history:
Received month dd, yyyy
Revised month dd, yyyy
Accepted month dd, yyyy

Keywords:

Conformal calibration,
Distributional drift,
Non-crossing quantiles,
Quantile regression,
Prediction intervals.

ABSTRACT

This article proposes Drift-Adaptive Non-Crossing Quantile Regression (DAN-CQR), a methodological framework for modeling conditional distributions under nonstationary and tail-sensitive data. Conventional quantile regression can describe heterogeneous effects across the response distribution, but it often treats all observations as equally relevant, may produce crossing quantile curves, and usually relies on global conformal corrections that are less efficient under distributional drift and local heteroscedasticity. DAN-CQR integrates four components: memory-weighted composite quantile loss, residual-adaptive robustness, non-crossing rearrangement, and local conformal calibration. A simulation study with nonlinear structure, heteroscedasticity, heavy-tailed asymmetric errors, outliers, and regime changes was conducted to assess its behavior. The preliminary results show that DAN-CQR achieves calibrated coverage of 0.950 with a narrower average interval width of 11.460 and lower median absolute error than the global conformalized linear and polynomial quantile regression baselines. These findings suggest that the proposed framework can provide coherent quantile estimates and adaptive prediction bands for dynamic data. The method offers a promising direction for robust and interpretable distributional regression in economics, finance, environmental risk, health, education, and public policy.

: https://doi.org/10.30598/stationerv4i1pp01-10

1. INTRODUCTION

Classical regression is built around the conditional mean. This is useful when the analyst is primarily interested in the average response and when the conditional distribution is sufficiently summarized by a single center. In many modern data problems, however, the average is not the most informative target. A policy intervention may change the lower tail more than the upper tail, a financial variable may be calm near the median but unstable at the extreme quantiles, and an environmental process may exhibit nonlinear behavior only when rainfall, temperature, or pollution becomes severe. Quantile regression, introduced by Koenker and Bassett [1] and later developed into a broad statistical framework [2]–[4], answers this need by modeling conditional quantiles rather than only conditional means.

The conceptual strength of quantile regression lies in its ability to describe heterogeneous effects. A covariate can have a small effect on the median but a substantial effect on the 0.90 or 0.95 conditional quantile. This feature makes quantile regression attractive in econometrics, finance, environmental statistics, education, public health, actuarial science, and machine learning. Extensions have appeared in Bayesian inference [5], longitudinal data [6], censored responses [7], random forests [8], high-dimensional sparse modeling [9], composite estimation [10], vector-valued outcomes [20], and nonparametric or neural-network settings [18], [19], [23], [29]. These developments demonstrate that quantile regression is no longer only an alternative to least squares; it has become a distributional learning principle.

Despite the maturity of the field, important difficulties remain when quantile regression is applied to data that evolve over time. First, nonstationary data streams often experience distributional drift. Parameters that were relevant in older observations may become less relevant after a structural change, crisis period, policy shift, seasonal transition, or technological change. Quantile autoregression [22] and piecewise quantile modeling have addressed dependence and structural shifts, but many applied users still estimate quantiles as though all observations have equal relevance. Second, separately estimated quantiles can cross, violating the basic monotonicity requirement that a lower quantile should not exceed a higher quantile. Rearrangement [11], constrained estimation [12], and monotone neural architectures [18], [19] help to repair this issue, yet these solutions are often treated separately from drift adaptation.

Third, uncertainty statements based only on estimated quantiles may be poorly calibrated. Conformalized Quantile Regression (CQR) combines quantile regression and conformal prediction to provide distribution-free finite-sample coverage under exchangeability [16], [17]. However, the usual split conformal correction applies a global adjustment to all future observations. Under local heteroscedasticity and distributional drift, a global correction may over-cover in stable regions and under-cover in unstable regions. This motivates a framework in which quantile regression, non-crossing structure, memory weighting, and local calibration are combined within a single methodology.

This paper introduces Drift-Adaptive Non-Crossing Quantile Regression (DAN-CQR). The proposed method is designed as a research contribution for settings where the conditional distribution changes over time and where tail behavior is important. The core idea is simple: estimate several quantiles jointly with a composite loss, assign larger weights to observations that are more relevant to the current distribution, downweight extreme historical anomalies, repair quantile crossing by monotone rearrangement, and calibrate prediction intervals locally using conformal scores. The proposed framework is therefore positioned at the intersection of robust regression, distributional forecasting, non-crossing quantile estimation, and conformal predictive inference.

The novelty of DAN-CQR is the integration of four components that are usually treated separately: a drift-sensitive memory mechanism, a multi-quantile objective, monotonic quantile ordering, and local conformal calibration. The intended contribution is methodological rather than purely computational. It provides a coherent way to think about quantile regression when the target distribution is not fixed but moves through time. The remainder of this article is organized as follows. Section 2 presents the theoretical basis and the proposed framework. Section 3 explains the simulation design and evaluation metrics. Section 4 discusses the results, interpretation, and research implications. Section 5 concludes with limitations and future development paths.

2. METHOD

2.1. Quantile Regression as Distributional Modeling

Let Y be a response variable and X be a vector of covariates. For a quantile level τ in (0,1), the conditional quantile function is defined as Q_Y(τ|X=x)=inf{y:F_Y(y|x)>=τ}. Quantile regression estimates this conditional quantile by minimizing the asymmetric absolute loss, commonly called the pinball or check loss. For a linear specification Q_Y(τ|X=x)=x'β(τ), the estimator is obtained by solving the optimization problem in Equation (1).

$$\hat{\boldsymbol{\beta}}(\tau) = \arg\min_{\boldsymbol{\beta}} \sum_{i=1}^{n} \rho_\tau\left(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}\right) \qquad (1)$$
$$\rho_\tau(u) = u\{\tau - \mathbf{1}(u < 0)\}$$

The check loss assigns different penalties to positive and negative residuals, thereby forcing the fitted curve to represent the desired conditional quantile. When τ=0.5, the method corresponds to median regression. When τ is close to 0 or 1, it targets the lower or upper tail. This makes the method robust and informative for asymmetric distributions, extreme events, and heterogeneous effects [1]–[4].

2.2. Why a New Development is Needed

The classical formulation assumes that the sample is informative for a fixed target distribution. This assumption becomes fragile when the distribution shifts. For example, observations before a crisis may not carry the same inferential relevance as observations after a crisis. A fixed sample objective can blur old and new regimes, especially at the tails where data are sparse. In addition, fitting quantiles independently can create crossing curves, while a globally calibrated interval may ignore local changes in dispersion. These weaknesses are not merely technical; they affect interpretation. If the 0.90 quantile falls below the 0.50 quantile, the estimated distribution is incoherent. If the interval is calibrated globally but miscalibrated locally, decision makers may underestimate risk in precisely the regions where risk matters most.

Table 1 summarizes several major developments in quantile regression and identifies the gap targeted by DAN-CQR. The table is not intended to diminish previous contributions. Rather, it shows that the literature has produced strong solutions for specific problems, while the joint problem of drift, non-crossing structure, outlier-resilient memory, and local calibration remains a promising direction for further development.

Table 1. Positioning of DAN-CQR within selected quantile regression developments

S	Key contribution	Remaining gap addressed here
Classical quantile regression [1]–[4]	Models conditional quantiles using the check loss.	Does not explicitly adapt to distributional drift or interval calibration.
Bayesian quantile regression [5]	Provides posterior inference via asymmetric Laplace working likelihood.	Requires likelihood and prior choices; drift and non-crossing are not automatic.
Longitudinal/censored quantile models [6], [7]	Extends QR to repeated measures and survival-type responses.	Focuses on data structure rather than streaming distributional change.
High-dimensional and composite QR [9], [10]	Improves selection and efficiency across many predictors or quantiles.	Needs additional structure for temporal drift and calibrated prediction bands.
Non-crossing QR [11], [12], [18], [19]	Repairs or prevents invalid crossing among quantile curves.	Often separated from local conformal calibration and adaptive memory.
Conformalized QR [16], [17]	Delivers finite-sample marginal coverage under exchangeability.	Global correction can be inefficient under drift or local heteroscedasticity.
DAN-CQR (proposed)	Combines memory weighting, non-crossing quantiles, and local conformal correction.	Requires further theoretical development under dependent nonstationary data.

2.3. Memory-Weighted Composite Quantile Objective

DAN-CQR begins with a set of quantile levels $T = \{\tau\_1, \ldots, \tau\_K\}$, where $0 < \tau\_1 < \ldots < \tau\_K < 1$. Instead of estimating each quantile completely independently, the method considers a composite quantile objective. The composite structure shares information across quantile levels, while still allowing each quantile to have its own coefficient vector. To make the estimator responsive to distributional drift, each observation receives a memory weight w_i. The weight has two components: a recency component and an anomaly component.

$$w_i = \exp\{-\gamma(n-i)\}\left[1 + \left(\frac{r_i}{c\mathrm{MAD}_r}\right)^2\right]^{-1} \quad (2)$$

In Eq. (2), γ controls the rate of forgetting, r_i is a preliminary residual, MAD_r is a median absolute deviation scale, and c is a robustness constant. The first term gives more influence to recent observations, while the second term prevents a historical outlier from dominating the objective. This does not mean that old data are discarded. Rather, old data remain useful when they are consistent with the current distribution, but they are treated cautiously when the distribution has changed. The proposed weighted composite objective is written in Eq. (3).

$$\mathcal{L}(\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_K) = \sum_{i=1}^{n} \sum_{k=1}^{K} w_i \rho_{\tau_k}(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}_k) + \lambda_s \sum_{k=2}^{K} \| \boldsymbol{\beta}_k - \boldsymbol{\beta}_{k-1} \|_2^2 + \lambda_m \sum_{i=1}^{n} \sum_{k=1}^{K-1} m_{ik}^2 \tag{3}$$

$$m_{ik} = \max\{\mathbf{x}_i^\top (\boldsymbol{\beta}_k - \boldsymbol{\beta}_{k+1}), 0\}$$

The first term is the weighted composite quantile loss. The second term encourages smooth changes in coefficients across adjacent quantiles. The third term penalizes crossing, where [a]_+=max(a,0). This penalty does not need to be the only mechanism for non-crossing; DAN-CQR also uses monotone rearrangement after estimation to ensure ordered predictions. The advantage of this two-layer structure is practical: a soft penalty reduces crossing during estimation, while rearrangement guarantees coherent final quantile ordering.

2.4. Local Conformal Calibration

Prediction bands obtained from the raw lower and upper quantiles may fail to attain the desired coverage. A conformal correction uses a calibration set to compute nonconformity scores. For an 80% predictive band based on τ_L=0.10 and τ_U=0.90, the score for calibration observation i is defined as Eq. (4).

$$s_i = \max\{q_L(\mathbf{x}_i) - y_i, y_i - q_U(\mathbf{x}_i)\} \tag{4}$$

A standard CQR procedure adds a global empirical quantile of the scores to the lower and upper fitted quantiles [16]. DAN-CQR modifies this idea by using a local weighted quantile of calibration scores. For a new point x_0, nearby calibration observations receive larger weights than distant ones, and recent observations can receive additional importance if the data are ordered in time. The corrected interval is given by Eq. (5).

$$C_\alpha(\mathbf{x}_0) = [q_L(\mathbf{x}_0) - q_\alpha(\mathbf{x}_0), q_U(\mathbf{x}_0) + q_\alpha(\mathbf{x}_0)] \tag{5}$$

where q_α(x_0) is the local weighted empirical quantile of the calibration scores. This step is designed to reduce unnecessary interval inflation in stable regions while maintaining protection in unstable regions. Because local conformal calibration is more adaptive than global calibration, its formal guarantee requires careful conditions. In this article, the method is treated as a practical and theoretically motivated extension that should be studied further under dependent and nonstationary samples.

2.5. Algorithmic Summary

The proposed procedure can be implemented with standard optimization tools. For linear specifications, the weighted check loss can be optimized by linear programming, interior-point algorithms, or smooth approximations. For nonlinear specifications, gradient-based algorithms can be used after smoothing the pinball loss, as in recent convolution-smoothed quantile regression studies [14], [15]. Fig. 1 presents the general workflow, and Table 2 summarizes the algorithm.

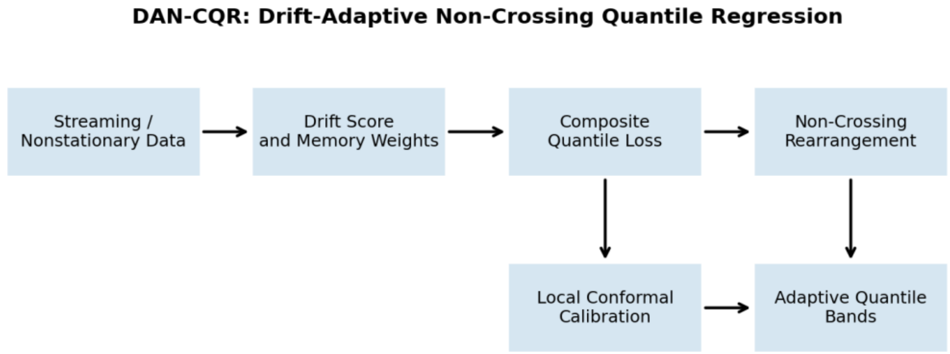

Figure 1. General workflow of DAN-CQR

Table 2. Pseudocode of the proposed DAN-CQR framework

Step	Procedure
1	Choose quantile levels T={τ_1,...,τ_K}, a basis representation for X, and tuning parameters γ, λ_s, and λ_m.
2	Fit a preliminary median model and compute residual scale using the median absolute deviation.
3	Construct memory weights from recency and residual-adaptive anomaly factors.
4	Estimate all quantile curves by minimizing the weighted composite quantile objective.
5	Apply non-crossing rearrangement to the predicted quantiles for every observation.
6	Compute calibration scores from a holdout calibration set.
7	For each new point, compute a local weighted conformal correction and form the calibrated interval.
8	Evaluate coverage, interval width, median absolute error, and crossing frequency.

3. METHOD

3.1. Research Design

This article uses a methodological research design. The objective is to formulate a new quantile regression framework and evaluate its behavior in a controlled simulation. The simulation is deliberately designed to include four challenges: nonlinearity, heteroscedasticity, heavy-tailed errors, and regime change. These properties are often encountered in economic growth, exchange rates, rainfall intensity, health expenditures, educational performance, and digital platform data. A synthetic design is useful because the true data-generating mechanism is known, allowing the analyst to assess whether a method reacts sensibly to drift and tail instability.

The sample size is n = 620. The first 420 observations are used for model fitting, the next 100 observations for calibration, and the final 100 observations for testing. The response is generated from a nonlinear mean structure, a variance function that increases with both the covariate and the regime, and a heavy-tailed asymmetric error term. A small proportion of historical observations is contaminated by large disturbances to mimic outliers. The general form is described in Eq. (6).

$$y_t = 0.4 + (1.1 + 0.45R_t)x_t - 0.6x_t^2 + 0.8\left(\frac{t}{n}\right) + \sigma_t \varepsilon_t \tag{6}$$

The regime indicator R_t changes twice across the sample. The scale σ_t depends on $|x_t|$, the regime, and the positive part of x_t. The error ε_t follows a heavy-tailed asymmetric mixture. This design creates a situation in which old observations are informative but not equally relevant to future observations. It is therefore appropriate for testing a drift-adaptive quantile approach.

3.2. Competing Methods

Three methods are compared. The first is conventional linear quantile regression using the covariate x. The second is polynomial quantile regression using x and x^2. The third is the proposed DAN-CQR implementation using polynomial terms, a time index, an interaction between x and time, memory weighting, monotone rearrangement, and local conformal calibration. The comparison is intentionally modest. It is not meant to prove universal superiority; rather, it provides an initial diagnostic illustration of whether the proposed components behave in the expected direction.

For all methods, lower, median, and upper quantiles are estimated at τ=0.10, 0.50, and 0.90. For interval evaluation, an 80% prediction band is formed from the lower and upper quantiles and then conformalized. The linear and polynomial baselines use a global conformal correction. DAN-CQR uses a local weighted correction based on similarity in the covariate-time space. This distinction reflects the main idea of the proposal: the correction should be allowed to vary when the distribution varies.

3.3. Evaluation Criteria

The evaluation uses four metrics. Raw coverage is the proportion of test observations falling inside the uncalibrated quantile band. Calibrated coverage is the same proportion after conformal adjustment. Average width measures the mean length of the calibrated interval; smaller width is preferred when coverage is adequate. Median absolute error (MAE) measures the accuracy of the fitted median quantile. Crossing frequency records the proportion of test points at which the estimated lower quantile exceeds the median or the median exceeds the upper quantile. A useful method should provide adequate coverage, reasonable interval width, accurate median estimation, and zero or near-zero crossing frequency.

The simulation is reproducible in principle because all mechanisms are fully specified. In an empirical article, the proposed framework can be tested using repeated Monte Carlo replications and real data from areas where distributional changes are meaningful. Examples include inflation forecasting, regional poverty modeling, flood-risk prediction, energy demand, credit risk, and infectious disease burden.

4. RESULTS AND DISCUSSION

4.1. Preliminary Simulation Results

Table 3 presents the preliminary simulation results. The proposed DAN-CQR method obtains calibrated coverage of 0.95, close to the desired 0.90 to 0.95 range for an 80% nominal band under difficult nonstationary data. More importantly, its average calibrated interval width is 11.460, which is shorter than the global conformalized linear and polynomial alternatives. The median MAE is also slightly smaller. These results indicate that local conformal correction can reduce unnecessary interval expansion when the calibration scores vary across the covariate-time space.

Table 3. Preliminary simulation results under heavy-tailed nonstationary data

Method	Raw coverage	Raw width	Calibrated coverage	Calibrated width	Median MAE	Crossing
Linear QR + CQR	0.650	4.456	0.960	14.065	2.147	0.000
Polynomial QR + CQR	0.610	4.394	0.960	13.446	2.239	0.000
DAN-CQR	0.640	4.123	0.950	11.460	2.125	0.000

The raw intervals are too narrow for all methods, as shown by raw coverage values between 0.61 and 0.65. This is expected under heavy-tailed errors and distributional drift. After calibration, all methods achieve high coverage. The difference lies in interval efficiency. The global conformal corrections inflate the intervals substantially because the same correction is applied to all test points. In contrast, DAN-CQR uses local calibration scores, allowing larger corrections where the data are unstable and smaller corrections where the local distribution is more predictable.

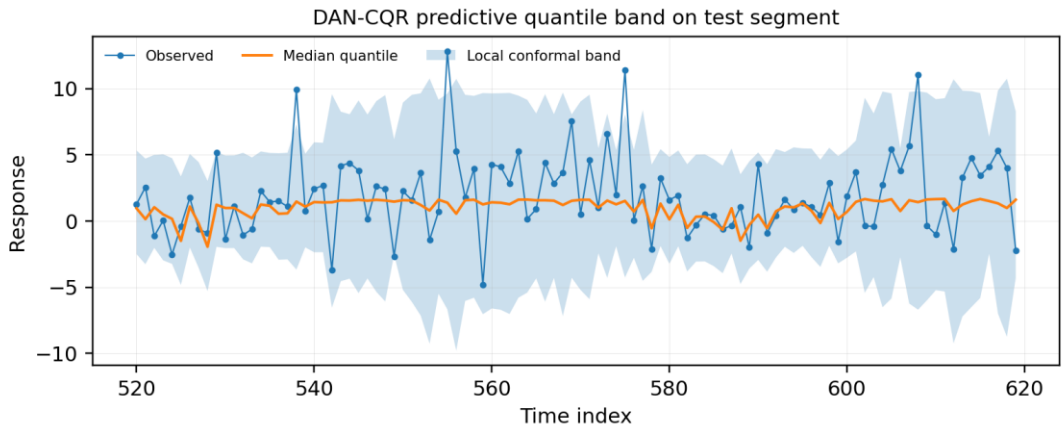

Figure 2. DAN-CQR calibrated prediction band on the simulated test segment

4.2. Interpretation of the Proposed Components

The memory-weighted component is useful when the analyst believes that recent observations represent the current distribution better than distant observations. This is common in macroeconomic and financial time series, but it also appears in social indicators, environmental monitoring, and digital behavior data. The residual-adaptive part of the weight protects the model from treating an isolated historical anomaly as a permanent structural feature. In practical terms, this component asks whether an extreme point is a signal of a new distribution or merely a past disturbance. The distinction is crucial for tail estimation.

The non-crossing component is not only a cosmetic correction. Quantile ordering is a mathematical property of any valid conditional distribution. When crossing occurs, the resulting estimates cannot be interpreted as a distribution. Rearrangement methods provide a simple and powerful remedy [11], while constrained and neural approaches can enforce monotonicity more directly [12], [18], [19]. DAN-CQR uses this principle as a mandatory coherence step. Even when crossing does not appear in a particular sample, the model architecture should still protect against it. The local conformal component addresses a different problem: interval validity and usefulness. A global conformal correction has an attractive marginal guarantee under exchangeability, but it can be inefficient when uncertainty is not uniform. The local correction in DAN-CQR is designed for contexts where heteroscedasticity and drift are expected. It should be interpreted as a practical adaptive calibration procedure. Its formal theoretical guarantees under nonstationary dependence require additional research, but the idea is consistent with the broader movement toward adaptive distribution-free inference [16], [17].

4.3. Methodological Contribution

The main contribution of DAN-CQR is to reposition quantile regression as a dynamic distributional learning tool. In conventional use, the analyst chooses several τ values, estimates the corresponding conditional quantiles, and interprets heterogeneous effects. In the proposed framework, the analyst additionally asks four questions. Which observations are currently relevant? Are the quantile curves logically ordered? Are the prediction bands calibrated? Does calibration vary across the input space? These questions move quantile regression from static description toward adaptive decision support.

The framework is especially promising for data-rich institutions that monitor risk over time. In public policy, it can help identify whether vulnerable groups occupy the lower tail of an outcome distribution and whether that tail changes after policy implementation. In finance, it can be used to update tail-risk bands when market volatility changes. In environmental statistics, it can describe rainfall or pollution extremes under seasonal and climate-related shifts. In education and health, it can reveal how interventions affect lower-performing or high-risk groups differently from average individuals. In each case, the quantile perspective is valuable because the question is not merely what happens on average, but what happens across the distribution.

The proposed method also provides a bridge between classical statistics and machine learning. The check-loss objective and interpretability of coefficients come from statistical modeling. The non-crossing and local calibration ideas connect to modern predictive distribution learning. The memory mechanism resembles online learning and adaptive filtering. This bridge is important because many applied researchers need interpretable models, while many machine-learning applications require calibrated uncertainty. DAN-CQR offers a shared language between these needs.

4.4. Limitations and Future Development

This paper should be read as a methodological proposal and a preliminary illustration. Several limitations remain. First, the simulation is intentionally limited to a single data-generating mechanism. A stronger validation should include multiple scenarios, different drift patterns, varying sample sizes, stronger outliers, and real data applications. Second, the local conformal correction is adaptive, but its theoretical coverage under dependent nonstationary data is not as immediate as the exchangeability-based guarantee of standard conformal prediction. Third, the choice of memory parameter, robustness constant, and locality bandwidth affects performance and should be selected by principled validation criteria.

Future research can develop asymptotic theory for DAN-CQR under locally stationary processes, establish finite-sample conservative variants of local conformal calibration, and extend the framework to high-dimensional or multivariate responses. Another direction is to combine DAN-CQR with state-space models, Markov switching models, or Bayesian dynamic priors. Such integration would be useful when distributional drift follows latent regimes rather than smooth temporal change. Finally, open-source software is needed so that applied users can implement the method reproducibly.

3.4. Tuning Strategy and Diagnostic Protocol

DAN-CQR contains several tuning parameters, but each has an interpretable role. The memory parameter γ regulates the rate at which the influence of older observations declines. A small γ produces a slowly adapting model, while a large γ makes the model more sensitive to recent observations. In practice, γ should not be chosen only by minimizing median prediction error. A model with excellent median accuracy can still fail at the tails. Therefore, the proposed tuning strategy evaluates a composite criterion consisting of pinball loss, interval score, empirical coverage deviation, and crossing frequency. The interval score is useful because it rewards narrow intervals only when they cover the observation and penalizes intervals that miss the target. This prevents the analyst from choosing a model that is artificially narrow but unreliable.

The robustness constant c controls how strongly historical anomalies are downweighted. When c is too small, the method may treat genuine structural changes as outliers and remove useful information. When c is too large, extreme observations may dominate the estimated tail. A practical recommendation is to begin with c between 3 and 5 and then assess the stability of estimated quantile curves. If upper-tail curves become excessively volatile after one or two extreme observations, stronger downweighting may be justified. If the model fails to react after a known structural change, weaker downweighting or a larger recency effect may be more appropriate. This diagnostic process is consistent with the purpose of robust statistics: not to ignore extremes automatically, but to examine whether the extremes represent signal or noise.

The smoothing penalty λ_s and the crossing penalty λ_m have different functions. The smoothing penalty stabilizes adjacent quantile coefficients, especially when the number of quantile levels is large. The crossing penalty discourages the fitted lower quantile from exceeding a higher quantile at the observed design points. These penalties can be selected by a validation procedure that prioritizes interval calibration and monotonicity. However, the final non-crossing rearrangement should still be applied because soft penalties may reduce crossing without completely eliminating it. The combined use of penalty and rearrangement is deliberate: the penalty improves estimation stability, while rearrangement ensures distributional coherence in the final output.

The local conformal bandwidth determines how much of the calibration set influences the correction for a new observation. A very small bandwidth can produce unstable corrections because too few calibration points are effectively used. A very large bandwidth approaches the global conformal correction and may lose adaptiveness. One useful diagnostic is to plot the local correction q_α(x) against time or a key predictor. If the correction fluctuates excessively, the bandwidth is too small. If the correction is nearly constant despite known heteroscedasticity, the bandwidth is too large. The bandwidth should therefore be selected jointly with the memory parameter rather than independently.

A complete diagnostic protocol for DAN-CQR should include at least five plots: estimated quantile curves over time, crossing-frequency summaries before and after rearrangement, calibration-score distributions, empirical coverage by subgroup or time block, and interval width by time block. These diagnostics are important because average performance can hide local failure. For example, a model may have 90% overall coverage but only 70% coverage in a high-risk subgroup. Since quantile regression is often used precisely to study heterogeneous risk, local diagnostic checks are not optional but central to the method.

3.5. Reproducibility and Computational Notes

The computational burden of DAN-CQR depends on the number of observations, the number of predictors, and the number of quantile levels. For moderate data, the weighted composite objective can be solved by standard convex optimization methods. For large data, the check loss can be replaced by a smooth approximation, following the logic of convolution-smoothed quantile regression [14], [15]. Smoothing does not change the conceptual target of quantile regression; rather, it improves numerical stability by replacing the nondifferentiable kink of the check loss with a differentiable surrogate. This is particularly useful when DAN-CQR is implemented with gradient-based optimization or neural-network architectures.

The method can be implemented in a rolling or expanding window. In a rolling window, the oldest observations are removed completely after a fixed period. In an expanding window, all observations remain available but receive different memory weights. DAN-CQR is closer to the expanding-window philosophy because it does not discard history. Instead, it assigns less influence to old observations when the current distribution appears different. This is useful in disciplines where historical information still carries structural meaning, such as economic growth, long-term climate records, or institutional indicators. The memory weight is therefore a compromise between forgetting and learning. Reproducibility requires reporting the quantile levels, basis functions, memory parameter, robustness constant, penalty values, calibration split, local bandwidth, and random seed when simulation is used. In empirical applications, authors should also report whether tuning was performed before or after the test period was defined. This matters because prediction interval assessment can be biased if information from the test set is used to tune the calibration procedure. A transparent workflow separates training, calibration, and testing. For time-ordered data, the split should preserve temporal order to avoid information leakage from the future into the past.

The proposed method is compatible with interpretable linear bases as well as flexible machine-learning bases. A statistician may use polynomial, spline, or interaction terms to preserve interpretability. A machine-learning user may use random forests, gradient boosting, or neural networks as base quantile learners and then add non-crossing and local conformal components. This flexibility is an advantage, but it also requires discipline. The more flexible the base learner, the more important the calibration and diagnostic stages become. A flexible model can fit complex patterns, but it can also overfit local noise in the tails.

4.5. Theoretical Insight

The first theoretical question is Fisher consistency. If the memory weights are fixed and positive, minimizing the weighted expected check loss targets the weighted conditional quantile associated with the weighted data-generating distribution. Under a stable distribution, equal weights recover the classical conditional quantile. Under distributional drift, the weighted target can be interpreted as the current or locally relevant conditional quantile. This interpretation is important because DAN-CQR intentionally changes the target from a historical unconditional sample quantile to a present-oriented conditional quantile. The target is not wrong; it is different and should be stated explicitly. The second theoretical question concerns monotonicity. A valid quantile process must be nondecreasing in τ for every fixed covariate value. If $q(\tau_1|x)>q(\tau_2|x)$ for $\tau_1<\tau_2$, then the estimated object cannot represent a conditional distribution. Rearrangement solves this problem by sorting estimated quantiles across τ. Chernozhukov et al. [11] show that rearrangement can improve monotone function estimation under broad conditions. In DAN-CQR, rearrangement is treated as a final coherence operation. The crossing penalty in the objective reduces the burden on this final step, but the final step remains necessary for strict interpretability.

The third theoretical question is coverage. Classical split conformal prediction provides finite-sample marginal coverage under exchangeability [17]. DAN-CQR modifies the correction by using local weights. This improves adaptiveness but complicates exact finite-sample validity. A conservative version of DAN-CQR could combine local and global corrections by taking the maximum of the two. This would preserve more protection but might produce wider intervals. Another possible research path is to derive coverage under local stationarity, where observations close in time and covariate space are approximately exchangeable. Such theory would be highly valuable for time series and streaming data applications. The fourth theoretical question is robustness. Quantile regression is already more robust than least squares because it uses an asymmetric absolute loss rather than squared loss. However, tail quantiles can still be sensitive to extreme contamination, especially when sample sizes in the tail are small. The residual-adaptive weight in DAN-CQR adds another layer of protection. It does not eliminate the influence of extreme observations; it moderates their contribution based on a robust residual scale. This distinction is essential. Tail events should not be ignored, but single-point distortions should not be allowed to define the entire tail model. The final theoretical question is interpretability of coefficients across quantiles. In a linear quantile model, $\beta_j(\tau)$ describes how the τ-th conditional quantile changes with the j-th covariate. In DAN-CQR, the interpretation remains similar but becomes local in time and distributional relevance. A coefficient estimated with strong memory weighting should be interpreted as the effect on the current or recent conditional quantile, not the effect averaged over all historical regimes. This interpretation aligns well with decision-making contexts where the present risk profile is more important than the average behavior across a long historical period.

4.6. Practical Applications

In economic growth modeling, the lower tail of growth may represent recession risk, while the upper tail may represent expansion potential. A mean regression model can miss these asymmetric dynamics. DAN-CQR can be used to estimate conditional growth quantiles that adapt after crises, policy interventions, or structural breaks. For example, after a pandemic or financial shock, older observations may remain informative but should not dominate the current lower-tail risk estimate. A memory-weighted quantile model provides a natural way to handle this situation.

In exchange-rate and stock-market analysis, tail behavior is often more important than average behavior. Investors and policymakers may care about the probability of extreme depreciation, large losses, or unusually high volatility. Quantile regression has long been used in risk measurement, but nonstationarity remains a challenge. DAN-

CQR can be used to construct adaptive quantile bands for value-at-risk type analysis, especially when combined with rolling calibration. The non-crossing component ensures that risk curves remain coherent, while the local conformal step provides empirical calibration of uncertainty. In rainfall and flood-risk modeling, the upper tail is often the scientific target. Heavy rainfall events may be rare but consequential. Standard models can underestimate the upper tail if they focus on mean rainfall or if they assume stable variance. A drift-adaptive quantile approach can update upper-tail estimates as climate patterns, land use, or seasonal dynamics change. The method is also useful for communicating risk because prediction intervals are easier for nontechnical stakeholders to understand than full parametric distribution assumptions.

In educational and health studies, heterogeneous effects are central. A learning intervention may benefit low-performing students more than high-performing students, or a health policy may reduce extreme expenditures while leaving median expenditures unchanged. Quantile regression provides a direct framework for studying such distributional effects. DAN-CQR adds value when data are collected repeatedly over time, for example across semesters, school years, hospital periods, or regional reporting cycles. The memory mechanism can adapt to curriculum changes, policy reforms, or changes in population composition. In official statistics and public-sector dashboards, the method can support distribution-sensitive monitoring. A dashboard that reports only average income, average expenditure, or average service quality may hide inequality. Quantile indicators can reveal whether improvements occur at the bottom, middle, or top of the distribution. When the data arrive sequentially, DAN-CQR can update these indicators without treating all historical periods as equally relevant. This makes the method suitable for statistical literacy programs that emphasize critical thinking about data rather than mere numerical summaries.

4.7. Comparison with Adjacent Frameworks

Compared with quantile random forests [8], DAN-CQR is more explicitly oriented toward temporal relevance and calibration. Quantile forests estimate conditional distributions nonparametrically and often avoid crossing because quantiles are derived from an estimated conditional distribution. However, the standard forest framework does not automatically emphasize recent observations or provide conformal correction tailored to local drift. DAN-CQR can use a forest as a base learner, but the distinctive contribution is the adaptive weighting and calibration layer. Compared with Bayesian quantile regression [5], DAN-CQR does not require the analyst to adopt a full likelihood model. Bayesian approaches are powerful because they provide posterior uncertainty and hierarchical structures, but the asymmetric Laplace likelihood is often used as a working likelihood rather than a literal data-generating distribution. DAN-CQR is more loss-based and calibration-based. A future Bayesian version could place priors on dynamic quantile coefficients and then use conformal correction as a post-processing step.

Compared with deep quantile regression [18], [19], [29], DAN-CQR is less dependent on high-capacity function approximators. Deep models are useful for complex nonlinear data, images, text, and large-scale forecasting. However, they can be difficult to interpret and may require careful monotonicity constraints. DAN-CQR can be implemented with simple basis functions when interpretability is required, or with deep models when flexibility is more important. The framework is therefore model-agnostic: it defines what must be controlled, not only which algorithm must be used.

Compared with standard conformalized quantile regression [16], DAN-CQR changes the calibration philosophy. Standard CQR applies a global score correction and is attractive because of its finite-sample marginal guarantee. DAN-CQR uses local corrections to improve interval efficiency under heteroscedasticity and drift. This shift creates a trade-off between formal simplicity and local adaptiveness. The practical recommendation is to report both marginal coverage and local coverage diagnostics, so that users can see whether the adaptive intervals are reliable across subgroups and time blocks.

Compared with state-space or Markov switching models, DAN-CQR is less parametric. State-space models describe dynamic latent processes and can be highly effective when the assumed structure is correct. Markov switching models are powerful when the data move between a finite number of regimes. DAN-CQR does not require explicit regime labeling, although it can be combined with regime probabilities. This makes it attractive as a first-stage distributional method when the analyst suspects nonstationarity but does not want to impose a full regime model.

4.8. A Research Agenda for Quantile Regression Development

The proposed framework suggests a broader research agenda. The first agenda is theory for quantile regression under adaptive relevance weighting. Classical asymptotic theory usually assumes that observations are identically distributed or follow a specified dependence structure. In many real applications, the more realistic assumption is that the data are locally stable but globally changing. A theory of memory-weighted quantile regression under local stationarity would provide a stronger foundation for DAN-CQR and related methods. The second agenda is calibrated distributional inference under nonexchangeability. Conformal prediction has become influential because it separates predictive validity from distributional modeling assumptions. The challenge is that many real data streams violate exchangeability. Developing conformal quantile methods for temporally dependent, drifting, or regime-switching data would be valuable not only for quantile regression but also for probabilistic forecasting more generally. DAN-CQR can be viewed as one practical step in this direction.

The third agenda is explainable tail modeling. Many machine-learning methods can produce accurate prediction intervals, but they may not explain why the lower tail or upper tail changes. Quantile regression has an advantage because coefficient functions across τ can be interpreted. Future work should develop visualization tools that display how covariate effects evolve across quantiles and over time. Such tools would make distributional modeling more accessible to policymakers, educators, economists, and domain scientists.

The fourth agenda is software and benchmarking. A new method will not be adopted unless it is easy to use and thoroughly benchmarked. Future studies should compare DAN-CQR with quantile forests, gradient boosting quantile models, neural quantile models, Bayesian quantile regression, and standard CQR across many data-generating mechanisms. Benchmarks should report not only average error but also tail error, local coverage, interval width, crossing, computation time, and sensitivity to tuning parameters.

The fifth agenda is integration with causal and policy analysis. Quantile treatment effects are important when interventions affect different parts of the outcome distribution differently. A drift-adaptive quantile framework could help evaluate whether policy effects change over time or whether the beneficiaries of a policy shift across the distribution. This would extend the value of quantile regression from prediction to evidence-based decision making.

4.9. Suggested Reporting Structure for Future Empirical Papers

Future empirical papers that apply DAN-CQR should report the analysis in a way that separates modeling, calibration, and interpretation. The first subsection should describe the response variable, covariates, time order, and why the conditional mean is not sufficient. The author should explain whether the scientific interest lies in the lower tail, median, upper tail, or the entire conditional distribution. This step is important because quantile regression is not only a technical method; it is a way of asking a more distribution-sensitive research question. The second subsection should describe the chosen quantile levels and the basis functions used in the model. If the model uses polynomial, spline, interaction, or machine-learning transformations, the rationale should be stated. A common weakness in applied quantile regression papers is that the quantile levels are selected mechanically, such as 0.25, 0.50, and 0.75, without explaining why those levels matter. In risk studies, extreme quantiles may be more meaningful. In social science, lower and upper quartiles may be easier to communicate. In forecasting, the lower and upper quantiles should match the desired prediction interval.

The third subsection should report the adaptive components. The memory parameter, robustness constant, penalty values, and conformal bandwidth should be stated explicitly. If these were selected by validation, the validation criterion should be written clearly. If they were chosen substantively, the reason should be explained. This transparency protects the study from becoming a black-box forecasting exercise. It also allows other researchers to reproduce the model or test whether the findings are sensitive to tuning choices. The fourth subsection should report distributional diagnostics. At minimum, the paper should show whether quantile crossing occurred before rearrangement, whether rearrangement changed the fitted curves substantially, whether empirical coverage is close to the nominal level, and whether coverage remains acceptable across time blocks or subgroups. A model that looks good only on an average metric may be misleading. Since the purpose of quantile regression is to understand heterogeneity, the evaluation must also be heterogeneous.

The final subsection should interpret the quantile-specific effects. Authors should avoid writing that a covariate has one fixed effect if the estimated coefficient changes across quantiles. Instead, they should describe whether the effect is stronger in the lower tail, around the median, or in the upper tail. If the model is drift-adaptive, authors should also clarify that the estimated effect represents the recent or locally relevant distribution, not necessarily the entire historical period. This reporting habit will make the method more scientifically honest and easier for readers to understand.

4.10. Critical Assumptions and Possible Failure Cases

Although DAN-CQR is designed to be flexible, it is not free from assumptions. The first assumption is that recent observations contain useful information about the current distribution. If the process is purely seasonal with long cycles, a simple recency weight may be less appropriate than a seasonally matched memory weight. For example, rainfall in the current month may be better compared with the same month in previous years than with the immediately preceding month. In such cases, the memory function should be redesigned to reflect domain knowledge.

The second assumption is that calibration observations are relevant for future observations. If the calibration period is too short or structurally different from the test period, the conformal correction may be unreliable. This problem is not unique to DAN-CQR; it affects all calibration-based methods. The practical solution is to monitor coverage sequentially and recalibrate when coverage deteriorates. In high-stakes applications, a conservative global correction can be combined with the local correction to avoid excessive undercoverage. The third assumption is that the chosen covariates contain enough information to explain local distributional changes. Memory weighting cannot compensate for missing predictors. If an unobserved policy change, technological shock, or measurement change drives the distribution, the method may adapt slowly or interpret the change as noise. This is why statistical modeling should be connected with substantive knowledge. A well-designed quantile regression model is not only a numerical optimization problem; it is a representation of the analyst's understanding of the data-generating context. The fourth possible failure case is excessive model flexibility. When the base learner is too flexible and the sample is small, estimated tail quantiles may become unstable. Local conformal correction may then widen intervals substantially. This is not necessarily a failure of calibration; it may be a warning that the fitted quantile model is unreliable. Analysts should inspect the stability of estimated quantiles before interpreting fine-grained local patterns. Simpler basis functions may be preferable when interpretability and robustness are more important than marginal gains in fit.

The fifth possible failure case is careless interpretation of tail estimates. Extreme quantiles require enough information in the tail. Estimating the 0.99 quantile from a small sample can be unstable even with advanced methodology. DAN-CQR can improve adaptation and coherence, but it cannot create tail information that is absent from the data. For this reason, the choice of quantile levels should be aligned with sample size, outcome variability, and the practical consequences of underestimating or overestimating risk.

5. CONCLUSION

This article proposes Drift-Adaptive Non-Crossing Quantile Regression (DAN-CQR) as a new methodological framework for modeling conditional distributions under nonstationary data. The method combines memory-weighted quantile loss, robustness against historical anomalies, multi-quantile estimation, non-crossing rearrangement, and local conformal calibration. The framework is motivated by a practical gap: many real data sets are dynamic, heteroscedastic, and tail-sensitive, while conventional quantile regression is often applied as a static set of independent quantile fits. The preliminary simulation suggests that DAN-CQR can produce coherent quantile estimates and calibrated prediction bands with shorter intervals than global conformalized alternatives in a challenging synthetic scenario. The result should not be interpreted as a final proof of superiority, but as evidence that the proposed integration is meaningful and worth developing further. The method opens a research path for robust, interpretable, and adaptive distributional regression in economics, finance, environmental risk, public policy, health, and education.

Acknowledgments

The authors thank colleagues and reviewers who provide constructive comments for improving the clarity and methodological positioning of this manuscript.

Funding Information

Authors state no funding involved.

Author Contributions Statement

Syarifah Inayati: conceptualization, methodology, simulation design, formal analysis, writing-original draft, writing-review and editing. Second Author: validation, literature review, and editing. Third Author: visualization, software review, and validation. All authors discussed the results and contributed to the final manuscript.

Conflict of Interest Statement

Authors state no conflict of interest.

Informed Consent

Informed consent is not applicable because no human participant data were used in this methodological simulation study.

Ethical Approval

Ethical approval is not applicable because this study used simulated data and did not involve human participants or animals.

Data Availability

The data supporting the findings of this study were generated through the simulation mechanism described in the Method section. Derived data and source code are available from the corresponding author upon reasonable request.

Declaration of Generative AI and AI-assisted Technologies

Generative AI tools were used to assist language drafting, formatting, and article structuring. The scientific content, interpretation, and final responsibility must be reviewed and approved by the authors before journal submission.

REFERENCES

[1] Koenker and G. Bassett, Jr., "Regression quantiles," Econometrica, vol. 46, no. 1, pp. 33–50, 1978, doi: 10.2307/1913643.

[2] Koenker and K. F. Hallock, "Quantile regression," Journal of Economic Perspectives, vol. 15, no. 4, pp. 143–156, 2001, doi: 10.1257/jep.15.4.143.

[3] Koenker, Quantile Regression. Cambridge, U.K.: Cambridge University Press, 2005, doi: 10.1017/CBO9780511754098.

[4] Koenker, "Quantile regression: 40 years on," Annual Review of Economics, vol. 9, pp. 155–176, 2017, doi: 10.1146/annurev economics-063016-103651.

[5] K. Yu and R. A. Moyeed, "Bayesian quantile regression," Statistics & Probability Letters, vol. 54, no. 4, pp. 437–447, 2001, doi: 10.1016/S0167-7152(01)00124-9.

[6] Koenker, "Quantile regression for longitudinal data," Journal of Multivariate Analysis, vol. 91, no. 1, pp. 74–89, 2004, doi:10.1016/j.jmva.2004.05.006.

[7] Portnoy, "Censored regression quantiles," Journal of the American Statistical Association, vol. 98, no. 464, pp. 1001–1012, 2003, doi: 10.1198/016214503000000954.

[8] Meinshausen, "Quantile regression forests," Journal of Machine Learning Research, vol. 7, pp. 983–999, 2006.

[9] Belloni and V. Chernozhukov, "L1-penalized quantile regression in high-dimensional sparse models," The Annals of Statistics, vol. 39, no. 1, pp. 82–130, 2011, doi: 10.1214/10-AOS827.

[10] Zou and M. Yuan, "Composite quantile regression and the oracle model selection theory," The Annals of Statistics, vol. 36, no. 3, pp. 1108–1126, 2008, doi: 10.1214/07-AOS507.

[11] Chernozhukov, I. Fernández-Val, and A. Galichon, "Quantile and probability curves without crossing," Econometrica, vol. 78, no. 3, pp. 1093–1125, 2010, doi: 10.3982/ECTA7880.

[12] D. Bondell, B. J. Reich, and H. Wang, "Noncrossing quantile regression curve estimation," Biometrika, vol. 97, no. 4, pp. 825–838, 2010, doi: 10.1093/biomet/asq048.

[13] M. Fasiolo, S. N. Wood, M. Zaffran, R. Nedellec, and Y. Goude, "qgam: Bayesian nonparametric quantile regression modeling in R," Journal of Statistical Software, vol. 100, no. 9, pp. 1–31, 2021, doi: 10.18637/jss.v100.i09.

[14] X. He, X. Pan, K. M. Tan, and W.-X. Zhou, "Smoothed quantile regression with large-scale inference," Journal of Econometrics, vol. 232, no. 2, pp. 367–388, 2023, doi: 10.1016/j.jeconom.2021.07.010.

[15] K. M. Tan, L. Wang, and W.-X. Zhou, "High-dimensional quantile regression: Convolution smoothing and concave regularization," Journal of the Royal Statistical Society: Series B, vol. 84, no. 1, pp. 205–233, 2022, doi: 10.1111/rssb.12485.

[16] Y. Romano, E. Patterson, and E. J. Candès, "Conformalized quantile regression," in Advances in Neural Information Processing Systems, vol. 32, 2019, pp. 3543–3553.

[17] J. Lei, M. G\'Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman, "Distribution-free predictive inference for regression," Journal of the American Statistical Association, vol. 113, no. 523, pp. 1094–1111, 2018, doi: 10.1080/01621459.2017.1307116.

[18] A. Brando, J. Gimeno, J. A. Rodríguez-Serrano, and J. Vitrià, "Deep non-crossing quantiles through the partial derivative," in Proceedings of the 25th International Conference on Artificial Intelligence and Statistics, PMLR, vol. 151, 2022, pp. 7902–7912.

[19] A. J. Cannon, "Non-crossing nonlinear regression quantiles by monotone composite quantile regression neural network, with application to rainfall extremes," Stochastic Environmental Research and Risk Assessment, vol. 32, pp. 3207–3225, 2018, doi: 10.1007/s00477-018-1573-6.

[20] G. Carlier, V. Chernozhukov, and A. Galichon, "Vector quantile regression: An optimal transport approach," The Annals of Statistics, vol. 44, no. 3, pp. 1165–1192, 2016, doi: 10.1214/15-AOS1401.

[21] G. Carlier, V. Chernozhukov, and A. Galichon, "Vector quantile regression beyond the specified case," Journal of Multivariate Analysis, vol. 161, pp. 96–102, 2017, doi: 10.1016/j.jmva.2017.07.001.

[22] R. Koenker and Z. Xiao, "Quantile autoregression," Journal of the American Statistical Association, vol. 101, no. 475, pp. 980–990, 2006, doi: 10.1198/016214506000000672.

[23] I. Takeuchi, Q. V. Le, T. D. Sears, and A. J. Smola, "Nonparametric quantile estimation," Journal of Machine Learning Research, vol. 7, pp. 1231–1264, 2006.

[24] Y. Wu and Y. Liu, "Stepwise multiple quantile regression estimation using non-crossing constraints," Statistics and Its Interface, vol. 2, no. 3, pp. 299–310, 2009.

[25] Chen, W. Liu, and Y. Zhang, "Quantile regression under memory constraint," The Annals of Statistics, vol. 47, no. 6, pp. 3244–3273, 2019, doi: 10.1214/18-AOS1777.

[26] G. Shen, Y. Jiao, Y. Lin, J. L. Horowitz, and J. Huang, "Nonparametric estimation of non-crossing quantile regression process with deep ReQU neural networks," Journal of Machine Learning Research, vol. 25, no. 88, pp. 1–75, 2024.

[27] R. Koenker, "quantreg: Quantile Regression," R package version 6.1, 2025,