

JIRSS (2021)

Vol. 20, No. 01, pp 247-267

DOI:10.52547/jirss.20.1.247

Finite Sample Properties of Quantile Interrupted Time Series Analysis: A Simulation Study

Rahim Moineddin¹, Christopher Meaney¹, and Sumeet Kalia¹

¹ Department of Family and Community Medicine, Faculty of Medicine, University of Toronto, 500 University Avenue, Toronto, Ontario M5G 1V7, Canada.

Received: 30/10/2020, Revision received: 10/02/2021, Published online: 03/04/2021

Abstract. Interrupted Time Series (ITS) analysis represents a powerful quasi-experimental design in which a discontinuity is enforced at a specific intervention point in a time series, and separate regression functions are fitted before and after the intervention point. Segmented linear/quantile regression can be used in ITS designs to isolate intervention effects by estimating the sudden/level change (change in intercept) and/or the gradual change (change in slope). To our knowledge, the finite-sample properties of quantile segmented regression for detecting level and gradual change remains unaddressed. In this study, we compared the performance of segmented quantile regression and segmented linear regression using a Monte Carlo simulation study where the error distributions were: IID Gaussian, heteroscedastic IID Gaussian, correlated AR(1), and T (with 1, 2 and 3 degrees of freedom, respectively). We also compared segmented quantile regression and segmented linear regression when applied to a real dataset, employing an ITS design to estimate intervention effects on daily-mean patient prescription volumes. Both the simulation study and applied example illustrate the usefulness of quantile segmented regression as a complementary statistical methodology for assessing the impacts of interventions in ITS designs.

Corresponding Author: Rahim Moineddin (Rahim.moineddin@utoronto.ca)
Christopher Meaney (Christopher.Meaney@utoronto.ca)
Sumeet Kalia (Sumeet.Kalia@utoronto.ca)

Keywords. Interrupted Time-Series, Segmented Linear Regression, Segmented Quantile Regression, Monte Carlo Simulation Study.

MSC: 62M10, 62G08, 91B84.

1 Introduction

A time series is a sequence of observations taken at successive points in time. The measurement occasions may be equally or unequally spaced in time. Time may be conceptualized as discrete or continuous. Descriptive analysis of a time-series may include the decomposition of time-series into trend (i.e. long term changes in the population mean), seasonal effects (i.e. cyclic/periodic changes with respect to calendar time) and systematic or non-systematic fluctuations (i.e. residual components). A combination of these factors may exist in an interrupted time series design, on an additive or multiplicative scale in the presence some external intervention.

According to the Cambridge dictionary ‘intervention’ is defined as “*the action of becoming intentionally involved in a difficult situation, in order to improve it or prevent it from getting worse*”. Interrupted Time Series (ITS) is a powerful quasi-experimental time series tool for evaluating effects of interventions on an outcome of interest Cook et al. (1979); Wagner et al. (2002); Gillings et al. (1981); Campbell and Stanley (2015). ITS designs are often analyzed using a form of segmented regression in which a discontinuity is enforced at the intervention point and a separate regression curve is fitted before and after the intervention point. ITS analysis examines the effects of an intervention using the pre-intervention observations to establish a baseline and then investigate if the intervention effects of post-intervention observations deviate from the expectations derived from baseline projections Ferron and Rendina-Gobioff (2005). In particular, segmented regression can isolate the intervention effects by estimating both sudden/level change (change in intercept) and gradual change (change in slope); while taking into account the effects of seasonality, and serial correlation Wagner et al. (2002).

The least squares method is commonly used to estimate the parameters of the segmented linear regression by which the regression lines fit the data set as closely as possible by minimizing the sum of squares of the differences between observed and fitted observations. Quantile regression, which was introduced by Koenker and Bassett Jr (1978), extends the modelling of the conditional mean to the modeling of the conditional quantiles. Furthermore, quantile regression may also be useful for

modeling heteroskedastic and heavy tailed data Koenker (2005). In the past decade, quantile regression has become a widely used technique to model a response variable where the assumption of normally distributed errors is untenable Xiong and Tian (2019); Geraci and Bottai (2007); Honda (2013); Yang and Zhang (2008). Austin et al. (2005) provides a review of quantile regression in the context of biostatistics applications. A comprehensive coverage of quantile regression is given in Koenker (2005). Koenker et al. (2017) provides a solid reference on current extensions of classical quantile regression (including quantile regression for time series designs).

Classic linear regression models focus on modelling the central tendency of the conditional distribution of a response variable given covariates. Quantile regression allows one to investigate how covariates effect any conditional quantile of the response variable. To our knowledge, the finite-sample properties of a quantile segmented regression for detecting level and gradual change in ITS designs remains unaddressed. Hence the purpose of this study is to compare the performance of segmented quantile regression with segmented linear regression for modelling ITS designs commonly encountered in epidemiology and health services research. We conduct an extensive Monte Carlo simulation study, comparing segmented linear regression and segmented quantile regression when the errors are standard normal $N(0, 1)$, heteroskedastic ($N(0, \log^2(\text{time} + 1))$), auto-correlated ($AR(1)$) and non-normally distributed (t -distribution with degrees of freedom 1, 2, and 3). We compare models in terms of bias, MSE, type-1 error and power. We also compare segmented linear regression and segmented quantile regression on a real ITS dataset involving mean daily prescriptions across patients, modelled over-time, involving two separate intervention points.

2 Method

In linear regression the conditional expectation of the dependent variable (Y_t at time t) is a linear function of the parameters and defined as $\mu(x) = E(Y_t|X = x)$. The linear regression parameters are estimated by minimizing the squared error $E((Y_t - \mu(x))^2|X = x)$ and are known as ordinary least square estimates. In linear quantile regression the conditional quantile which is defined as $Q_\tau(Y_t|X = x)$ where $0 < \tau < 1$ and $Pr(Y_t \leq Q_\tau) = \tau$ is a linear function of the parameters. The parameters $\beta(\tau)$ of a quantile regression defined by $Q_\tau(Y_t|X = x) = x^T\beta(\tau)$ are estimated by minimizing $\sum_t \rho_\tau(y_t - x^T\beta)$ where ρ_τ is the check function defined as $\rho_\tau(z) = z(\tau - I(z < 0))$, and $I(\cdot)$ is indicator function Koenker and Hallock (2001). In the special case when the quantile is restricted to the median, regression based on absolute deviation loss function may

estimate the conditional median of Y_t given X Koenker et al. (2017).

2.1 Segmented Regression

Wagner et al. (2002) described a segmented regression model to estimate the sudden change and gradual changes after a single intervention at time t using the following expression

$$y_t = \beta_0 + \beta_1 * (time)_t + \beta_2 * Intervention_t + \beta_3 * time.after.intervention_t + \epsilon_t$$

where y_t is dependent variable; time is a discrete variable starting from 1 for the beginning of the observation period; intervention is a dummy variable being zero for time before intervention and 1 for time after intervention; time after is also a discrete variable being zero before intervention and starting from 1 counting the number of time points after the intervention. The intercept β_0 denotes the mean outcome at baseline (time zero), β_1 denotes the pre-intervention slope, β_2 denotes the sudden change due to intervention at time t , β_3 denotes the gradual change (i.e. difference between pre-intervention and post-intervention slope) due to intervention. We assume that the error ϵ_t is distributed as $F(\cdot)$ with mean zero.

2.2 Segmented Quantile Regression

As an extension to earlier work on segmented regression by Wagner et al. (2002), we formulate segmented quantile regression model as

$$Q_{Y_t}(\tau|\mathcal{F}_{t-1}) = \beta_0(\tau) + \beta_1(\tau) * (time)_t + \beta_2(\tau) * Intervention_t + \beta_3(\tau) * time.after.intervention_t + F_\tau^{-1}$$

where \mathcal{F}_{t-1} the σ -field containing information up to time $t - 1$. We assumed that the conditional distribution of Y_t given past information is location shift with respect to the inverse of the error distribution F_τ^{-1} evaluated at some fixed quantile (τ). As an example, for $\tau = 0.50$ the parameter $\beta_0(\tau)$ estimates the outcome median (τ^{th} quantile) at time 0 (intercept before intervention); $\beta_1(\tau)$ the slope of the median (τ^{th} quantile) regression line before intervention estimates the change in the median (τ^{th} quantile) of the outcome with each time point increase before intervention; $\beta_2(\tau)$ estimates the level change in the median (τ^{th} quantile) of the outcome immediately after the intervention (sudden/level change); $\beta_3(\tau)$ estimates the change in the slope of the median (τ^{th} quantile) of the outcome (gradual change) after intervention.

2.3 Simulation

In the Monte Carlo simulation study, we restricted our attention to the special case of segmented quantile regression where the median slopes were the same as mean slopes. In particular, we compared the finite sample performance of segmented quantile regression against the segmented linear regression using the location shift of symmetric error distribution. The error terms are standard normal $N(0,1)$, heteroskedastic ($N(0, \log^2(\text{time} + 1))$), auto-correlated ($AR(1)$), non-normally distributed (t -distribution with degrees of freedom 1, 2, and 3). Given estimates of the parameters of the segmented quantile and linear regression models, we compare statistical properties of the estimated parameters, such as: Type 1 error, bias, power, and MSE. The heteroskedastic residual is generated using $\log^2(\text{time} + 1) * \epsilon_t$ where $\epsilon_t \sim N(0,1)$ and the heavy tail errors are generated using t -distribution with 1, 2, and 3 degrees of freedom.

We investigated finite sample properties of the estimated parameters by varying the sample size within each unique simulation experiment. In all scenarios, there is one intervention and the sample size before and after intervention are the same. The sample sizes in terms of total data points with respect to time t were specified to be: 20 (10 before and 10 after intervention), 40, 60, and 80. For each scenario, we used 5000 Monte Carlo replicates to ensure that the type I error rates can be evaluated under tight bounds (0.05 ± 0.0031) Burton et al. (2006). For sudden change scenario the coefficient for the intervention $\beta_2(\tau)$ is non-zero and the coefficient for time-after $\beta_3(\tau)$ is zero. For gradual change scenario the coefficient for intervention $\beta_2(\tau)$ is zero and the coefficient for time-after $\beta_3(\tau)$ is non-zero. For both sudden and gradual change both intervention $\beta_2(\tau)$ and time-after coefficients $\beta_3(\tau)$ are nonzero. The parameter configurations of $\beta_0(\tau), \beta_1(\tau), \beta_2(\tau)$ and $\beta_3(\tau)$ are (i) (0, 0, 0, 0) for no overall effect; (ii) (0.5, -0.3, 0, 0) for no level and gradual change; (iii) (0.5, -0.3, 1.5, 0) for level change; (iv) (0.5, -0.3, 0, 0.5) for gradual change; (v) (0.5, -0.3, 1.5, 0.5) for both level and gradual changes..

The visual representation of these five scenarios is shown in Figure 1. All computations are performed in SAS version 9.4 M6. PROC REG is used to fit linear segmented regression models. PROC QUANTREG is used to fit quantile regression models, with CI=Sparsity option for estimating the standard errors. The sparsity function is the reciprocal of the density function which is estimated by assuming that the conditional quantile function is locally linear Astivia and Zumbo (2019). As part of the sensitivity analysis, confidence intervals were also computed using the resampling (bootstrap-based) option in PROC QUANTREG. The complete set of simulation results for all investigated scenarios (including the sensitivity analysis) are provided in the supplementary Table 1.




Figure 1: Five scenarios of simulation study

3 Results

The results of the simulation study based on five parameter configurations are given in Figure 3 through Figure 8 that contains type 1 error, power, bias and mean square error (MSE). Type-1 error, power, bias, and MSE are estimated by averaging over 5,000 replicated simulation experiments.

3.1 Type 1 Error

We evaluate type 1 error rates when the true parameter values are assumed to be zero. Figure 3 describes the type I error rate (with nominal level $\alpha=0.05$) for segmented linear regression and segmented quantile regression. We ignored the estimates of the intercept because this parameter is of least interest when we fit a segmented linear or quantile regression. The lines L-time and Q-time represent the slope before intervention, L-intervention and Q-intervention represent the level change, and L-time-after and Q-time-after represent the gradual changes in linear and quantile regressions.

When the errors are IID standard normal, type 1 errors of both linear and quantile regression models are close to nominal 5 percent. When the errors are heteroskedastic, type 1 errors for quantile regression are below 5 percent while the estimate for trend in linear model has the worse underestimated type 1 error. When the error distribution follows an AR(1) process both linear and quantile regression over estimate type 1 error; with over-estimation being more pronounced in the linear regression model. When residuals are heavy tailed quantile regression slightly underestimates the type 1 errors;

whereas, linear regression shows type-1 error rates closer to 5% level.

3.2 Power

To further assess the performance of segmented linear and quantile regression we considered power under the alternative hypothesis (i.e. when regression parameters are non-zero in the simulation experiment).

3.2.1 Level Change Only

When the true parameter for level change is not zero the power of linear and quantile segmented regression for detecting non-null pre-intervention slope and level change is shown Figure 4. When the errors are IID standard normal the linear regression model is more powerful than the quantile regression model. When the errors are heteroskedastic both linear and quantile regression had low power for detecting the level change while their power for estimating trend before intervention is very similar. When errors follow an AR(1) model linear regression appears to perform better than quantile regression. The apparent power gains for linear regression (relative to quantile regression) are possibly a result of an under-estimation of standard errors (as apparent from the inflated type-1 errors, shown in Figure 3). Quantile regression appears to have increased power over linear regression at detecting intervention effects, when data arise from a heavy tailed generating mechanism.

3.2.2 Gradual change only

When the true parameter for gradual change is not zero then power of linear and quantile segmented regression for detecting the gradual change is shown in Figure 5. When the errors are IID standard normal, heteroskedastic or AR(1) the linear model performs slightly better than quantile for small sample sizes (i.e. $N=20$, $N=40$). The higher power for linear regression when errors are AR(1) may be related to underestimated variances, and inflated type-1 errors (as seen in Figure 3). When the distribution of the errors follow a t-distribution with one degree of freedom the quantile regression has increased power compared to the linear regression model.

3.2.3 Both Level and Gradual Changes

When both the true parameters for level change and gradual change are not zero we can assess and compare the power of linear and quantile regressions as shown in

Figure 6. When the errors are IID standard normal linear regression demonstrates increased power compared to quantile regression (especially for small sample sizes). When the errors are heteroskedastic both linear and quantile regressions have low power for detecting level change effects and modest power for estimating gradual change. The linear model with AR(1) errors appears to perform better than quantile regression (especially at small sample sizes). Again, the apparent power gains for the linear regression model compared to the quantile regression model may be related to variance under-estimation and type-1 error inflation (as seen in Figure 3). The quantile regression models has superior power compared to the linear regression model when the error distribution is $T(df=1)$. Power appears similar for linear regression and quantile regression when the errors are distributed $T(df=3)$.

3.3 Bias

Bias is defined as the difference between the average of the estimated parameter and the true parameter. When the errors are $T(df=1)$ distributed, parameter estimates from the linear segmented regression model are severely biased; whereas parameter estimates for the quantile segmented regression model are unbiased. We excluded the results for t-distribution with one degrees of freedom and plotted the biases for the remaining simulation experiments in Figure 7. Considering the five simulation experiments presented in Figure 7 below, the absolute biases are less than 0.05 for level and gradual change using linear and quantile regression. For normal and AR(1) residuals linear model has smaller bias compare to quantile regression. When the error distribution is heteroskedastic quantile regression performs better than linear regression. For t-distribution with 2 degrees of freedom linear regression has worse bias compare to quantile regression. In all scenarios the bias improves as the sample size increases.

3.4 Mean Square Error

The mean square error (MSE) is defined as the sum of bias squared and variance of the estimator parameter. Again, MSE is severely distorted for estimating a linear model, when data are simulated according to a $T(df=1)$ distribution. Figure 8 shows the MSE for parameter estimates and error distributions excluding t-distribution with 1 degrees of freedom for the scenario of level and gradual change.

As sample size increases the MSE for both linear and quantile segmented regression decreases. The MSE for level change in linear and quantile segmented regression were

higher than the MSE for the other estimated parameters. The MSE of the segmented quantile regression estimated parameters were lower compared to the MSE of the estimated parameters of segmented linear regression when the distribution of the errors were t-distribution. The magnitude of the MSE decreased as the degrees of the freedom of the t-distribution increased. The MSE of the quantile segmented regression were higher than those of linear segmented regression for heteroskedastic residuals.

4 Application

Wagner et al. (2002) used the New Hampshire Medicaid programme interventions data for illustrating how linear segmented regression can estimate intervention effects in ITS designs. We explore the same data set, using segmented quantile regression. The outcome measure is the monthly mean number of prescriptions with two known intervention points, the first intervention occurring after 21 months and second intervention occurring after 32 months with a total of 48 months follow-up. We can assess the impacts of these two interventions using the following quantile segmented regression model:

$$Q_{Y_t}(\tau) = \beta_0(\tau) + \beta_1(\tau) * (time)_t + \beta_2(\tau) * Intervention.1_t + \beta_3(\tau) * time.after.intervention.1_t + \beta_4(\tau) * Intervention.2_t + \beta_5(\tau) * time.after.intervention.2_t$$

The estimated parameters and associated p-values for both the segmented linear regression model and segmented quantile regression model ($\tau=\{0.10, 0.25, 0.50, 0.75, 0.90\}$) are given in Table 1 and displayed in Figure 2. Point estimates and inferences regarding the statistical significance of parameter estimates are similar for linear regression and median quantile regression ($\tau=0.50$). This suggests similar impacts of time/intervention on the conditional mean and conditional median functions, respectively. Noticeable differences in estimated regression parameters are observed, depending on whether the analyst is modelling extreme vs. central conditional quantiles (e.g. the intercept, intervention-1 and intervention-2 parameters). Differences in estimated time slopes are more modest, depending on the conditional quantile of interest (e.g. time, time-after-1, and time-after-2 parameters). Practically speaking, if we consider the estimated quantile regression parameters for intervention-1 and intervention-2, we notice a trend where covariate effects are larger in magnitude for larger values of τ . This is suggestive that the impact of interventions varies, depending on whether one has a small vs. large mean daily volume of prescriptions.

Table 1: Parameters estimates and their p-values for different values of quantiles

Model	Intercept	Time	Intervention 1	Time After 1	Intervention 2	Time After 2
Linear	5.05(<0.0001)	0.02(0.1166)	-2.65(<0.0001)	0.00(0.9842)	0.54(0.0077)	0.06(0.0366)
$\tau = 0.10$	4.75(<0.0001)	0.01(0.0063)	-2.42(<0.0001)	0.00(0.8865)	0.21(0.4656)	0.08(0.0056)
$\tau = 0.25$	5.07(<0.0001)	0.00(0.7757)	-2.33(<0.0001)	0.02(0.2198)	0.45(0.1067)	0.05(0.0390)
$\tau = 0.50$	5.17(<0.0001)	0.00(0.8540)	-2.44(<0.0001)	0.02(0.1333)	0.57(<0.0001)	0.05(0.0043)
$\tau = 0.75$	5.14(<0.0001)	0.03(0.2332)	-2.94(<0.0001)	-0.01(0.7711)	0.49(0.0115)	0.07(0.0094)
$\tau = 0.90$	5.12(<0.0001)	0.05(0.0003)	-3.25(<0.0001)	-0.05(0.2618)	0.83(0.0187)	0.07(0.1711)

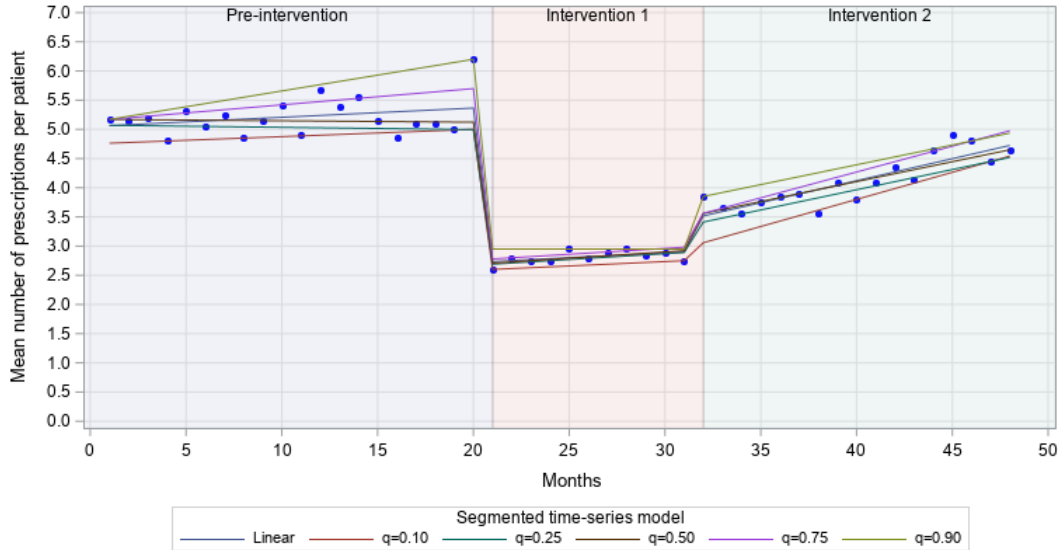
This applied data analytic example nicely illustrates the utility of quantile regression for assessing covariate effects on any particular conditional quantile distribution. And further, how using segmented quantile regression under an ITS design can afford improved understanding on how interventions impact these particular conditional quantile distributions. In this case, inferences regarding time and intervention effects are similar when comparing segmented linear regression and segmented median regression ($\tau=0.50$); however, a more holistic understanding of intervention effects is achieved when one investigates impacts across a family of conditional quantile functions (e.g. $\tau=\{0.10, 0.25, 0.50, 0.75, 0.90\}$). This demonstrates that using segmented quantile regression in ITS designs can yield useful and complementary estimates of intervention effects, compared to what would be understood if the analyst were to fit only a segmented linear regression model.

5 Discussion

In both our Monte Carlo simulation study and our applied data analytic illustration we have demonstrated the feasibility and usefulness of the quantile segmented regression as an alternative and complementary statistical tool for assessing the impacts of interventions by modelling the conditional median of the response, given measured features/-covariates. Applying quantile regression in the context of interrupted time series designs allows the analyst to estimate the impact of covariates on any conditional quantile of the outcome/response variable. This may permit more nuanced understanding of the impact of various interventions, then can be obtained through modelling of the conditional mean, or conditional median function alone.

When errors are distributed IID according to a Gaussian process, we observe that segmented linear regression and segmented quantile (median) regression perform

Figure 2: Segmented time-series models applied to Wagner et al. (2002) data set



similarly well. Further, the conditional mean of the response is quite similar to the conditional median of the response, as the error variables are unimodal and symmetric. When data are generated from an IID normal process, it can be shown analytically that the sample mean is more efficient than the sample median (with respect to estimation of the location parameter in the Normal model). In other words, the variance of the sample median is larger than variance of the sample mean, asymptotically. In finite samples, the efficiency gains of the sample mean relative to the sample median are further exaggerated Sherman (1997). We observe similar findings in our simulation; however, from a practical point of view, both segmented linear regression and segmented quantile regression perform similarly well in our simulation experiments.

When errors are distributed according to an AR(1) process we note that both segmented linear regression and segmented quantile regression struggle with inflated type-1 error rates in our Monte Carlo simulation study. The problem is more pronounced in the case of the linear regression models, as compared to the quantile regression models. When dealing with any time series regression models, it is prudent to post-hoc investigate the distribution of the models residuals. This can be done visually by inspecting the auto-correlation function (ACF) and the partial auto correlation

function (ACF) plots. In addition, hypothesis tests for assessing the extent of auto-correlation in the residuals exist, for example, the Durbin-Watson test, and others. If these post-hoc tests/visualizations are suggestive of auto-correlated residuals, then the analyst likely needs to consider more complicated models than the simple segmented linear/quantile regression models presented in this paper. For example, the segmented linear regression model can be easily extended with modern software to account for auto-correlated errors (e.g. AR, MA, ARMA errors). Extending quantile regression to account for auto-correlated errors requires more attention. Theoretical work extending quantile regression to deal with correlated errors is discussed in chapter 17 of Koenker et al. (2017). As far as the authors are aware, readily available software implementations of the cited methods are not currently available for segmented quantile regression.

When errors are heteroskedastic, naïve implementations of segmented linear regression perform poorly as compared with segmented quantile regression. Dealing with heteroskedastic error distributions is one of the earliest applications of quantile regression, and is discussed in (Koenker et al. , 2017, p. 12). If post-hoc investigation of residuals are suggestive of heteroskedastic errors, then one might extend the linear segmented regression model by i) weighting influential observations, or ii) incorporating robust standard error estimators Kezdi (2003).

When errors follow a heavy-tailed t-distribution we observe that segmented quantile regression often outperforms segmented linear regression. These results corroborate earlier findings in the literature where the inefficient estimation of segmented linear regression for heavy tail distribution is well documented Geraci and Bottai (2007). Again, the ability to make sound inferences in the presence of heavy tailed distributions is an advantage of quantile regression compared to linear regression. Albeit, the segmented linear regression model may be adapted by considering trimmed/windsorized loss functions, which down-weight the impact of outlying observations on fitted model parameters.

Sensitivity analyses were performed to compare the performance of sparsity-based variance estimation versus resampling-based (i.e. bootstrap-based) variance estimation using PROC QUANTREG. Although sparsity based estimation is computationally more efficient than the bootstrap-based resampling method, it was clearly not robust to data violating the assumption of IID errors (e.g. AR(1) errors). In contrast to the inflated Type I error rates of sparsity method, the Type I error rates were close to the 5% nominal level for small number of time points (i.e. $N=20$) using the bootstrap-based estimation when the error terms were AR(1). The results of this simulation study also supported the earlier findings of improved empirical coverage probabilities of

the resampling methods as compared to direct sparsity method under smaller sample sizes Tarr (2012). In the case of heavy tail distribution (with $t\text{-df}=1,2,3$), the coverage probabilities of resampling method were close to 5% nominal level under larger sample sizes ($N>150$) but substantial deviation was evident with smaller samples. Since the focus of this study was restricted to assessing the finite properties of segmented “median” regression, we conjecture that the coverage probabilities at the extreme ends of quantile segmented regression may suffer from significant deviations from pre-specified nominal level as pointed out by Tarr (2012). Similar arguments may also hold for the power functions of estimated quantities of segmented quantile regression under different scenarios of level and/or gradual change.

In this paper, we have reviewed interrupted time series analysis as a powerful quasi-experimental design for investigating intervention effects. We have reviewed the literature related to segmented linear regression for ITS designs and suggested that segmented quantile regression models can be formulated in a similar manner (i.e. the same design matrix which yielded segmented linear regression can be used to estimate segmented quantile regression models). We conducted a Monte Carlo simulation experiment to investigate the finite sample properties of both segmented linear regression and segmented quantile regression under a variety of data generating mechanisms. When the data generating mechanism is IID Gaussian noise, both segmented linear/quantile regression perform satisfactorily. However, when error terms are auto-correlated, heteroskedastic, or heavy-tailed we notice that segmented linear regression can struggle. As such, the analyst who is comfortable fitting segmented linear regression, must check that certain model assumptions are valid before making particular conclusions/inferences. Further, we investigate whether segmented quantile regression may act as a blanket solution to the problems encountered in these scenarios. In general, when the errors are heteroskedastic or heavy tailed we note that segmented quantile regression out-performs segmented linear regression. When the errors are auto-correlated neither the segmented linear regression model nor the segmented quantile regression model perform well and model extensions are needed to deal with the auto-correlated errors. In general, segmented quantile regression offers the unique opportunity to assess the variation of the impacts of the interventions across the population quantiles (not simply the population mean/median); for certain subject matter problems this may be a major advantage of segmented quantile regression.

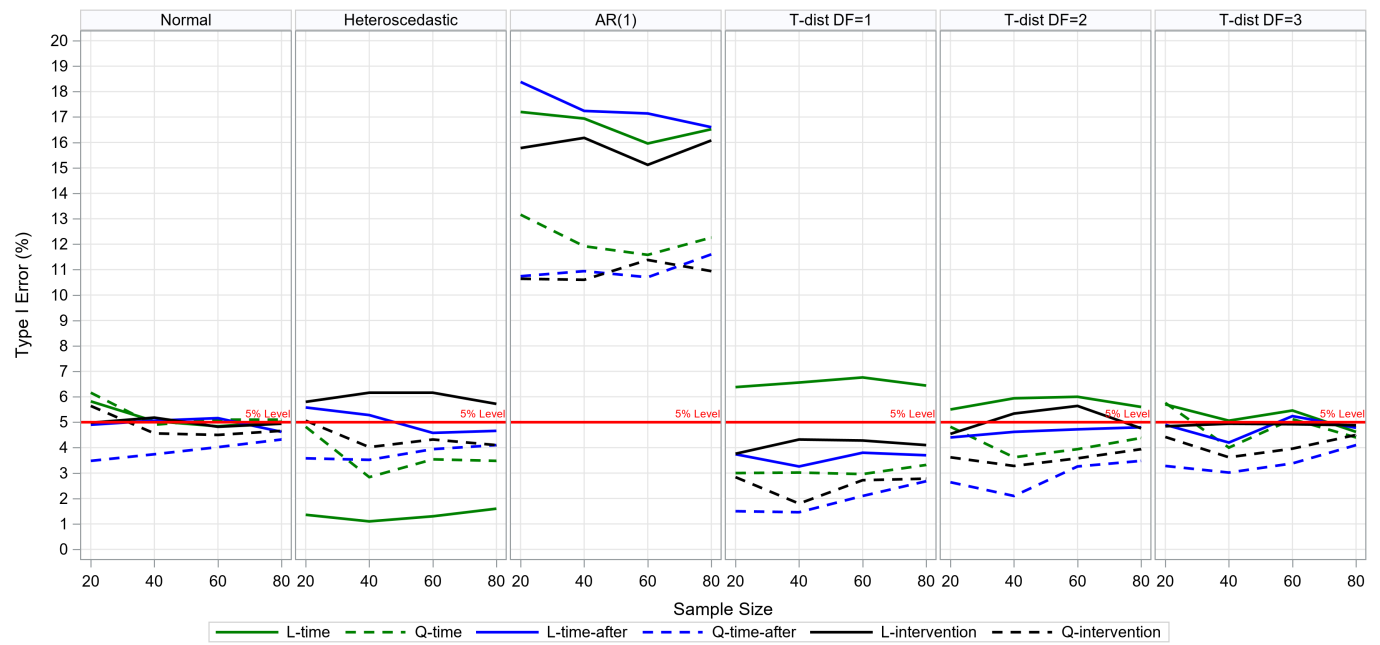


Figure 3: Type I error rates for linear and quantile segmented regression models




Figure 4: The performance metric of power for the scenario of level change




Figure 5: The performance metric of power for the scenario of gradual change




Figure 6: The performance metric of power for the scenario of level and gradual change

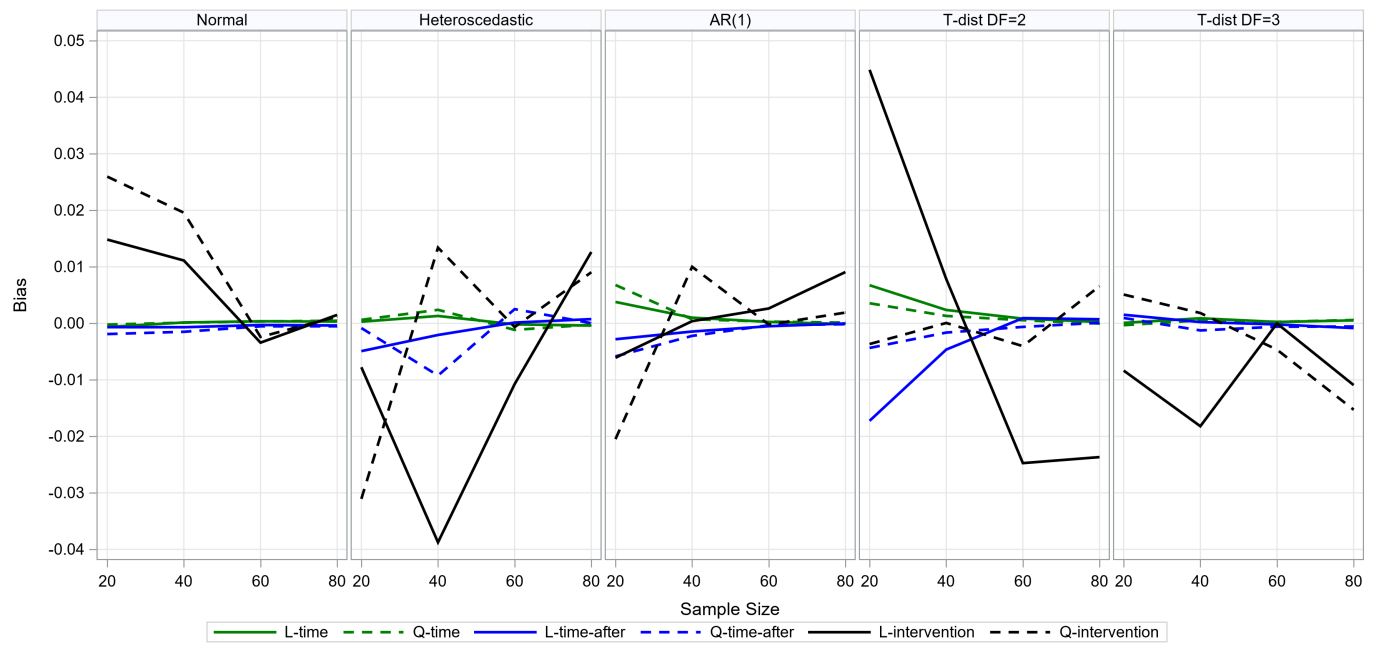


Figure 7: The performance metric of bias for the scenario of level and gradual change

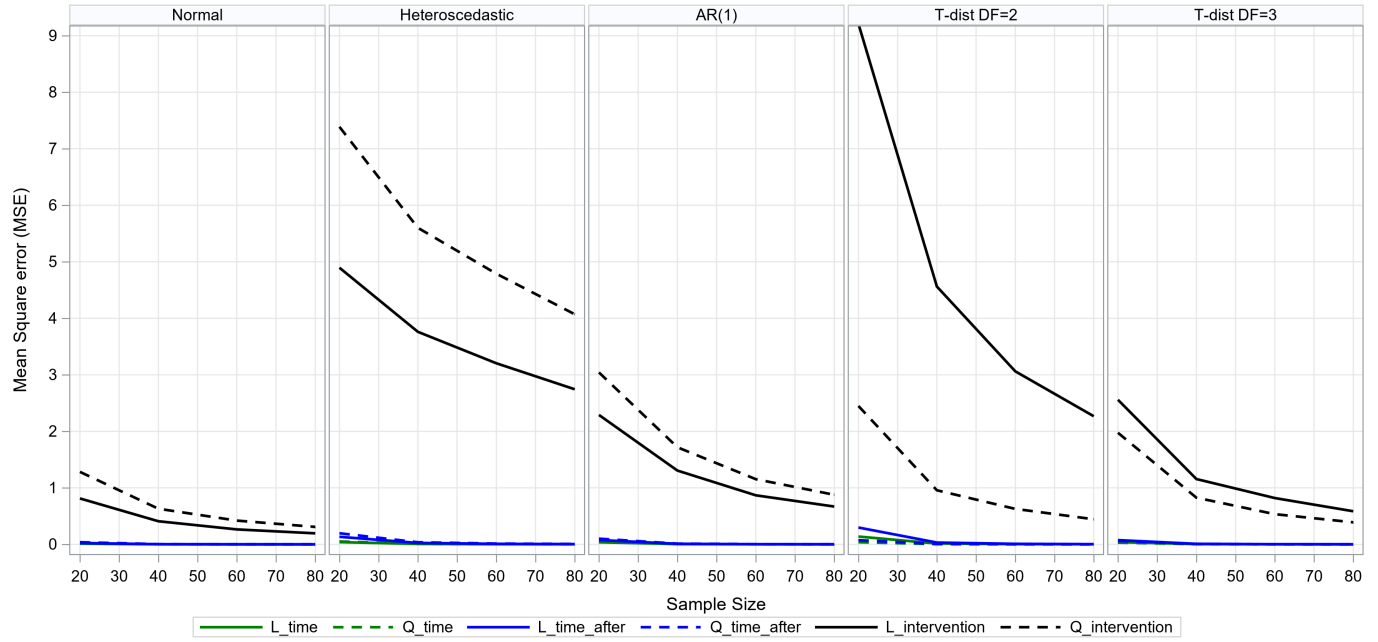


Figure 8: The performance metric of Mean Square error for the scenario of level and gradual change

6 Conclusion

Segmented quantile regression is a useful and powerful statistical method for assessing the impact of interventions in ITS designs. Segmented quantile regression models can be easily fit and interpreted, affording a unique understanding of the impacts of interventions on different conditional quantiles of the outcome variable. Future work may focus on the estimation of quantile segmented regression models using the rank-score method which relies on the estimation of confidence interval by inverting the appropriate test statistics Kocherginsky et al. (2005).

References

- Astivia, O. L. O., and Zumbo, B. D. (2019), Heteroskedasticity in Multiple Regression Analysis: What it is, How to Detect it and How to Solve it with Applications in R and SPSS. *Practical Assessment, Research, and Evaluation*, **24**(1).
- Austin, P. C., and Tu, J. V., and Daly, P. A., and Alter, D. A. (2005), The use of quantile regression in health care research: a case study examining gender differences in the timeliness of thrombolytic therapy. *Statistics in medicine*, **24**(5), 791-816.
- Baltagi, B. H. (2010), *Econometrics*. Springer.
- Burton, A., and Altman, D. G., and Royston, P., and Holder, R. L. (2006), The design of simulation studies in medical statistics. *Statistics in medicine*, **25**(24), 4279-4292.
- Campbell, D. T., and Stanley, J. C. (2015), *Experimental and quasi-experimental designs for research*. Ravenio Books.
- Cook, T. D., and Campbell, D. T., and Day, A. (1979), *Quasi-experimentation: Design & analysis issues for field settings*. Houghton Mifflin Boston.
- Ferron, J., and Rendina-Gobioff, G. (2005), Interrupted time series design. *Encyclopedia of Statistics in Behavioral Science*.
- Geraci, M., and Bottai, M. (2007), Quantile regression for longitudinal data using the asymmetric Laplace distribution. *Biostatistics*, **8**(1), 140-154.
- Gillings, D., and Makuc, D., and Siegel, E. (1981), Analysis of interrupted time series mortality trends: an example to evaluate regionalized perinatal care. *American journal of public health*, **71**(1), 38-46.

- Honda, T. (2013), Nonparametric quantile regression with heavy-tailed and strongly dependent errors. *Annals of the Institute of Statistical Mathematics*, **65**(1), 23-47.
- Kezdi, G. (2003), Robust standard error estimation in fixed-effects panel models. Available at SSRN 596988.
- Kocherginsky, M., and He, X., and Mu, Y. (2005), Practical confidence intervals for regression quantiles. *Journal of Computational and Graphical Statistics*, **14**(1), 41-55.
- Koenker, R. (2005), *Quantile Regression*. Cambridge.
- Koenker, R., and Bassett Jr, G. (1978), Regression quantiles. *Econometrica: journal of the Econometric Society*, 33-50.
- Koenker, R., and Chernozhukov, V., and He, X., and Peng, L. (2017), *Handbook of quantile regression*. CRC press.
- Koenker, R., and Hallock, K. F. (2001), Quantile regression. *Journal of economic perspectives*, **15**(4), 143-156.
- Nolan, J., P and Ojeda-Revah, D. (2013), Linear and nonlinear regression with stable errors. *Journal of Econometrics*, **172**(2), 186-194.
- Sherman, M. (1997), Comparing the sample mean and the sample median: An exploration in the exponential power family. *The American Statistician*, **51**(1), 52-54.
- Tarr, G. (2012), Small sample performance of quantile regression confidence intervals. *Journal of Statistical Computation and Simulation*, **82**(1), 81-94.
- Wagner, A. K., and Soumerai, S. B., and Zhang, F., and Ross-Degnan, D. (2002), Segmented regression analysis of interrupted time series studies in medication use research. *Journal of clinical pharmacy and therapeutics*, **27**(4), 299-309.
- Xiong, W., and Tian, M. (2019), Weighted quantile regression theory and its application. *Journal of Data Science*, **17**(1), 145-160.
- Yang, X. R., and Zhang, L. X., (2008), A note on self-weighted quantile estimation for infinite variance quantile autoregression models. *statistics & probability letters*, **78**(16), 2731-2738.