

# UC Irvine

## UC Irvine Previously Published Works

### Title

Assessing health care interventions via an interrupted time series model: Study power and design considerations

### Permalink

<https://escholarship.org/uc/item/21n68797>

### Journal

Statistics in Medicine, 38(10)

### ISSN

0277-6715

### Authors

Cruz, Maricela  
Gillen, Daniel L  
Bender, Miriam  
[et al.](#)

### Publication Date

2019-05-10

### DOI

10.1002/sim.8067

Peer reviewed



Published in final edited form as:

*Stat Med.* 2019 May 10; 38(10): 1734–1752. doi:10.1002/sim.8067.

## Assessing health care interventions via an interrupted time series model: Study power and design considerations

Maricela Cruz<sup>1</sup>, Daniel L. Gillen<sup>1</sup>, Miriam Bender<sup>2</sup>, Hernando Ombao<sup>1,3</sup>

<sup>1</sup>Department of Statistics, University of California, Irvine, California

<sup>2</sup>Sue and Bill Gross School of Nursing, University of California, Irvine, California

<sup>3</sup>Statistics Program, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

### Abstract

The delivery and assessment of quality health care is complex with many interacting and interdependent components. In terms of research design and statistical analysis, this complexity and interdependency makes it difficult to assess the true impact of interventions designed to improve patient health care outcomes. Interrupted time series (ITS) is a quasi-experimental design developed for inferring the effectiveness of a health policy intervention while accounting for temporal dependence within a single system or unit. Current standardized ITS methods do not simultaneously analyze data for several units nor are there methods to test for the existence of a change point and to assess statistical power for study planning purposes in this context. To address this limitation, we propose the “Robust Multiple ITS” (R-MITS) model, appropriate for multiunit ITS data, that allows for inference regarding the estimation of a global change point across units in the presence of a potentially lagged (or anticipatory) treatment effect. Under the R-MITS model, one can formally test for the existence of a change point and estimate the time delay between the formal intervention implementation and the over-all-unit intervention effect. We conducted empirical simulation studies to assess the type one error rate of the testing procedure, power for detecting specified change-point alternatives, and accuracy of the proposed estimating methodology. R-MITS is illustrated by analyzing patient satisfaction data from a hospital that implemented and evaluated a new care delivery model in multiple units.

### Keywords

change-point detection; complex interventions; patient satisfaction; power analysis; segmented regression; time series

## 1 | INTRODUCTION

The delivery and assessment of quality health care is increasing in complexity. Now, more than ever, patients, providers, resources and contexts of care interact in dynamic ways to produce various measurable health outcomes that, oftentimes, do not align with

expectations.<sup>1</sup> Assessing the impact of health interventions on patient health outcomes is therefore inherently difficult with regard to research design and statistical analysis.<sup>2,3</sup> Interrupted time series (ITS) designs borrow from traditional case-crossover designs and function as quasi-experimental methodology that allows each sampling unit to serve as its own control without stripping contextual and temporal factors from the analysis.<sup>4</sup> Current standardized methods for analyzing ITS designs do not borrow information across units. This is a serious limitation because it does not take advantage of all available data that may provide information on the lag associated with a given intervention. A main contribution of the work presented here are the empirical power studies that illustrate the gain in efficiency obtained by borrowing information across units.

The methodology presented in this paper is motivated by our interest in estimating the lagged effect of an intervention on average patient satisfaction survey scores, recorded monthly at five clinical care units. A time series plot of patient satisfaction scores from January 2008 to December 2012 at two hospital units (the Stroke and Surgical units) is given in Figure 1. There seems to be a change in the mean functions of the Stroke and Surgical units around the middle of the time series, slightly before the formal implementation of the intervention on July 2010. The time series data are from a study aimed to assess the impact of a new nursing care delivery system on publicly recorded standardized quality and safety metrics.<sup>5</sup> These metrics are a central area for improvement because the Center for Medicaid and Medicare Services (CMS) Value-Based Purchasing Program utilizes them for health systems' care services reimbursement.<sup>6</sup>

The intervention was the implementation of Clinical Nurse Leader (CNL) integrated care delivery, a nursing model that embeds a master prepared nurse into the front lines of care.<sup>7</sup> The nurses, referred to as CNLs, have advanced competencies in clinical leadership, care environment management, and clinical outcomes management.<sup>7</sup> The CNLs, conducting their master's level microsystem change project, were introduced into their respective hospital units on January 2010, six months prior to the formal intervention implementation time. This may or may not have influenced the change point of the intervention effect. Namely, because of this early introduction, the estimated change point may have occurred up to 6 months prior to the formal intervention time point. We are therefore interested in estimating the time lag (or delay) between the onset of the intervention—ie, the change point—and the effect on patient satisfaction. Our model assumes a global change point rather than unit-specific change points (1) to pool information across hospital units and increase efficiency and (2) to reduce the impact of unit-specific high-leverage points around the CNL and formal intervention implementation time points. Importantly, we are interested in examining whether or not a change point actually exists, thereby deducing whether or not the intervention impacts patient satisfaction. Our interest is not solely on properly modeling the CNL intervention; we are also interested in future study designs, and so, focus on power.

The most utilized statistical methodology for analyzing ITS data in the health care literature is segmented regression.<sup>8-11</sup> Segmented regression restricts the analysis to one health care outcome for one unit (group or cluster). In the context of assessing the above intervention, perhaps, a severe drawback of segmented regression is that it restricts the interruption to a pre-determined time point in the series or censors data by removing the set of time points for

which the intervention effects may not be realized. Additionally, segmented regression neglects the plausible differences in autocorrelation and variability between the pre- and post-intervention phases present in the data. The Robust-ITS model proposed by Cruz et al treats the change point as variable, appropriate for situations where the data warrants such treatment, and tests for differences in autocorrelation and variability pre- and post-change point.<sup>12</sup> Nevertheless, Robust-ITS and segmented regression both neglect shared information across hospital units and inherently assume a change point *exists*.

Assessing the impact of an intervention with traditional segmented regression or Robust-ITS on these data requires a separate analysis for each individual. We expect many of the units to share several characteristics—ie, abide by the same regulations, have similar schedules, hire staff based on the same criteria, etc—because the units are housed within one hospital. Moreover, we expect the CNL “training” or education for each of the CNL students to include commonalities, such as course work and care delivery ideology. Assessing the intervention impact on multiple units via current segmented regression methods ignores shared characteristics across units, in particular, the similarity between characteristics influencing the change point.

Inherently assuming that a change point exists, as in segmented regression and Robust-ITS, may lead to erroneous results when there is no actual a change point. Change-point models will forcefully quantify a change in the outcome regardless of the presence of a true change point. This is a problem whether the change point is determined a priori or estimated over a set of possible change points. Assuming a change point exists, when it truly does not, will force a model to provide an estimate of an artificial difference in the outcome. To avoid incorrectly specifying an unnecessary change point and regression to the mean phenomena, we focus on formally testing for the existence of a change point.

In this paper, we develop the Robust Multiple ITS model (R-MITS), a novel extension of Robust-ITS, appropriate for multiple independent ITS. Furthermore, we present the supremum Wald test (SWT), able to test for the existence of a change point across units. Importantly, we provide empirical type one error, power, and accuracy studies assessing the operating characteristics of our developed methodology. The proposed method (a) borrows information across hospital units to increase efficiency, (b) estimates a global change point of an instituted intervention, (c) formally tests for the existence of a change point in the unit-specific mean functions, and (d) allows for changes in the mean functions and autocorrelation structures across units.

We go on to describe our proposed R-MITS model and provide details on the estimation and inference procedures. In our model specification, we outline the SWT used to determine the necessity of a change point. Next, we present empirical simulations to assess the type one error, power for detecting specified change-point alternatives, and accuracy of the change-point estimation procedure. We then analyze the impact of CNL integrated care delivery on patient satisfaction. Lastly, we present a summary of our developed methodology and briefly describe future work.

## 2 | THE ROBUST MULTIPLE ITS MODEL

Our proposed model tests for the existence—rather than merely assume—of a change point and adequately manages multiple units/time series. A noteworthy feature of our approach is the clear distinction between the time of intervention and the change point, as in Robust-ITS. Setting the change point to a pre-determined time may lead to incorrect measures of the intervention's effect on the system; particularly when set to the intervention time, because that does not necessarily represent the reality that complex interventions may have varied effects and take time to manifest change. Prevalent approaches to overcoming this limitation are to remove, or censor, a specific set of time points from the analysis.<sup>8,9</sup> R-MITS borrows information from all microsystems to estimate a global change point; ie, determines the time point at which the effect of the intervention initiates for the entire health system. Moreover, detecting differences in autocorrelation and variances pre- and post-intervention is critical in evaluating the effectiveness of an intervention. The R-MITS model allows for two completely different data dependency and variability structures to exist prior to the intervention and post-intervention within each unit.

To prelude model development, we plot the outcome against time to (a) illuminate the functional form of the longitudinal mean over time, (b) determine the presence of seasonality, and (c) further investigate the set of plausible change points and the necessity of a change point. If the functional form of the longitudinal mean is not linear, we transform the outcome to obtain a linear pattern, or apply a different segmented regression model appropriate for the pattern present within the ITS design. When needed, we account for seasonality via traditional statistical methods concisely described in the work of Bhaskaran et al.<sup>13</sup> Although not used in the analyses here, variance stabilizing transforms can be applied on the outcomes of interest if necessary. In our ITS data, the longitudinal mean functions are relatively linear in time with no apparent seasonality. Thus, no transformations are applied on the outcomes of interest.

### 2.1 | Description of R-MITS

Denote  $t^*$  as the time point at which the intervention is introduced and  $\tau$  as the time point at which the effect of the intervention initiates (the change point) for the outcome of interest. Sometimes, it may indeed be true that  $t^* = \tau$ , but this may not necessarily be true for all outcomes. Often, it is entirely possible that the time of effect of the intervention differs from the time of intervention introduction (ie, either  $\tau > t^*$  or  $\tau < t^*$ ). If  $\tau > t^*$ , then the effect of the intervention on the outcome is not realized until after the formal intervention time point. As it might be the case, when a learning effect exists with regard to the intervention, thereby leading to a delay in the realization of the full intervention impact. When  $\tau < t^*$ , there is an anticipatory intervention effect on the outcome. This may be the case in our motivating study, where CNLs are introduced into units prior to the formal intervention start time. We propose a data adaptive procedure for estimating and determining the existence of  $\tau$ , discussed in Section 2.1.3. Many change-point detection methods in time series exist, but often deal only with changes in the mean functions and variance (not the autocorrelation structure itself), and may not work well in shorter time series.<sup>14,15</sup> The method proposed in

this paper can suitably manage changes in the autocorrelation structure, as well as in the mean functions and volatility.

Define  $y_{jt}$  as the outcome of interest for hospital unit  $j$  at time  $t$  (where  $j = 1, \dots, J$  and  $t = 1, \dots, T$ ). For example,  $y_{jt}$  may be patient satisfaction for the Stroke unit at time  $t$ . The general regression is defined as

$$y_{jt} = \mu_{jt} + \epsilon_{jt}, \quad (1)$$

where  $\mu_{jt}$  is the mean function and  $\epsilon_{jt}$  is the stochastic process that models that fluctuations around the mean function. The mean component,  $\mu_{jt}$ , characterizes the mean function of the response for unit  $j$  during the pre-intervention and post-intervention phases. The stochastic process,  $\epsilon_{jt}$ , accounts for the variability and correlation of the outcome in the  $j$ th unit. In the following discussion, we define the mean functions and stochastic components for the R-MITS model, and the estimation procedures.

**2.1.1 | The pre- and post-intervention mean function**—The mean function of the outcome for hospital unit  $j$  at time  $t$  is

$$\mu_{jt} = \begin{cases} \beta_{j0} + \beta_{j1}t, & t < \tau \\ (\beta_{j0} + \delta_j) + (\beta_{j1} + \Delta_j)t, & t \geq \tau. \end{cases} \quad (2)$$

The parameters in  $\mu_{jt}$  are as follows: (1)  $\beta_{j0}$ , the intercept of the mean function prior to the change point; (2)  $\beta_{j1}$ , the slope of the outcome prior to the change point; (3)  $\beta_{j0} + \delta_j$ , the intercept of the post-intervention phase; (4)  $\beta_{j1} + \Delta_j$ , the slope of the post-intervention phase, for the outcome in unit  $j$ ; and (5)  $\tau$ , the global over-all-unit change point of the response. Thus,  $\delta_j = \Delta_j = 0$  implies that there is no change in the mean structure before and after time  $\tau$ . Health care specialists are primarily interested in testing for the intervention lag (delay in the effect of the intervention), and the differences in the outcome means between the pre- and post-change-point phases.

*Remark 1.* The metrics adopted by the health policy evaluation literature to assess the effect size of an intervention via ITS designs are the change in level and change in trend (or slopes). While the level change identifies the size of an intervention's effect, the change in trend quantifies the impact of the intervention on the overall mean function. It is necessary to report both level change and change in trend to interpret the results of an ITS study accurately.<sup>16</sup>

*Remark 2.* The level change is interpreted as the change in the anchored intercept (anchored at the change point) and is therefore the jump between the projected mean function based on the pre-change-point phase and the estimated mean function post-change point. In our model, the unit-specific change in level is defined mathematically as  $\delta_j + \beta_{j1}\tau$  and is graphically depicted in Figure 2. Trend change, or slope change, is denoted by  $\Delta_j$  in the mean function, Equation (2).

The mean function parameters are estimated simultaneously with the stochastic component parameters and change point, via maximizing the conditional likelihood given in Equation (4) of Section 2.1.3, with the autoregressive coefficients' estimator accounting for the volatility of the shifted series. An algorithm on how to precisely estimate the parameters is provided in Section 2.1.3.

**2.1.2 | Stochastic properties pre- and post-intervention**—The stochastic component in Equation (1),  $\epsilon_{jt}$ , captures the autocorrelation structure of the outcome variable across time for unit  $j$  and may change as a result of the intervention; the  $\epsilon_{jt}$  are zero-mean random fluctuations around the mean function of unit  $j$ . One goal of the CNL intervention is to increase the consistency of care delivery and hence patient assessment outcomes, (ie, decreasing variability of the outcomes). We therefore include separate stochastic components for the pre- and post-change-point phases, to allow for a change in outcome variability.

Due to the impact of the intervention, the stochastic process pre-intervention might differ from the stochastic process post-intervention. That is,  $\epsilon_{jt}$  for  $t \in \{1, \dots, \tau - 1\}$  may be a different stochastic process than  $\epsilon_{jt}$  for  $t \in \{\tau, \dots, T\}$ . Note that the length of the time series is denoted by  $T$ . Hence, the autocorrelation and variance might differ pre- and post-change point. Here, the stationarity requirement is satisfied if the variance, mean function, and autocorrelation are constant within each stochastic process, not constant across all time points.

In order to fit stationary AR or ARMA processes to the stochastic components, one should first confirm that there are no striking signs of nonstationarity. That is, the mean and variance of the residuals (obtained from modeling and removing the mean function as in the previous section) must be relatively constant. If the mean function is not misspecified, then the residuals should be fluctuating around zero without trend. Moreover, the residuals should be stationary within each of the pre- and post-intervention phases.<sup>17</sup> Our analysis of patient satisfaction suggests that it is reasonable to assume stationarity within each phase, and hence, we proceed with the assumption of stationarity.

In this work, we use the AR(1) process to model the stochastic component,  $\vec{\epsilon}_j = Y_j - \vec{\mu}_j$ , where  $Y_j = [y_{j2}, \dots, y_{jT}]'$  and  $\vec{\mu}_j = [\mu_{j2}, \dots, \mu_{jT}]'$  for unit  $j$ . Note that  $y_{j1}$  is not included in  $Y_j$  and  $\mu_{j1}$  is not included in  $\vec{\mu}_j$ , because we condition on the first observation. The AR(1) coefficient is estimated by maximizing the conditional likelihood with the denominator of the estimator averaging the volatility of the shifted AR(1) series. We therefore condition on the first observation  $y_{j1}$ . Since the mean function  $\vec{\mu}_j$  is not known (we only have its estimate,  $\widehat{\vec{\mu}}_j$ ), the stochastic component is not directly observed. Hence, we use the residuals  $R_j = Y_j - \widehat{\vec{\mu}}_j \equiv [r_{j2}, \dots, r_{jT}]'$  in place of  $\vec{\epsilon}_j$ . The residuals are modeled as

$$r_{jt} = \begin{cases} \phi_{j1} r_{j,t-1} + e_{jt,1}, & 1 < t \leq \hat{\tau} - 1 \\ \phi_{j2} r_{j,t-1} + e_{jt,2} & \hat{\tau} - 1 < t \leq T. \end{cases} \quad (3)$$

To ensure causality in the time series sense,  $\phi_{j1}$  and  $\phi_{j2}$  must lie in the interval  $(-1, 1)$  for all  $j$ . Note that the autoregressive coefficient prior to the change point,  $\phi_{j1}$ , is the correlation between time point  $t$  and  $t + 1$  (the adjacent correlation or autocorrelation) where  $t$  and  $t + 1$  belong to the pre-change-point phase ( $t$ , and  $t + 1 \in \{1, \dots, \tau - 1\}$ ), and  $\phi_{j1}^{|h|}$  is the correlation between two time points  $h$  units away (say  $t$  and  $t + h$ , both in the pre-change-point phase) of the outcome. The autoregressive coefficient post-change point,  $\phi_{j2}$ , has a similar interpretation. The zero-mean random fluctuations of model (3) are white noise,  $e_{jt,i} \stackrel{\text{iid}}{\sim} N(0, \sigma_{jw,i}^2)$  for  $i \in \{1, 2\}$ . The variance of the distribution of the response at any time point  $t$  is  $\sigma_{ji}^2 = \frac{\sigma_{jw,i}^2}{1 - \phi_{ji}^2}$  for  $i \in \{1, 2\}$ .

The variance and autoregressive coefficients in the AR(1) setting can be estimated by maximizing the conditional likelihood provided in Equation (4) of Section 2.1.3, with the autoregressive coefficients' estimator accounting for the volatility of the shifted series. The structure of the variance-covariance matrix and the estimators of the autoregressive coefficients and white noise standard deviations are given in the Appendix.

To determine whether the stochastic process differs as a result of the change point for each unit, one can test the hypothesis that  $v_j \equiv \phi_{j2} - \phi_{j1}$  equals zero. This can be tested by either estimating  $v_j$  directly or by conducting an F-test for nested models. The F-test for nested models for the AR(1) scenario is described by Cruz et al.<sup>12</sup>

**2.1.3 | Estimation of the change point and model parameters**—In this paper, we propose a conditional likelihood procedure for estimating the global change point. The set of possible change points is established by the researcher. We estimate the change point and therefore all of the parameters, both from the mean functions and stochastic components, simultaneously by obtaining the generalized least squares (GLS) estimates. Then, we test for the existence of a change in the mean functions—ie, we test the null hypothesis that there is no change in any of the mean functions versus the alternative that there is a change in at least one of the mean functions—at each possible change point by applying the SWT, described in Section 2.1.4.

Define the length of the time series as  $T$ , the number of units as  $J$ , the vector of mean function parameters as  $\theta_j = [\beta_{j0}, \beta_{j1}, \delta_j, \dots, \gamma_j]'$ , and  $\Sigma_j$  as the variance-covariance matrix of the response in unit  $j$ . The structure of the variance-covariance matrices is included in the Appendix.

Let  $q$  be a candidate change point in the set of possible change points  $Q$ , where  $Q = \{t^* - m, \dots, t^*, \dots, t^* + k\}$  for positive integer values of  $m$  and  $k$  set by the researcher. Recall the response vector for unit  $j$  is  $Y_j = [y_{j2}, \dots, y_{jT}]'$ . Note that  $y_{j1}$  is not included in  $Y_j$  because we model the zero-mean random fluctuations around the mean functions as AR(1) processes. For each candidate change point  $q \in Q$ , we derive the conditional likelihood function, conditional on the first observations,



$$L(\theta_1, \Sigma_1, \dots, \theta_J, \Sigma_J | q, Y_1, \dots, Y_J) \equiv \prod_{j=1}^J \left( \frac{1}{\sqrt{2\pi}} \right)^{T-1} |\Sigma_j|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (Y_j - \mathbf{X}(q)_j \theta_j)' (\Sigma_j)^{-1} (Y_j - \mathbf{X}(q)_j \theta_j) \right\}, \quad (4)$$

where

$$\mathbf{X}(q)_j \equiv \begin{bmatrix} 1 & 2 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & q-1 & 0 & 0 \\ 1 & q & 1 & q \\ \vdots & \vdots & \vdots & \vdots \\ 1 & T & 1 & T \end{bmatrix}.$$

We iteratively estimate  $\theta_j$  and  $\Sigma_j$  for all  $j$ , as in Algorithm 1.

Define

$$L(q) = \max_{(\theta_1, \Sigma_1, \dots, \theta_J, \Sigma_J)} L(\theta_1, \Sigma_1, \dots, \theta_J, \Sigma_J | q, Y_1, \dots, Y_J),$$

then the estimated change point is

$$\hat{\tau} = \arg \max_{q \in Q} L(q).$$

---

**Algorithm 1** Estimating  $\theta_j$  and  $\Sigma_j$  iteratively

---

- 1: **for**  $j \in \{1, \dots, J\}$  **do**
  - 2:   set  $\zeta = 1$
  - 3:   set  $\hat{\theta}_j^0$  to OLS estimates
  - 4:   from residuals  $R_j^0$  calculate  $\hat{\phi}_{j1}^0$ ,  $\hat{\phi}_{j2}^0$ ,  $\hat{\sigma}_{j1}^0$ , and  $\hat{\sigma}_{jw,2}^0$  and generate  $\hat{\Sigma}_j^0$
  - 5:   **while**  $\zeta > \text{tol}$  **do**
  - 6:     set  $i$  to the iteration
  - 7:     calculate  $\hat{\theta}_j^i$  based on  $\hat{\Sigma}_j^{i-1}$
  - 8:     use residuals  $R_j^i$  to estimate  $\hat{\phi}_{j1}^i$ ,  $\hat{\phi}_{j2}^i$ ,  $\hat{\sigma}_{jw,1}^i$ , and  $\hat{\sigma}_{jw,2}^i$
  - 9:     obtain  $\hat{\Sigma}_j^i$ ,
  - 10:     set  $\zeta$  to the Euclidean distance between  $[\hat{\phi}_{j1}^{i-1}, \hat{\phi}_{j2}^{i-1}]$  and  $[\hat{\phi}_{j1}^i, \hat{\phi}_{j2}^i]$
  - 11:   **end while**
  - 12: **end for**
-

The estimates of  $\theta_1, \dots, \theta_J$  are the GLS estimates obtained, after the desired tolerance level is reached, conditional on  $\hat{\tau}$ . The GLS estimates of  $\theta_j$  and  $\Sigma_j$  for all  $j$  given  $\hat{\tau}$  are

$$[\{\hat{\theta}_1, \hat{\Sigma}_1\}, \dots, \{\hat{\theta}_J, \hat{\Sigma}_J\}] = \underset{\{[\theta_1, \Sigma_1], \dots, [\theta_J, \Sigma_J]\}}{\arg \max} L(\theta_1, \Sigma_1, \dots, \theta_J, \Sigma_J | \hat{\tau}, Y_1, \dots, Y_J).$$

The presence of  $\tau$  does not restrict the model to a fixed interruption with an instantaneous effect. In fact,  $\tau$  allows the design matrix and estimates to transform based on the information the data provides. Importantly, the inclusion of an over-all-unit change point allows us to utilize information from all available units to determine when the intervention begins to affect the outcome globally. This flexibility of the model can be helpful in minimizing misleading results from an assumed change point.

**2.1.4 | Multivariate Wald test for the existence of a change point**—The change point is estimated by maximizing the conditional likelihood over the set  $Q$  and, thus, concurrently estimates all other model parameters at each possible change point. Since we test for the existence of a change point at each  $q \in Q$ , multiple testing bias exists if one utilizes standard critical values. As such, it is necessary to apply a correction to control the familywise type one error rate. To this end, we calculate the multivariate Wald test statistic for every  $q \in Q$ . We apply the Benjamini-Hochberg method to adjust for the total number of tests conducted—the total number of tests is equal to the cardinality of  $Q$ . The Benjamini-Hochberg method controls the false discovery rate; control of the false discovery rate weakly implies control of the familywise type one error rate for an  $\alpha = 0.05$  level.<sup>18</sup> In this case, a binary decision of whether a change point exists or not corresponds to a rejection of the null hypothesis for any one of the tests conducted.

We focus on determining the existence of a change point across the unit-specific mean functions, ie, for each  $q \in Q$ , we test

$$H_0: \delta_j = \Delta_j = 0 \quad \forall j, j = 1, \dots, J \text{ (no change point)} \\ \text{vs } H_a: \delta_j \neq 0 \text{ and/or } \Delta_j \neq 0, \text{ for some } j, j = 1, \dots, J, \quad \text{(a change point at } q\text{)}.$$

Even though our model assumes a global change point to pool information across units for efficiency, a rejection of the null hypothesis for our Wald test implies a change point in at least one of the hospital units. A rejection does not imply that a change point exists across all units and is the same in all hospital units. Moreover, we do not restrict the impact of the change point at each unit—ie, we allow the change in level and change in slope to differ across units as in R-MITS. We borrow information across units for the estimation of the global change point, but we do not force the impact on the outcome to be the same in each unit. Clearly, if one wanted to establish the existence of a change point for a particular unit, enough data would have to be gathered within that single unit to detect and estimate (with high enough precision) the unit-specific change point.

The hypotheses can be written in terms of full and reduced mean function models. Define the full and reduced mean function models as

$$\mu_{jt}^1 = \beta_{j,0} + \beta_{j,1}t + (\delta_j + \Delta_j t)I(t \geq q), \tag{5}$$

$$\mu_{jt}^0 = \beta_{j,0} + \beta_{j,1}t, \tag{6}$$

respectively. The full model, appropriate under the alternative hypothesis, is essentially the model of Equation (2) and models a change in the mean functions at  $q$ . The reduced model, appropriate under the null hypothesis, assumes the same mean functions across the entire observational period. Based on our model specification, testing the above hypotheses is equivalent to testing

$$H_0: \begin{bmatrix} \delta_1 \\ \Delta_1 \\ \vdots \\ \delta_J \\ \Delta_J \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix} \text{ vs } H_a: \begin{bmatrix} \delta_1 \\ \Delta_1 \\ \vdots \\ \delta_J \\ \Delta_J \end{bmatrix} \neq \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}. \text{ Let } \mathbf{C} = \begin{bmatrix} 0 & 0 & 1 & 0 & & \\ & 0 & 0 & 0 & 1 & & \mathbf{0} \\ & & & & & \ddots & \\ & & & & & & 0 & 0 & 1 & 0 \\ & & & & & & & \mathbf{0} & & 0 & 0 & 0 & 1 \end{bmatrix}.$$

$\vec{\beta}^1 = [\beta_{1,0}^1 \beta_{1,1}^1 \delta_1 \Delta_1 \dots \beta_{J,0}^1 \beta_{J,1}^1 \delta_J \Delta_J]'$  and  $\vec{\beta}^0 = [\beta_{1,0}^0 \beta_{1,1}^0 \dots \beta_{J,0}^0 \beta_{J,1}^0]'$ . Then, the hypotheses can be written as

$$H_0: \mathbf{C} \vec{\beta}^1 = \vec{0} \text{ vs } H_a: \mathbf{C} \vec{\beta}^1 \neq \vec{0},$$

The multivariate Wald test statistic is given by

$$W = \left( \mathbf{C} \widehat{\vec{\beta}}^1 \right) \left[ \mathbf{C} \widehat{\mathbf{V}}(\widehat{\vec{\beta}}^0) \mathbf{C}' \right]^{-1} \left( \mathbf{C} \widehat{\vec{\beta}}^1 \right)' \stackrel{H_0}{\sim} \chi_{2J}^2, \tag{7}$$

where  $\widehat{\mathbf{V}}(\widehat{\vec{\beta}}^0)$  is the block diagonal estimator of the variance covariance matrix of  $\widehat{\vec{\beta}}^0$ . We specify  $\widehat{\mathbf{V}}(\widehat{\vec{\beta}}^0)$  in the Appendix. Note that we allow  $\delta_j$  and  $\Delta_j$  to differ for each  $j$ , ie, for each unit.

We calculate the multivariate Wald statistic for each  $q \in Q$ . Then, we apply the Benjamini-Hochberg procedure to obtain corrected critical values. The Benjamini-Hochberg procedure is fully described by Benjamini and Hochberg.<sup>18</sup> If any of the multivariate Wald tests provide significant results, when compared to the corrected critical values, we conclude that a change point exists for at least one of the units. The resulting ‘supremum Wald test’ (SWT) is appropriate for detecting a change in any of the mean functions over a set of possible change points. Our test accounts for the heterogeneity of the mean functions and autocorrelation structures across units. In the following sections, we illustrate that the SWT has empirically high power under specified change-point alternatives.

### 3 | EMPIRICAL TYPE ONE ERROR AND POWER SIMULATIONS

Prior to analyzing the outcome of interest, we conduct simulations to (1) examine the type one error rate and (2) determine the power and accuracy of our proposed methodology to detect a global change point in the mean functions of the response. These simulations examine the operating characteristics of our proposed SWT under various conditions. We continue to test

$$H_0: \delta_j = \Delta_j = 0 \quad \forall j, j = 1, \dots, J \text{ (no change point)}$$

$$\text{vs } H_a: \delta_j \neq 0 \text{ and/or } \Delta_j \neq 0, \text{ for some } j, j = 1, \dots, J, \quad (\text{a change point at } q)$$

with  $q \in Q$  (the set of possible change points specified by the researcher). The full and reduced models of these simulations are those of Equations (5) and (6), respectively. We have additionally examined the scenario with standardized quadratic time (and standardized linear time) in the mean functions of the reduced and full models. We focused on standardized time, as opposed to untransformed time, to avoid collinearity between the two time terms. These simulations are omitted for brevity, though we note that we obtain similar results as those discussed in the following sections. In both sets of simulations, we assume an autocorrelation structure that remains constant over the entire duration of the observational period, since the focus is on testing for the existence of a change point in the mean functions.

The outcome of interest is recorded for 60 time periods, in five units, with adjacent correlation estimates smaller than  $\phi = 0.1$ , and the set of possible change points equal to  $\{25, \dots, 34\}$ . Thus, we chose parameters similar to these values for our simulations. We consider two values of the time series length,  $T \in \{60, 120\}$ , two values of the adjacent correlation,  $\phi \in \{0.1, 0.6\}$ , and three values for the total number of units,  $J \in \{1, 3, 5\}$ . When the length of the time series is  $T = 60$ , we allow the set of possible change points to be  $Q_{60} = \{25, \dots, 34\}$ , as with the patient satisfaction data. In this situation, we conduct 10 total tests, since there are 10 elements in  $Q_{60}$ . When the length of the time series is  $T = 120$  we allow the set of possible change points to be  $Q_{120} = \{50, \dots, 69\}$ ; a total of 20 tests are conducted for  $Q_{120}$ .

We choose to compare two values of the time series length to illustrate the possible gain in efficiency longer time series provide with regard to power. We illustrate the gain that may come from doubling the length of the time series. The length of the time series can be increased in two ways: (1) increase the observational period, say, from 5 years to 10 years and/or (2) increase the resolution of recordings, ie, record patient satisfaction bimonthly, as opposed to monthly. The two values of  $\phi$  examined are larger adjacent correlation values than what we estimate for the patient satisfaction data. The largest unit-specific adjacent correlation estimates obtained for the patient satisfaction data (when information is not borrowed across units) is 0.09, and so, 0.1 is an upper bound for the adjacent correlation in our setting. The value  $\phi = 0.6$  represents an upper bound for the correlation between repeated measurements in the literature. The estimated adjacent correlations for patient satisfaction are smaller than either 0.1 and 0.6. Our simulation results are conservative because power decreases for ITS designs as the adjacent correlation increases.<sup>19</sup>

Importantly, we conduct type one error and power simulations for the SWT with one, three, and five units. We examine the case with a single unit,  $J=1$ , to illustrate the performance of our SWT in the traditional ITS analysis setting. We explore the value  $J=3$  to depict the healthy gain in efficiency that borrowing information across a small number of units yields. Lastly, we consider  $J=5$  because patient satisfaction is recorded at five units. Our aim is to highlight the improvement in power that borrowing information across units can provide.

### 3.1 | Empirical type one error for SWT

We provide the empirical type one error rates when testing for the existence of a change point via the SWT. Four different scenarios are considered for one unit, three units, and five units. We generate 10 000 time series for each scenario under the reduced model—ie, from one overall regime where there is no change point present in either the mean functions or stochastic processes.

For the case when there is only one unit, we set the mean function parameters to  $\vec{\beta} = [65, 0.5]'$ . When there are three units or five units, the mean function parameters vary slightly across individual units. The white noise standard deviation,  $\sigma_w$ , is always set to 3.38, regardless of the number of units in the simulation. The value  $\sigma_w = 3.38$  is approximately the average of the single-unit estimates of the white noise standard deviation for patient satisfaction. The response standard deviation,  $\sigma$ , is 4.23 when  $\phi = 0.6$  and 3.40 when  $\phi = 0.1$  for all individual units. The mean function parameter values mimic results obtained from the patient satisfaction data.

The empirical type one error rates for each scenario are provided in Table 1. As expected, the empirical type one error rate is smaller for the longer time series and for smaller values of the adjacent correlation. Larger values of adjacent correlation imply a smaller number of effective independent statistical information. The larger adjacent correlation quantity corresponds to higher type one error rates exclusively. In the scenario with the shorter time series and high adjacent correlation value, it is more difficult to control the familywise type one error rate, even as we increase the number of units. The lack of type one error rate control in short time series with high correlation values is exacerbated in the simulations with quadratic time in the mean functions. In fact, for that particular setting, the type one error rate becomes worse as the number of units increases. This is primarily attributable to the increased dependency in an already short time series that reduces the information in the time series. Because of this, we recommend our proposed procedure when the length of the time series is at least 120 time points in cases with complex mean functions and/or high correlation values. For all other scenarios considered, the type one error rate is well controlled, though slightly conservative because of the Benjamini-Hochberg multiplicity correction.

### 3.2 | Empirical power for SWT

We conduct simulation-based power calculations when testing for the existence of a change point via the SWT. Time series are generated under the alternative model appropriate in our setting, ie, generated with a global change point in the mean functions. The change point is set at the middle of the time series; cases with the change point at the boundary or close to

the boundary have been considered and yield similar, yet slightly less powerful, results. We focus on providing power as a function of the slope change. Simulation-based power calculations with power as a function of the autoregressive coefficient for ITS designs are provided in the work of Zhang et al.<sup>19</sup>

Power is examined as a function of the slope change, with a change in baseline intercept ( $\delta_j$ ) set to zero. Note that estimates of  $\delta_j$  obtained from the patient satisfaction data are not statistically different from zero. The range of values for the change in slope,  $\{0, 0.01, \dots, 0.24, 0.25, 0.30, \dots, 0.40, 0.45\}$ , encompass the estimated quantities of the change in slope for the patient satisfaction data. Similarly to the type one error simulations, the white noise standard deviation is set to 3.38 for all units, yielding a response standard deviation of 4.23 when  $\phi = 0.6$  and 3.40 when  $\phi = 0.1$ .

Simulated power curves are provided in Figure 3, with each subfigure corresponding to a separate data generation regime. As expected, power increases as the slope change and length of the time series increases, and power decreases for the larger adjacent correlation value. Power is consistently higher for the larger number of units across the four scenarios, thereby illustrating that the SWT gains power as the number of units increases by borrowing information across units. Analyzing multiple time series data (or data from multiple hospital units) jointly results in higher power.

**3.2.1 | Accurate estimation of the change point**—The power simulation results, provided in Figure 3, suggest that the SWT has reasonable power to detect an existing global change point and that power increases as the number of units increase. We are not simply interested in power by itself. We are also interested in whether R-MITS will provide the correct global change point estimate when our SWT concludes that a change point exists. Figure 4 illustrates the proportion of simulations that correctly estimate the true change point as a function of the slope change for one, three, and five units. Similar to the empirical power, the proportion of correctly estimated change points increases as the number of units and the length of the time series increases. We also calculated the proportion of simulations that exactly estimate the true change point for change points not in the middle of the time series—ie, with a change point on the boundary or near the boundary—and obtained comparable results.

## 4 | MULTIUNIT ANALYSIS OF THE CNL INTERVENTION

We assess the impact of the CNL integrated care delivery intervention on average patient satisfaction at five hospital units. Average patient satisfaction is the mean of patient satisfaction survey scores for seven indicators, shown for the Stroke and Surgical units in Figure 1. The seven patient satisfaction indicators are as follows: effective nurse communication, nurses treated me with courtesy/respect, responsiveness of hospital staff, effective physician communication, staff did everything to help control your pain, effective communication about medicines, and discharge information provided. We refer to the average patient satisfaction scores simply as patient satisfaction.

We are interested in estimating the time lag (or delay) between the onset of the intervention and the effect on patient satisfaction. In practice, the change point may occur either *before* or *after* the official intervention time. An intervention intended to improve care quality requiring a training over several months or weeks (such as the CNL intervention) may already produce a change in the outcome, even before the official start of intervention, if the trainees execute their training as they learn.

Inference on the global change point and time lag between the onset of the intervention and the intervention's effect is included in Table 2. In addition, Table 2 provides the (a) global change-point estimate, (b) SWT p-value, (c) time point of CNL integration into their respective hospital units, (d) formal intervention implementation time, and (e) lag between formal intervention implementation and estimated change point. The SWT concludes that a change point exists over the set of possible change points for patient satisfaction in at least one of the units at the  $\alpha = 0.05$  level. The p-value associated with the test for the existence of a change point is 0.003, which is less than the respective Benjamini-Hochberg corrected critical value of 0.005. R-MITS estimates a preemptive effect of the CNL integrated care delivery intervention on patient satisfaction. The global change point is estimated to occur on May 2010, while the formal intervention implementation occurs on July 2010. Estimating an anticipatory effect (from the expected and a priori specified change point) is not feasible with standard segmented time series regression. Segmented regression methodology requires clearly separated pre- and post-intervention phases, often with an assumed change point greater than or equal to the formal intervention implementation time point.

Although the CNL integrated care delivery is officially implemented on July 2010, it was unofficially being practiced prior to July 2010. Nurses put into practice the new concepts they learned from their "training." It is completely realistic that many of the CNLs implemented their training prior to July 2010, particularly, if they believed it would be beneficial. Thus, the anticipatory effect of the CNL integrated care delivery intervention (of 2 months, provided in the "Lag" column of Table 2) is consistent with the integration of the CNLs on January 2010. In fact, the estimated global change point for patient satisfaction occurs 4 months after the CNLs introduction into their respective units. The CNL care delivery intervention requires a restructuring of patient care and care delivery, likely to manifest itself to patients after a time lag from the CNLs introduction. This time lag and the behavioral component of the intervention may explain why the global change point occurs 4 months after the CNLs integration into the hospital units and 2 months prior to the formal intervention time point.

Estimates of the R-MITS mean function parameters are provided in Tables 3 and 4, and estimates of the stochastic process parameters are included in Table 5. Estimates and 95% confidence intervals of the two standardized effect sizes used in the health care literature, change in level and change in trend/slope,<sup>16</sup> are provided in Table 4. The level and trend change are not statistically significant for any unit. The estimated level change tends to be positive for the majority of hospital units, indicating an initial drop of the outcome level, as in Figure 2. This may be due to the adjustment period associated with the intervention. Moreover, it may occur as an artifact of the regression itself, particularly for a bounded outcome such as patient satisfaction.

Trend (slope) change is negative for patient satisfaction, suggesting a decrease in the slope of patient satisfaction post-change point. Due to the nature of patient satisfaction as a percentage—and thus as a bounded outcome—the change in slope must be interpreted with caution. Patient satisfaction cannot continue to grow at a rapid rate because the mean patient satisfaction function at the estimated change point is already relatively close to 100, the maximum patient satisfaction value. This is evident in Figure 5, in which the estimated mean functions for all hospital units are plotted, particularly for the Stroke, Surgical, and Mother/Baby units.

The estimated volatility of patient satisfaction, given by the “Standard Deviation” column of Table 5, is smaller post-change point for 3 out of the 5 units, and the adjacent correlation is larger post-intervention in 4 out of the 5 units. The Medical Surgical and Mother/baby units estimated standard deviations increase post-estimated change point, increasing from 4.84 and 2.97 to 5.06 and 3.72, respectively; while in the Stroke, Surgical, and Cardiac units, the estimated standard deviation decreases from 3.15, 4.37, and 5.28 to 3.01, 2.35, and 3.76, respectively. After the estimated change point, the patient satisfaction scores are observed to be less volatile for the Stroke, Surgical, and Cardiac units and, hence, may be more predictable. The adjacent correlation estimates mainly move from negative to positive post-estimated change point, indicating a more stationary patient satisfaction score post-intervention. These are positive results of the CNL intervention. It is important for hospitals to have patients that are generally satisfied over patients who range from extremely satisfied to extremely dissatisfied. Patient satisfaction scores that are more dependent, closely related, and less volatile result in a more predictable outcome.

#### 4.1 | Doubly robust ITS

R-MITS pools information across units to estimate a global change point, thereby increasing efficiency and reducing the impact of misleading influential points. Reducing the effect of influential points is desirable in our patient satisfaction data, for which the change point search space consists of only a few time points. We illustrate the gravity of influential points on the estimated change point for the single-unit analyses of patient satisfaction at the Medical Surgery and Cardiac units. To model patient satisfaction for a single unit, we implement the Robust-ITS model. The estimated mean functions and change point estimates are included in Figure 6 for two cases. The plots on the left of Figure 6 correspond to single-unit analyses including all observations, while the plots on the right pertain to the single-unit analyses without observation  $t = 25$  (January 2010). When all the observations are included, Robust-ITS estimates the change point to be February 2010 for both the Medical Surgical and Cardiac units. However, for the analyses without January 2010, the estimated change points are October 2010 and April 2010 for the Medical Surgical and Cardiac units, respectively. One single time point has the ability to perturb the estimated change point by 6 months in the Medical Surgical unit and by 2 months in the Cardiac unit. Our proposed R-MITS model guards against these influential points by borrowing information across hospital units. Pooling data across hospital units in the estimation of a global change point automatically reduces the impact of spurious influential points, resulting in robust mean function estimates.



## 5 | CONCLUSION AND FUTURE WORK

Our proposed R-MITS model is appropriate for multiple time series, able to estimate a global change point rather than assume it a priori, and can model differences in both the mean functions and stochastic components. R-MITS borrows information across units to estimate a global change point and to estimate the mean functions and stochastic processes separately for each unit. The proposed model does not assume that the impact of the global change point on the outcome is equivalent for all units. That is, although R-MITS borrows information across units to estimate an over-all-unit change point, the level change and trend change are allowed to vary for each unit. R-MITS further allows the autocorrelation and variability during pre- and during post-intervention to differ across units.

Importantly, our proposed SWT formally tests for the *existence* of a change point in at least one of the mean functions, rather than merely assuming and requiring a change. Now, researchers will be able to formally test whether an intervention is associated with a change in the mean functions of a health outcome. Erroneous inference regarding the response's mean functions may result from incorrectly assuming—both the existence and placement—of the change point. Assuming a change point when no change point truly exists forces the estimation of an artificial change. Our SWT will test for the existence of a change in the response over a pre-determined set of possible change points. As demonstrated by our simulation studies, the operating characteristics of R-MITS and the SWT are well behaved with regard to power and type one error. Moreover, the empirical power of the SWT and accuracy of the change point estimates—and so the accuracy of the estimated time delay between an intervention and the intervention's effect on an outcome—increase as the number of units increases.

The R-MITS model and the SWT provide researchers with insight to re-address hypothesis generation for future study design. The methodology better informs researchers of the likely lag that may be realistic for a similar intervention. We note that, in our application example, nurses finishing their masters' thesis project (in a program that trained them to implement the CNL intervention) were introduced into their respective hospital units 6 months prior to the formal intervention. The nurses' integration potentially changed practice as soon as they were introduced. In fact, the estimated change point occurs between the introduction of the nurses to the hospital unit and the formal intervention. A primary utility of R-MITS is that, through exploration of the change point, we are able to observe this and provide direction for future study planning.

R-MITS estimates the global change point via a grid search over a pre-determined set of possible change points. Researchers must specify the set of possible change points with care since, as with traditional ITS designs, we must be cautious of competing intervention effects. The set of possible change points must adequately capture the time points during which the intervention of interest plausibly impacted the outcome, yet simultaneously exclude time periods affected by another intervention. This is to avoid the risk of competing interventions. Parsing out the effect of competing interventions is a concern in general with ITS designs. Ideally, the entire observational period (both the pre- and post-intervention phases) of an ITS

design should be solely affected by the intervention of interest. Although theoretically simple, in practice, this requires careful consideration and expertise.

Identification of a change point via our proposed procedure relies upon detection of a difference in either the mean level of the response and/or the slope of the response, comparing the pre- and post-intervention effect periods across units. As such, if no change point in the time series truly exists, this would indicate that there is no difference in the mean function of the response over time. Most researchers would consider this absence of a difference in the mean function to be the absence of an intervention effect. One could argue that, if the pre-intervention slope was positive (indicating improvement in outcomes) and if the slope remained constant during the post intervention, then this could have been solely attributable to the intervention. In this case, the counterfactual may have revealed a decline (or an increase) in the slope if the intervention had not been instituted. Of course, such a counterfactual could never be observed in practice but certainly should be considered in theory.

Currently, the SWT focuses on changes solely in the mean functions. We are currently working on extensions of the SWT to accurately detect changes that are both the mean functions and stochastic components, to better handle the nuances of the autocorrelation structures across units. In addition, we are considering more efficient multiple testing corrections that utilize information obtained in the autocorrelation structure. We plan on developing the theory needed for our SWT (and its future derivatives) to guarantee consistency.

It is paramount to note that the current status of the R-MITS model is for continuous-valued outcomes only. We will soon expand this class of models to handle counts and rates data (eg, infection rates and counts of accidental falls). Lastly, R-MITS does not provide inference on the overall population of hospital units—the population of hospital units that gives rise to the units we observe. Particularly, R-MITS does not account for heterogeneity of change points across units for situations where the data warrants such treatment. We will develop an ITS mixed effects model as an alternative to R-MITS, able to detect unit-specific change points and borrow information across units while allowing for change point heterogeneity.

## ACKNOWLEDGEMENTS

This work was supported in part by the Eugene Cota-Robles Fellowship at the University of California, Irvine, by the National Science Foundation Graduate Research Fellowship Program under grant DGE-1321846, by the National Science Foundation under grants MMS 1461534 (National Science Foundation Division of Mathematical Sciences) and 1509023 (National Institute on Aging of the National Institutes of Health), R01AG053555 and P50AG16573, and by the National Institute of Mental Health MH115697. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors(s) and do not necessarily reflect the views of the National Science Foundation. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

### Funding information

Eugene Cota-Robles Fellowship at the University of California, Irvine; National Science Foundation Graduate Research Fellowship Program, Grant/Award Number: DGE 1321846; National Science Foundation, Grant/Award Number: MMS 146153; National Science Foundation Division of Mathematical Sciences, Grant/Award Number: 1509023; National Institute on Aging of the National Institutes of Health, Grant/Award Number: R01AG053555 and P50AG16573; National Institute of Mental Health, Grant/Award Number: MH115697

## APPENDIX A

### A.1 | Estimators of the mean function parameters

The GLS estimators for the mean function parameters of  $\theta_1, \dots, \theta_j$  given  $q \in Q$ , obtained in step (6) of Algorithm 1 of Section 2.1.3, are

$$\hat{\theta}_j = \begin{bmatrix} \hat{\beta}_{j0} \\ \hat{\beta}_{j1} \\ \hat{\delta}_j \\ \hat{\Delta}_j \end{bmatrix} = \left[ \mathbf{X}^{(q)j} \hat{\Sigma}_j^{-1} \mathbf{X}^{(q)j} \right]^{-1} \left[ \mathbf{X}^{(q)j} \hat{\Sigma}_j^{-1} \mathbf{Y}_j \right],$$

$$\text{where } \mathbf{X}^{(q)j} \equiv \begin{bmatrix} 1 & 2 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & q-1 & 0 & 0 \\ 1 & q & 1 & q \\ \vdots & \vdots & \vdots & \vdots \\ 1 & T & 1 & T \end{bmatrix},$$

and  $\hat{\Sigma}_j$  given in the subsequent section. Hence, for unit  $j$ , the estimator of (a) the intercept pre-change point is  $\hat{\beta}_{j0}$ ; (b) the slope pre-change point is  $\hat{\beta}_{j1}$ ; (c) the change in level (post-change-point intercept anchored at  $\hat{\tau}$ ) is  $\hat{\delta}_j + \hat{\Delta}_j \hat{\tau}$ ; and the change in slope is  $\hat{\Delta}_j$ .

### A.2 | Estimators of the AR(1) processes parameters

In steps (4) and (7) of the iterative estimation process, provided in Section 2.1.3, the residuals,  $r_{jt} = y_{jt} - \hat{\mu}_{jt}$ , are modeled as AR(1) processes

$$r_{jt} = \begin{cases} \phi_{j1} r_{j,t-1} + e_{jt,1}, & 1 < t < q, \\ \phi_{j2} r_{j,t-1} + e_{jt,2}, & q \leq t \leq T, \end{cases}$$

with  $e_{jt,i} \stackrel{\text{iid}}{\sim} N(0, \sigma_{jw,i}^2)$  for  $i \in \{1, 2\}$ . Recall, to ensure causality in the time series sense,  $\phi_{j1}$  and  $\phi_{j2}$  must lie in the interval  $(-1, 1)$  for all  $j$ . The variance-covariance matrix,  $\Sigma_j$ , is therefore equal to

$$\begin{bmatrix} \frac{\sigma_{jw,1}^2}{1-\phi_{j1}^2} \begin{bmatrix} 1 & \phi_{j1} & \dots & \phi_{j1}^{q-2} \\ \phi_{j1} & 1 & \dots & \phi_{j1}^{q-3} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{j1}^{q-2} & \phi_{j1}^{q-2} & \dots & 1 \end{bmatrix} & \mathbf{0} \\ \mathbf{0} & \frac{\sigma_{jw,2}^2}{1-\phi_{j2}^2} \begin{bmatrix} 1 & \phi_{j2} & \dots & \phi_{j2}^{T-q} \\ \phi_{j2} & 1 & \dots & \phi_{j2}^{T-q-1} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{j2}^{T-q} & \phi_{j2}^{T-q-1} & \dots & 1 \end{bmatrix} \end{bmatrix}$$

and completely determined by  $\phi_{j1}$ ,  $\phi_{j2}$ ,  $\sigma_{jw,1}$ , and  $\sigma_{jw,2}$ . We therefore only provide the estimators of  $\phi_{j1}$ ,  $\phi_{j2}$ ,  $\sigma_{jw,1}$ , and  $\sigma_{jw,2}$ , conditional on  $q \in Q$ . Define

$$\begin{aligned} \bar{r}_{(1a)} &= \frac{1}{q-2} \sum_{t=1}^{q-2} r_t, & \bar{r}_{(1b)} &= \frac{2}{q-1} \sum_{t=2}^{q-1} r_t, \\ \bar{r}_{(2a)} &= \frac{1}{T-q-1} \sum_{t=q-1}^{T-1} r_t, & \bar{r}_{(2b)} &= \frac{1}{T-q-1} \sum_{t=q}^T r_t, \\ \sigma_{r_1}^2 &= \frac{\sum_{t=2}^{q-1} (r_t - \bar{r}_{(1b)})^2 + \sum_{t=2}^{q-1} (r_{t-1} - \bar{r}_{(1a)})^2}{2}, \\ \text{and } \sigma_{r_2}^2 &= \frac{\sum_{t=q}^T (r_t - \bar{r}_{(2b)})^2 + \sum_{t=q}^T (r_{t-1} - \bar{r}_{(2a)})^2}{2}. \end{aligned}$$

**The estimators of  $\phi_{j1}$ ,  $\phi_{j2}$ ,  $\sigma_{jw,1}$ ,  $\sigma_{jw,2}$ ,  $\sigma_{j,1}$ , and  $\sigma_{j,2}$ , conditional on  $q$ , are**

•

$$\hat{\phi}_{j1} = \frac{\sum_{t=2}^{q-1} (r_t - \bar{r}_{(1b)})(r_{t-1} - \bar{r}_{(1a)})}{\sigma_{r_1}^2}$$

•

$$\hat{\phi}_{j2} = \frac{\sum_{t=q}^T (r_t - \bar{r}_{(2b)})(r_{t-1} - \bar{r}_{(2a)})}{\sigma_{r_2}^2}$$

•

$$\hat{\sigma}_{jw,1}^2 = \frac{1}{q-2} \sum_{t=2}^{q-1} [(r_t - \bar{r}_{(1b)}) - \hat{\phi}_{j1}(r_{t-1} - \bar{r}_{(1a)})]^2$$

•

$$\hat{\sigma}_{jw,2}^2 = \frac{1}{T-q+1} \sum_{t=q}^T [(r_t - \bar{r}_{(2b)}) - \hat{\phi}_{j2}(r_{t-1} - \bar{r}_{(2a)})]^2$$

$$\hat{\sigma}_{j,1} = \frac{\hat{\sigma}_{jw,1}}{\sqrt{1 - (\hat{\phi}_j)^2}}$$

$$\hat{\sigma}_{j,2} = \frac{\hat{\sigma}_{jw,2}}{\sqrt{1 - (\hat{\phi}_j)^2}}.$$

### A.3 | Covariance matrix of the full model mean function parameters for the SWT

The supremum Wald statistic of Section 2.1.4 depends on  $\widehat{\mathbf{V}}_0(\widehat{\beta}^0)$ , the block diagonal estimator of the variance covariance matrix of  $\widehat{\beta}^0$ . Each block of  $\widehat{\mathbf{V}}_0(\widehat{\beta}^0)$  corresponds to  $\widehat{\mathbf{V}}_0(\widehat{\beta}^0_j)$ , the estimated variance-covariance matrix of the mean function parameters for unit  $j$ . Note that

$$\widehat{\mathbf{V}}_0(\widehat{\beta}^0_j) = (\mathbf{X}^1{}' (\widehat{\Sigma}_j)^{-1} \mathbf{X}^1)^{-1},$$

with  $\mathbf{X}^1$  as the design matrix of the **full** model (model of Equation (6)) and the variance-covariance matrix under the **reduced** model (model of Equation (5)) as  $\widehat{\Sigma}_j$ . Since the aim of the SWT is to test the existence of a change point in the mean, we assume an autocorrelation structure that remains constant over the entire duration of the observational period. Thus, for unit  $j$ ,

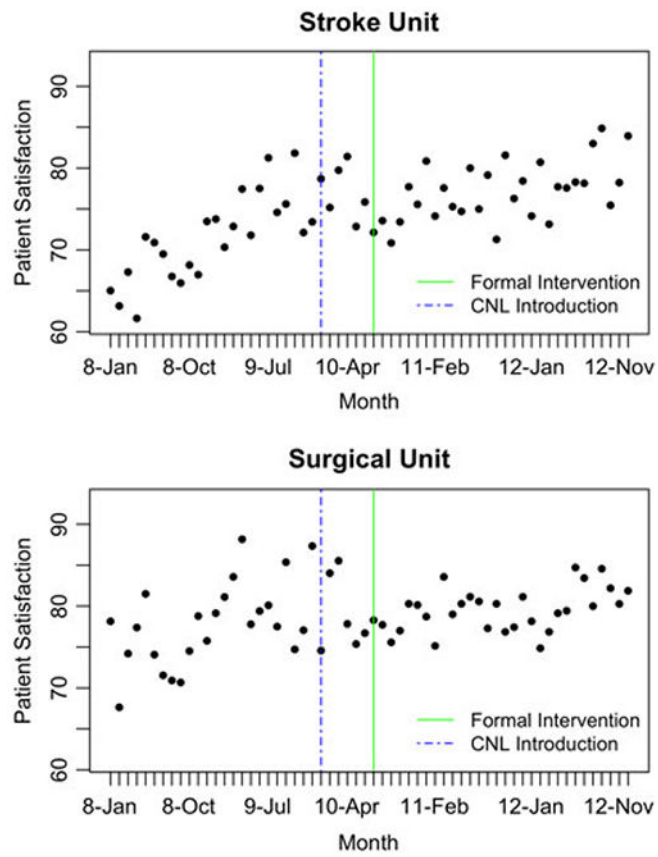
$$\widehat{\Sigma}_j = \frac{(\hat{\sigma}_{jw})^2}{1 - (\hat{\phi}_j)^2} \begin{bmatrix} 1 & \hat{\phi}_j & \dots & (\hat{\phi}_j)^{T-2} \\ \hat{\phi}_j & 1 & \dots & (\hat{\phi}_j)^{T-3} \\ \vdots & \vdots & \ddots & \vdots \\ (\hat{\phi}_j)^{T-2} & (\hat{\phi}_j)^{T-3} & \dots & 1 \end{bmatrix},$$

where  $\hat{\phi}_j$  and  $(\hat{\sigma}_{jw})^2$  are estimated under the reduced model.

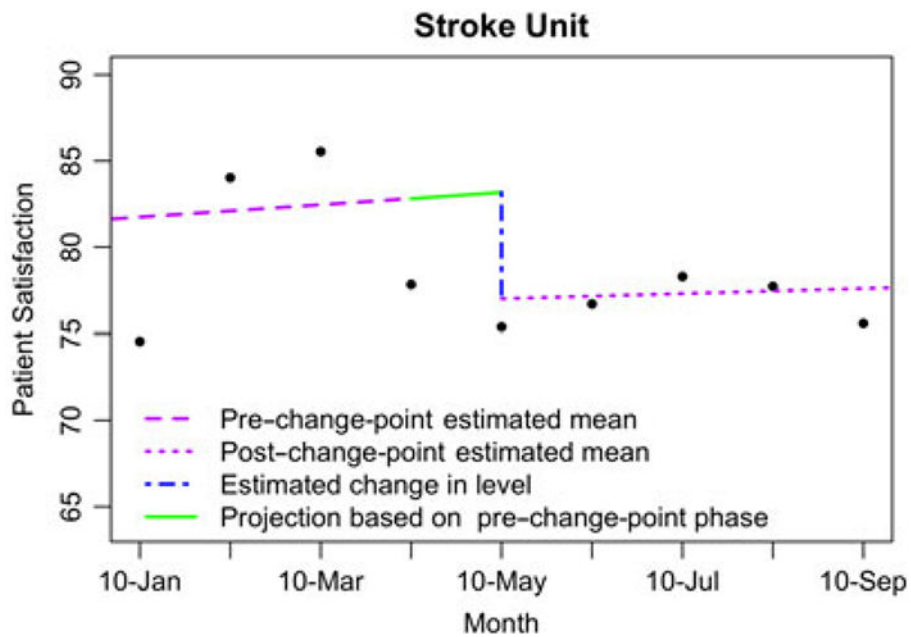
## REFERENCES

1. Hawe P Lessons from complex interventions to improve health. *Annu Rev Public Health*. 2015;36:307–323. [PubMed: 25581153]
2. Datta J, Petticrew M. Challenges to evaluating complex interventions: a content analysis of published papers. *BMC Public Health*. 2013;13(1):1. [PubMed: 23280303]
3. Campbell M, Fitzpatrick R, Haines A, et al. Framework for design and evaluation of complex interventions to improve health. *BMJ*. 2000;321(7262):694–696. [PubMed: 10987780]
4. Shadish WR, Cook TD, Campbell DT. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston, MA: Houghton Mifflin and Company; 2002.

5. Bender M, Murphy E, Thomas T, Kaminski J, Smith B. Clinical nurse leader integration into care delivery microsystems: quality and safety outcomes at the unit and organization level. Paper presented at: Academy Health Annual Research Meeting; 2015; Minneapolis, MN.
6. Kavanagh KT, Cimiotti JP, Abusalem S, Coty M. Moving healthcare quality forward with nursing-sensitive value-based purchasing. *J Nurs Scholarsh.* 2012;44(4):385–395. [PubMed: 23066956]
7. Bender M, Williams M, Su W, Hites L. Refining and validating a conceptual model of clinical nurse leader integrated care delivery. *J Adv Nurs.* 2017;73(2):448–464. [PubMed: 27555500]
8. Taljaard M, McKenzie JE, Ramsay CR, Grimshaw JM. The use of segmented regression in analysing interrupted time series studies: an example in pre-hospital ambulance care. *Implement Sci.* 2014;9(1):1. [PubMed: 24398253]
9. Penfold RB, Zhang F. Use of interrupted time series analysis in evaluating health care quality improvements. *Acad Pediatr Assoc.* 2013;13(6):S38–S44.
10. Wagner AK, Soumerai SB, Zhang F, Ross-Degnan D. Segmented regression analysis of interrupted time series studies in medication use research. *J Clin Pharm Ther.* 2002;27(4):299–309. [PubMed: 12174032]
11. Linden A Conducting interrupted time-series analysis for single-and multiple-group comparisons. *Stata J.* 2015;15(2):480–500.
12. Cruz M, Bender M, Ombao H. A robust interrupted time series model for analyzing complex health care intervention data. *Statist Med.* 2017;36(29):4660–4676.
13. Bhaskaran K, Gasparini A, Hajat S, Smeeth L, Armstrong B. Time series regression studies in environmental epidemiology. *Int J Epidemiol.* 2013;42(4):1187–1195. [PubMed: 23760528]
14. Davis RA, Lee TCM, Rodriguez-Yam GA. Structural break estimation for nonstationary time series models. *J Am Stat Assoc.* 2006;101(473):223–239.
15. Kirch C, Muhsal B, Ombao H. Detection of changes in multivariate time series with application to EEG data. *J Am Stat Assoc.* 2015;110(511):1197–1216.
16. Effective Practice and Organisation of Care (EPOC). Interrupted time series (ITS) analyses. EPOC Resources for review authors. Oslo, Norway: Norwegian Knowledge Centre for the Health Services; 2015.
17. Shumway RH, Stoffer DS. *Time Series Analysis and Its Applications: With R Examples.* New York, NY: Springer Science + Business Media; 2011.
18. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J Royal Stat Soc Ser B Methodol.* 1995;57(1):289–300.
19. Zhang F, Wagner AK, Ross-Degnan D. Simulation-based power calculation for designing interrupted time series analyses of health policy interventions. *J Clin Epidemiol.* 2011;64(11):1252–1261. [PubMed: 21640554]



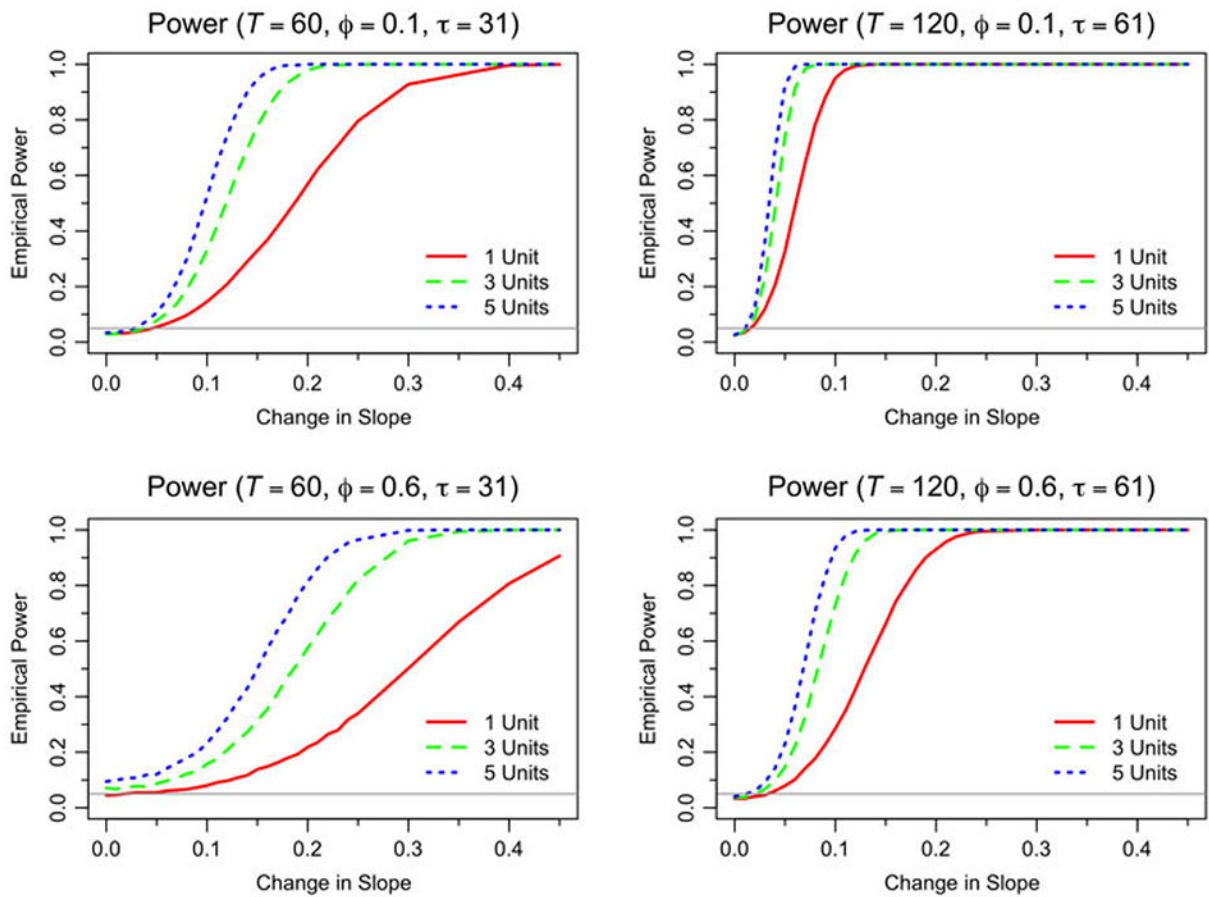
**FIGURE 1.** Plots the time series of observed average patient satisfaction for the Stroke and Surgical units. CNL, Clinical Nurse Leader



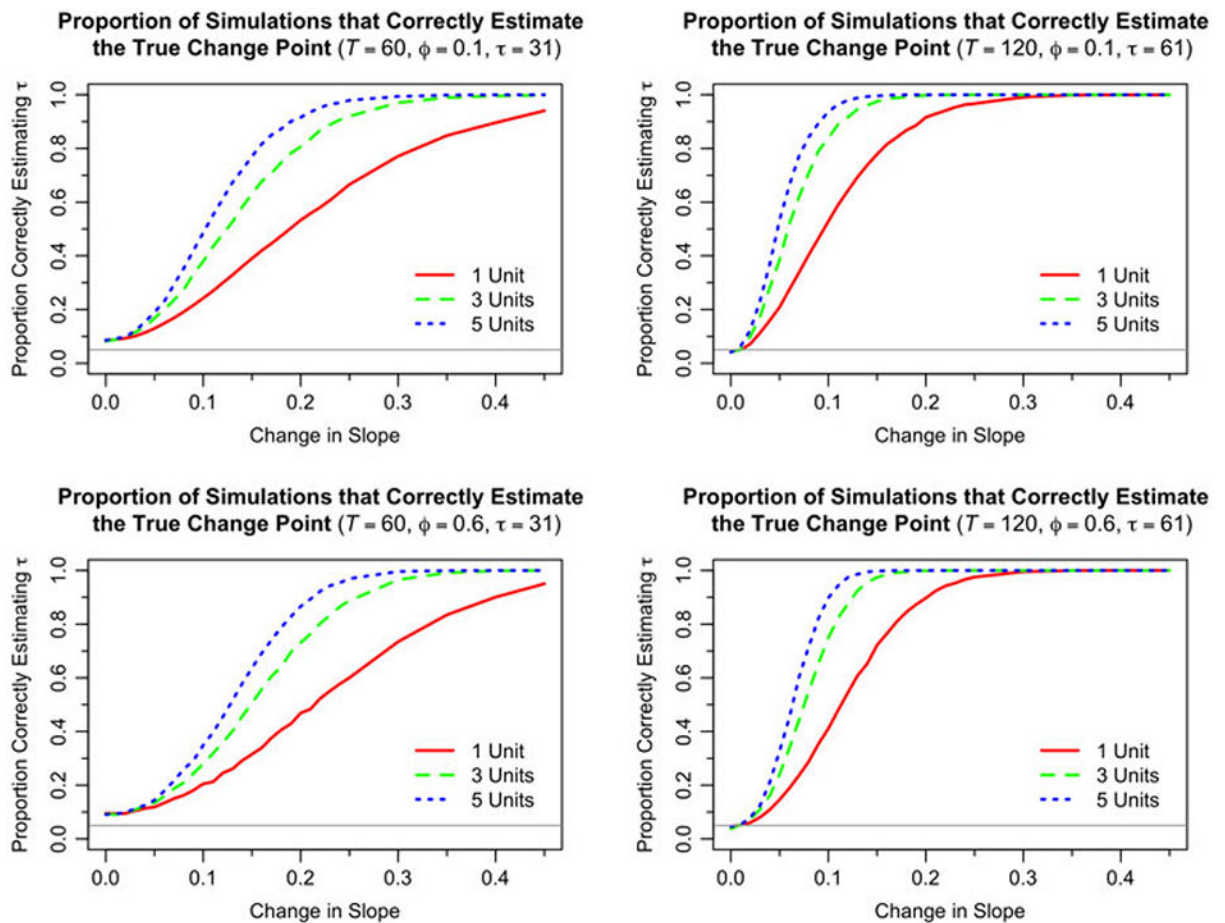
**FIGURE 2.**

An example of a segmented regression model fit for the Stroke unit. The plot depicts (1) the segmented regression lines fit to the pre- and post-change-point phases, (2) the projection of the mean at the change point based on the pre-change-point regression, and (3) the change in level as defined here. The plot contains data from January 2010 to September 2010, instead of the entire observational period, to clearly illustrate the level change

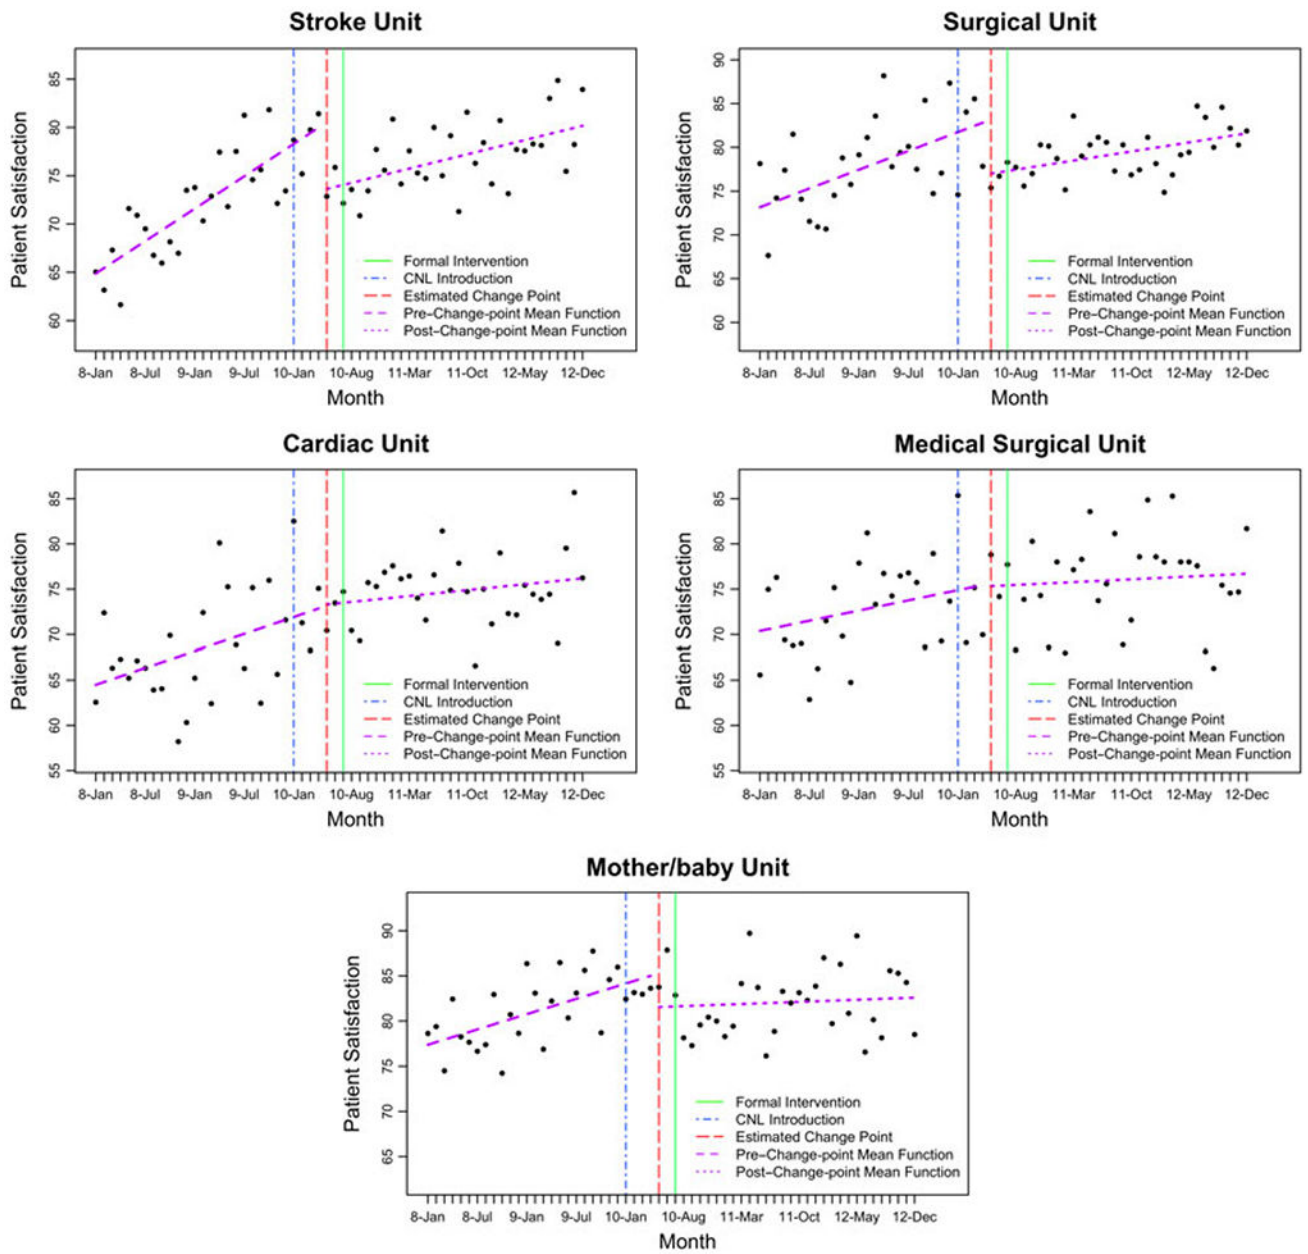


**FIGURE 3.**

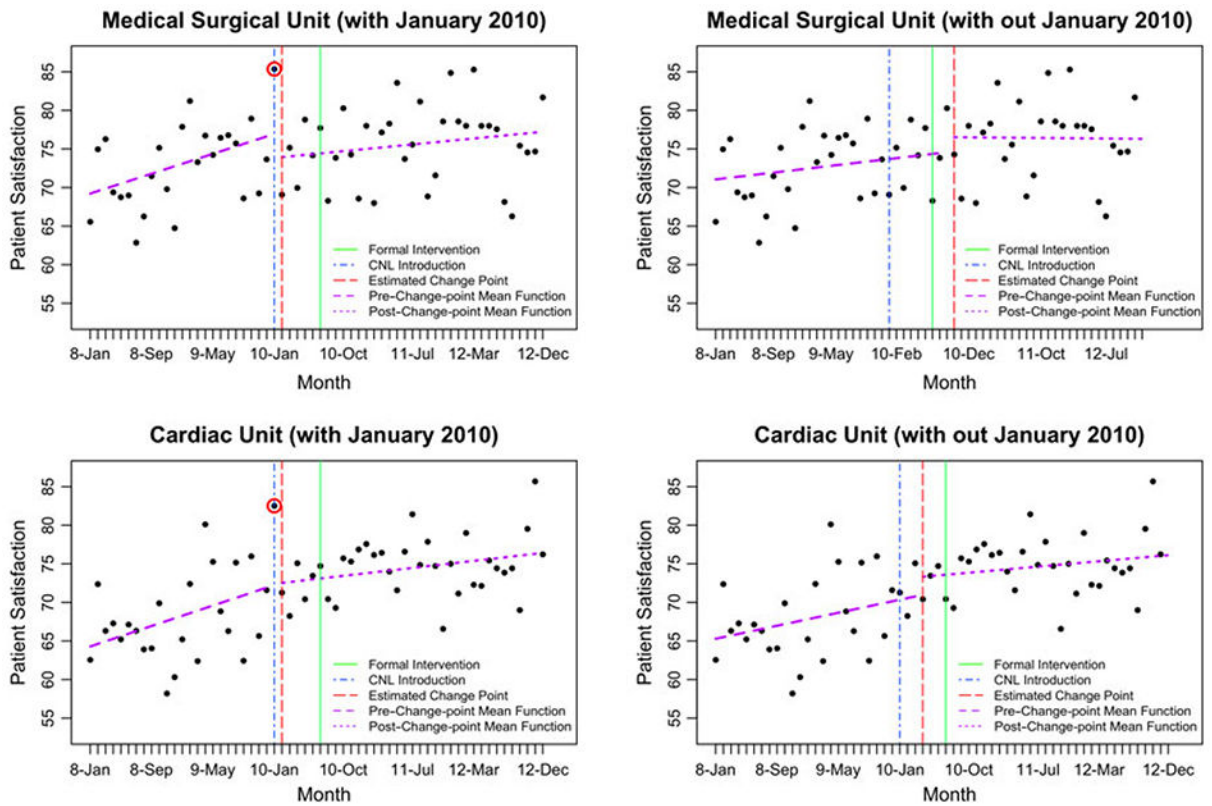
Empirical power, over 10 000 iterations, for various number of units and for 4 regimes. The empirical power increases as the number of units and the length of time series increases, and the power increases as the adjacent correlation decreases

**FIGURE 4.**

The proportion of estimated change points exactly equal to the true change point, over 10 000 iterations, for various number of units and for 4 regimes. Similar to the empirical power, the proportion of correctly estimated change points increases as the number of units and the length of time series increases



**FIGURE 5.** Plots the time series of observed average patient satisfaction for all hospital units, along with the estimated change point, estimated mean functions, and formal intervention time



**FIGURE 6.** Plots the time series of observed average patient satisfaction, along with the estimated change point, estimated mean functions, and formal intervention time for the Medical Surgery and cardiac units with and without observation  $t = 25$  (January 2010), obtained by using Robust-ITS to conduct the unit-specific analyses. *Note that the analysis with  $t = 25$  is on the left and the analysis without  $t = 25$  is on the right.* CNL, Clinical Nurse Leader

**TABLE 1**

The empirical type one error rate: the proportion of iterations for which we rejected the null hypothesis of no change point. The larger adjacent correlation quantity corresponds to higher type one error rates exclusively. The type one error rates are reasonable for almost all of the scenarios and stay reasonable as the number of units increases. However, it is slightly difficult to control the type one error rate at the desired  $\alpha = 0.05$  level with the smaller time series and high adjacent correlation value

Adjacent Correlation	Time Series of Length 60			Time Series of Length 120		
	One Unit	Three Units	Five Units	One Unit	Three Units	Five Units
$\phi = 0.1$	0.0295	0.0291	0.0342	0.0274	0.0265	0.0263
$\phi = 0.6$	0.0460	0.0704	0.1003	0.0299	0.0318	0.0436

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE 2**

Provides the estimated global change point, its p-value, the month Clinical Nurse Leaders were integrated into their respective units, the formal intervention time point, and the intervention lag. The Benjamini-Hochberg corrected p-value cut-off is 0.005. We conclude that there is a change point in patient satisfaction at the  $\alpha = 0.05$  level, because  $0.0003 < 0.005$ .

CNL Introduction *	Estimated Change Point $\hat{\tau}$	P-value <sup>+</sup>	Formal Intervention Implementation	Lag <sup>**</sup>
January 2010	May 2010	0.0003	July 2010	-2

<sup>+</sup>The p-value for the supremum Wald test, ie, the p-value for the existence of a change point.

\* All clinical nurse leaders were integrated into their respective hospital units on January 2010.

\*\* The intervention lag is the difference between the estimated change point and the formal intervention introduction time point

**TABLE 3**

The unit-specific pre-change-point intercepts and slopes

Hospital Unit	Intercept Pre-Change Point $\hat{\beta}_{j0}$			Slope Pre-Change Point $\hat{\beta}_{j1}$		
	Estimate	95% CI	p-val	Estimate	95% CI	p-val
Stroke	64.32	(46.34, 82.31)	0	0.56	(-0.52, 1.64)	0.3
Surgical	72.8	(47.72, 97.88)	0	0.36	(-1.05, 1.77)	0.61
Cardiac	64.17	(37.08, 91.27)	0	0.31	(-1.3, 1.92)	0.7
Medical Surgical	70.19	(41.77, 98.61)	0	0.19	(-1.53, 1.91)	0.83
Mother/baby	77.1	(63.1, 91.09)	0	0.28	(-0.58, 1.15)	0.52

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE 4**

The unit-specific change in levels and change in slopes

Hospital Unit	Change in Level $-\hat{\delta}_j - \hat{\Delta}_j \hat{\tau}$			Change in Slope $\hat{\Delta}_j$		
	Estimate	95% CI	p-val	Estimate	95% CI	p-val
Stroke	6.91	(-14.65, 28.46)	0.52	-0.35	(-1.59, 0.89)	0.58
Surgical	6.17	(-20.22, 32.56)	0.64	-0.21	(-1.87, 1.45)	0.8
Cardiac	-0.15	(-34.36, 34.06)	0.99	-0.22	(-2.25, 1.82)	0.83
Medical Surgical	0.3	(-40.57, 41.18)	0.99	-0.14	(-2.53, 2.24)	0.9
Mother/baby	3.73	(-22.1, 29.56)	0.77	-0.25	(-1.72, 1.23)	0.74

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript



**TABLE 5**

Estimates of the stochastic component parameters: the adjacent correlations and response standard deviations pre- and post-change points. All the adjacent correlations are relatively small and tend to switch from negative to positive post-intervention. The response standard deviations tend to decrease post-intervention

Hospital Unit	Pre-Change Point		Post-Change Point	
	Adjacent Correlation $\hat{\phi}_{j1}$	Standard Deviation $\hat{\sigma}_{j1}$	Adjacent Correlation $\hat{\phi}_{j2}$	Standard Deviation $\hat{\sigma}_{j2}$
Stroke	-0.06	3.15	-0.35	3.01
Surgical	-0.02	4.37	0.19	2.35
Cardiac	-0.16	5.28	0.10	3.76
Medical Surgical	-0.03	4.84	0.09	5.06
Mother/baby	-0.27	2.97	0.08	3.72

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript