Conducting interrupted time-series analysis with panel data: The `xtitsa` command

Ariel Linden
Linden Consulting Group, LLC
San Francisco, CA, USA
alinden@lindenconsulting.org

Abstract. Interrupted time-series analysis (ITSA) is a popular study design when conducting a randomized experiment is not feasible. The design is called an *interrupted time series* because the intervention is expected to "interrupt" the level and/or trend of the outcome variable —measured at equal intervals over time—subsequent to its introduction. The ITSA design is most commonly used when there is a single aggregated treated unit (e.g., hospital, county, state) and the outcome of interest is reported at the summary level (e.g., morbidity or mortality rates). This article introduces the `xtitsa` command which generalizes the ITSA design for situations in which data are available at a more granular level (e.g., wards within hospital, zip codes within county, counties within state), allowing for the variability within the treatment group to be accounted for during estimation. The `xtitsa` command estimates the effect of an intervention on an outcome variable either for a single treatment group or when compared with a control group, and can estimate effects for multiple sequential interventions. Finally, as `xtitsa` is a wrapper for `xtgee`, there is tremendous flexibility in modelling outcomes with various distributions and autocorrelation structures.

Keywords: interrupted time series analysis, quasi-experimental designs, natural experiment, causal inference

1 Introduction

Interrupted time series analysis (ITSA) is a popular study design for evaluating the effectiveness of large-scale interventions and policy changes, in which an aggregate entity (e.g., hospital, city, region) is the treatment unit, and accordingly, the outcome is reported at the summary level (e.g., morbidity or mortality rates). The design is called an *interrupted time series* because the intervention is expected to "interrupt" the level and/or trend of the outcome variable—serially measured over time—subsequent to its introduction (Campbell and Stanley 1966; Shadish, Cook, and Campbell 2002).

Specific features of an ITSA make this design particularly appealing, such as requiring only a single treatment unit (making this useful for n-of-1 trials), not requiring a control group for comparison (which is advantageous in situations where controls may be difficult to find), using aggregated data (which are often publicly available and therefore easier to obtain than individual-level data), and displaying the outcomes graphically as a complement to the statistical results. Currently, the ITSA design can be implemented in Stata using the `itsa` command (Linden 2015; Linden 2017a).

Given the appeal of the ITSA design, it is tempting for researchers to collapse individual-level data into a single aggregate treatment unit to facilitate the use of this analytic approach. However, in doing so, information is lost about the underlying individual-level units that is important for statistical inference. To illustrate, the estimates table below presents the results of a simple regression of an outcome on a covariate *t*, first estimated using data from 10 individuals, and then again after collapsing the 10 individuals into a single aggregate unit. As shown, the point estimates are the same for both models, but the individual-level model indicates a high degree of variability between individuals that results in a statistically non-significant estimate (p=0.0786). Conversely, the collapsed data have much less variability, and the estimate for *t* is statistically significant (p=0.0001).

```
. estimates table individual collapsed, se p
```

Variable	individual	collapsed
t	3.9487898	3.9487904
	2.218895	.46303329
	0.0786	0.0001
_cons	26.142995	26.142986
	35.961619	7.5043779
	0.4692	0.0102

legend: b/se/p

In this article I introduce the `xtitsa` command which generalizes the ITSA design for situations in which data are available at a more granular (individual) level (e.g., wards within hospital, zip codes within county, counties within state), allowing for the variability within the treatment group to be accounted for during estimation. The `xtitsa` command estimates the effect of an intervention on an outcome variable either for a single treatment group (pre- versus post-intervention) or when the treatment group is compared to a control group, and can estimate effects for multiple sequential interventions. Finally, as `xtitsa` is a wrapper for `xtgee`, outcomes can be modelled using the appropriate form for the distribution and autocorrelation structure.

Although `xtitsa` utilizes the same analytic framework as the `itsa` command (Linden 2015) it is not intended as a replacement for `itsa`. `itsa` should be used when data are aggregated to a single unit because the underlying models used in the command are designed to accommodate univariate time series data (see [TS] **`newey`** and [TS] **`prais`**). Applying `xtitsa` to data from a single unit will likely produce a non-convergence error in the underlying `xtgee` procedure. Additionally, the robust standard errors option cannot be specified with only one panel. Conversely, `xtitsa` offers great flexibility in estimating ITSA models with individual-level data,

including the choice of distribution and link function, and the ability to specify the appropriate time related within-group correlation structure (e.g., autoregressive, stationary or non-stationary processes [see [XT] **xtgee correlation**]).

2 Method and formulas

2.1 The single-group analysis

When only the treatment group is under study (no comparison group), and with only a single treatment period, the general ITSA regression model (Huitema and McKean 2000; Linden 2015; Linden 2017a; Simonton 1977a; Simonton 1977b) assumes the following form (modified for individual-level analysis):

$$Y_{ti} = \beta_0 + \beta_1 T_{ti} + \beta_2 X_{ti} + \beta_3 X_{ti} T_{ti} + \epsilon_{ti} \quad (1)$$

where Y_{ti} is the outcome variable measured at each equally spaced time point t for each individual-level i, T_{ti} is the time since the start of the study, X_{ti} is a dummy (indicator) variable representing the intervention (pre-intervention periods 0, otherwise 1), and $X_{ti}T_{ti}$ is an interaction term. In the case of a single-group study, β_0 represents the intercept or starting level of the outcome variable. β_1 is the slope or trend of the outcome variable until the introduction of the intervention. β_2 represents the change in the level of the outcome that occurs in the period immediately following the introduction of the intervention. β_3 represents the difference between pre-intervention and post-intervention slopes of the outcome. Thus, we look for significant p-values in β_2 to indicate an immediate treatment effect, or in β_3 to indicate a treatment effect over time (Linden 2015; Linden 2017a).

By design, a single-group ITSA has no comparable control group; rather, the preintervention trend projected into the treatment period serves as the counterfactual. We assume that any time-varying unmeasured confounder is relatively slowly changing so that it would be distinguishable from the sharp jump of the intervention indicator. This underscores the need for caution with these methods if there are multiple policy shifts occurring in the time window around the implementation of the intervention (Linden 2017b).

2.2 The multiple-group analysis

When one or more control groups are available for comparison, and there is only one intervention, the regression model in Equation 1 is expanded to include four additional terms (β_4 to β_7) (Linden 2015; Linden 2017a; Simonton 1977a; Simonton 1977b):

$$Y_{ti} = \beta_0 + \beta_1 T_{ti} + \beta_2 X_{ti} + \beta_3 X_{ti} T_{ti} + \beta_4 Z_i + \beta_5 Z_i T_{ti} + \beta_6 Z_i X_{ti} + \beta_7 Z_i X_{ti} T_{ti} + \epsilon_{ti} \quad (2)$$

where Z_i is a dummy variable to denote the individual's cohort assignment (treatment or control), and $Z_i T_{ti}$, $Z_i X_{ti}$, and $Z_i X_{ti} T_{ti}$ are all interaction terms among previously described variables. The

coefficients β_0 to β_3, represent the control group, and the coefficients β_4 to β_7, represent values of the treatment group. More specifically, β_4 represents the difference in the level (intercept) of the outcome variable between treatment and controls prior to the intervention, β_5 represents the difference in the slope (trend) of the outcome variable between treatment and controls prior to the intervention, β_6 indicates the difference between treatment and control groups in the level of the outcome variable immediately following introduction of the intervention, and β_7 represents the difference between treatment and control groups in the slope (trend) of the outcome variable after initiation of the intervention compared with the pre-intervention (Linden 2015; Linden 2017a).

Identification in both the single- and multiple-group models is driven by the functional-form assumptions of the ITSA model. In the multiple-group analysis, the control group serves as the counterfactual to the treatment group, and the key assumption is that the change in the level or trend in the outcome variable is presumed to be the same both for the control group and, counterfactually, for the treatment group had it not received the intervention. In other words, we assume that confounding omitted variables affect both treatment and control groups similarly. A major strength of the multiple-group ITSA is the ability to test for comparability between groups on observed covariates and in particular, the two parameters β_4 and β_5, which play a particularly important role in establishing whether the treatment and control groups are balanced on both the level and the trajectory of the outcome variable in the preintervention period (Linden 2015).

2.3 Data variables corresponding to model parameters

As with `itsa`, `xtitsa` generates all the variables used in regression models (1) and (2). Table 1 displays these variables, using an artificial example with one intervention period. There are two individuals in these data (ID = 1, 2) with seven observations each starting at $T = 0$. X indicates that there are two preintervention period observations, followed by four observations in the intervention period (the intervention commences when $T = 2$). XT is an interaction term of $X \times T$, which starts in the observation period immediately following the start of the intervention ($T = 3$) and runs sequentially until the last observation when $T = 6$ (see Huitema and McKean [2000] for an exposition on the appropriateness of commencing the sequence in the observation period after the start of the intervention). Here we transform $XT = (T - 2) \times X$ so that it runs sequentially starting at 1. Additional variables are required for a multiple-group analysis. Z indicates the treatment status, where $Z = 1$ for the treatment group and $Z = 0$ for the control group. ZT, ZX, and ZXT are additional interaction terms used in multiple-group comparisons, as described above in section 2.2. When multiple treatment periods are specified, additional variables are added to the dataset, corresponding to each treatment period respectively (see section 4.3 for an example).

The final variable added to the dataset is a clone of the dependent variable (Y). The user may find this helpful in cases when the dependent variable is specified using a time series operator (see [U] **11.4.4 Time-series varlists**).

Table 1: Covariates used in a single-group ITSA (*T*, *X*, *XT*) and multiple-group ITSA (*T*, *X*, *XT*, *Z*, *ZT*, *ZX*, *ZXT*) corresponding to regression models (1) and (2), respectively

ID	T	X	XT	Z	ZT	ZX	ZXT
1	0	0	0	0	0	0	0
1	1	0	0	0	0	0	0
1	2	1	0	0	0	0	0
1	3	1	1	0	0	0	0
1	4	1	2	0	0	0	0
1	5	1	3	0	0	0	0
1	6	1	4	0	0	0	0
2	0	0	0	1	0	0	0
2	1	0	0	1	1	0	0
2	2	1	0	1	2	1	0
2	3	1	1	1	3	1	1
2	4	1	2	1	4	1	2
2	5	1	3	1	5	1	3
2	6	1	4	1	6	1	4

3 The `xtitsa` command

This section describes the syntax of the `xtitsa` command and available options.

3.1 Syntax

```
xtitsa depvar [indepvars] [if][in] [weight], trperiod(numlist) [options]
```

indepvars may contain factor variables; see [U] **11.4.3 Factor variables**. *depvar* and *indepvars* may contain time-series operators; see [U] **11.4.4 Time-series varlists**. `iweight, fweight,` and `pweights` are allowed; see [U] **11.1.6 weight**.

The panel data must be strongly balanced and be declared to be time-series data by using either `tsset` *panelvar timevar* or `xtset` *panelvar timevar*. See [TS] **tsset** or [XT] **xtset**.

3.2 Options

`trperiod`(numlist) specifies the time period when the intervention begins. The value(s) entered for time period(s) must be in the same units as the panel time variable specified in `tsset timevar`; (see [TS] **tsset**). Dates should be specified as human readible dates using the respective pseudofunction (see [D] **datetime**), such as `trperiod`(2020) for a four-digit year,

or `trperiod`(2019m11) for quarterly data, etc. Multiple periods may be specified, separated with a semicolon as `trperiod`(2019m6; 2019m11); `trperiod`() is required.

`single` indicates that `xtitsa` will be used for a single-group analysis. Conversely, omitting single indicates that xtitsa is for a multiple-group comparison.

`treat`(varname) indicates the binary treatment variable (where the control group is equal to 0 and the treatment group is equal to 1). When the dataset contains data for only the treatment group, `treat`() must be omitted.

`posttrend` produces posttreatment trend estimates using lincom (see [R] **lincom**), for the specified model. In the case of a single-group ITSA, one estimate is produced. In the case of a multiple-group ITSA, an estimate is produced for the treatment group, the control group, and the difference. In the case of multiple treatment periods, a separate table is produced for each treatment period.

`prefix`(string) adds a prefix to the names of variables created by `xtitsa`. Short prefixes are recommended.

`replace` replaces variables created by xtitsa if they already exist. If `prefix`() is specified, only variables created by `xtitsa` with the same prefix will be replaced.

`figure`[(twoway_options)] produces a line plot of the average predicted depvar variable combined with a scatterplot of the average actual values of depvar over time. Specifying figure without options uses the default graph settings.

`model_options` specify all available options for xtgee (see [XT] **xtgee**).

3.3 Stored results

Because `xtitsa` passes all user-entered information to `xtgee`, all results stored by `xtgee` are available. Additionally, `xtitsa` generates several key time-series variables and adds them to the current dataset, as described in section 2.3. These additional variables allow the user to further estimate treatment effects using `arima` or other time-series models if desired.

Table 2 is a cross reference to default names for those variables that appear in the regression output tables (and used when `posttrend` is specified). Variables starting with `_z` are added to the dataset only when a multiple-group comparison is specified.

(`trperiod`) is a suffix added to certain variables indicating the start of the intervention period. This is particularly helpful for differentiating between added variables when multiple interventions are specified (see the example presented in section 4.3). If the user specifies a `prefix`(), it will be applied to all variables generated by `xtitsa`.

Table 2. Descriptions of default names for variables that appear in the regression output tables

Variable	Description
_*depvar*	dependent variable
_t	time since start of study
_x(trperiod)	dummy variable representing the intervention periods (preintervention periods 0, otherwise 1)
_x_ t(trperiod)	interaction of _x and _t
_z	dummy variable to denote the cohort assignment (treatment or control)
_z_x(trperiod)	interaction of _z and _x
_z_x_t(trperiod)	interaction of _z, _x, and _t
_s_*depvar*_pred	predicted value generated after running `xtitsa` for a single group
_m_*depvar*_pred	predicted value generated after running `xtitsa` for a multiple-group comparison

4 Examples

In this section, we demonstrate the use of `xtitsa` with an artificial dataset that includes 100 individuals (50 treated and 50 controls) followed over 20 months (January 2019 through August 2020). The data were generated to assume an intervention was initiated in November 2019, eliciting a sharp increase in the time-series in the month immediately following the initiation and continuing to rise monthly, thereafter. The control time-series was generated to show no change over time.

4.1 Single-group ITSA

In this example, we use `xtitsa` to assess the impact of the intervention, using a single-group design. More specifically, we assess whether the intervention resulted in a shift in the level and trend of the outcome compared with those of the preintervention period (as described in section 2.1).

First, we load the data and declare the dataset as panel:

```
. use "xtitsa_example.dta", clear

. tsset id month
       panel variable:  id (strongly balanced)
        time variable:  month, 2019m1 to 2020m8
                delta:  1 month
```

Next, we specify a single-group ITSA where `z` is the treatment group variable (and thus `z`=1 is treatment), 2019m1 is the start of the intervention, specify robust standard errors, request postintervention trend estimates, and plot the results. We leave all other default settings of `xtgee`.

```
. xtitsa y, single treat(z) trperiod(2019m11) vce(robust) posttrend figure replace

       panel variable:  id (strongly balanced)
        time variable:  month, 2019m1 to 2020m8
                delta:  1 month

Iteration 1: tolerance = 1.878e-13

GEE population-averaged model                   Number of obs     =      1,000
Group variable:                        id       Number of groups  =         50
Link:                            identity       Obs per group:
Family:                          Gaussian                     min =         20
Correlation:                  exchangeable                    avg =       20.0
                                                              max =         20
                                                Wald chi2(3)      =     341.81
Scale parameter:                  1632.008      Prob > chi2       =     0.0000

                                    (Std. Err. adjusted for clustering on id)
------------------------------------------------------------------------------
             |               Robust
           _y |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
           _t |   .6394942   .4271997     1.50   0.134    -.1978018     1.47679
    _x2019m11 |   21.00198   5.685869     3.69   0.000     9.857884    32.14608
 _x_t2019m11 |   5.347266   .6494279     8.23   0.000     4.074411    6.620121
        _cons |    44.9201   2.247136    19.99   0.000     40.51579     49.3244
------------------------------------------------------------------------------

                  Postintervention Linear Trend: 2019m11

Treated: _b[_t]+_b[_x_t2019m11]

------------------------------------------------------------------------------
 Linear Trend |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
--------------+---------------------------------------------------------------
      Treated |    5.98676   .4427186    13.52   0.000     5.119048    6.854473
------------------------------------------------------------------------------
```

As shown in the regression table, the starting level of the outcome (`_cons`) was estimated at 45 points, and appeared to increase slightly (but not statistically) every month (`_t`) prior to 2019m1, by 0.64 points ($P = 0.134$, CI = [−0.20, 1.48]). In the first month of the intervention (`_x2019m1`), there appeared to be a statistically significant jump in the time-series by 21.01 points ($P < 0.0001$, CI = [9.86, 32.15]), followed by a statistically significant increase in the monthly trend relative to the preintervention trend (`_xt2019m1`) of 5.35 points per month ($P < 0.0001$, CI = [4.07, 6.62]). We also see from the `lincom` estimate produced by specifying `posttrend`, that after initiation of the intervention, the time-series increased monthly at a rate of 5.99 points (95% CI = [5.12, 6.85]). Figure 1 provides a visual display of these results.

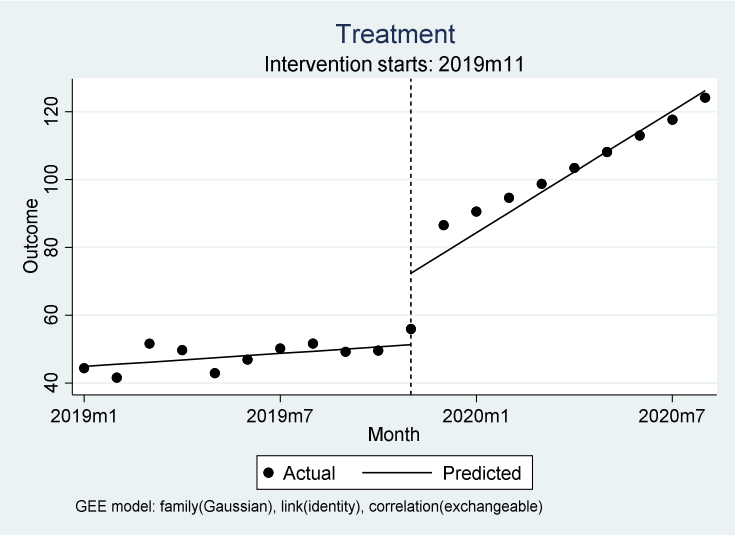

Figure 1. Single-group ITSA with default `xtgee` settings

To ensure that we fit a model that accounts for the correct autocorrelation structure, we use `actest` (Baum and Schaffer 2013) to test for autocorrelation.

```
. gen resid = y - _s__y_pred if e(sample)
(1,000 missing values generated)

. actest resid, lags(12) robust

Cumby-Huizinga test for autocorrelation
  H0: disturbance is MA process up to order q
  HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified				H0: q=specified lag-1 HA: s.c. present at lag specified			
lags	chi2	df	p-val	lag	chi2	df	p-val
1 - 1	175.256	1	0.0000	1	175.256	1	0.0000
1 - 2	175.272	2	0.0000	2	58.211	1	0.0000
1 - 3	178.540	3	0.0000	3	32.018	1	0.0000
1 - 4	178.645	4	0.0000	4	24.740	1	0.0000
1 - 5	179.634	5	0.0000	5	20.788	1	0.0000
1 - 6	179.676	6	0.0000	6	23.063	1	0.0000
1 - 7	181.309	7	0.0000	7	18.307	1	0.0000
1 - 8	181.491	8	0.0000	8	15.948	1	0.0001
1 - 9	181.578	9	0.0000	9	6.592	1	0.0102
1 - 10	181.633	10	0.0000	10	0.278	1	0.5977
1 - 11	181.635	11	0.0000	11	0.451	1	0.5017
1 - 12	181.636	12	0.0000	12	0.169	1	0.6810

Test robust to heteroskedasticity

As shown in the right-side panel of the output table, autocorrelation is present up to lag 9 but not at any higher lag orders (up to the twelve lags tested). Thus we re-estimate the model specifying `lag`(9) to correctly account for this autocorrelation.

```
. xtitsa y, single treat(z) trperiod(2019m11) vce(robust) posttrend figure replace corr(ar 9)

       panel variable:  id (strongly balanced)
        time variable:  month, 2019m1 to 2020m8
                delta:  1 month

Iteration 1: tolerance = .39506889
Iteration 2: tolerance = .02639033
Iteration 3: tolerance = .00137872
Iteration 4: tolerance = .00007392
Iteration 5: tolerance = 4.073e-06
Iteration 6: tolerance = 2.299e-07

GEE population-averaged model                   Number of obs     =      1,000
Group and time vars:              id month      Number of groups  =         50
Link:                             identity      Obs per group:
Family:                           Gaussian                    min =         20
Correlation:                         AR(9)                    avg =       20.0
                                                              max =         20
                                                Wald chi2(3)      =     311.22
Scale parameter:                  1646.467      Prob > chi2       =     0.0000

                                     (Std. Err. adjusted for clustering on id)
```

_y	Coef.	Robust Std. Err.	z	P>|z|	[95% Conf.	Interval]
_t	-.036334	.4571806	-0.08	0.937	-.9323915	.8597236
_x2019m11	18.28607	5.54179	3.30	0.001	7.424358	29.14778
_x_t2019m11	6.451831	.7877206	8.19	0.000	4.907927	7.995735
_cons	47.61581	2.787907	17.08	0.000	42.15161	53.08

Postintervention Linear Trend: 2019m11

Treated: _b[_t]+_b[_x_t2019m11]

Linear Trend	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Treated	6.415497	.5070516	12.65	0.000	5.421695	7.4093

While the actual estimates have changed somewhat accounting for the autocorrelation, we see that the treatment effects have remained consistent with those of the previous model.

4.2 Multiple-group ITSA

In this example, we use `xtitsa` to assess the impact of the intervention using a multiple-group design. More specifically, we now compare the treatment group's time-series (`z=1`) to that of the control group (`z=0`).

```
. xtitsa y, treat(z) trperiod(2019m11) vce(robust) posttrend figure replace

      panel variable:  id (strongly balanced)
       time variable:  month, 2019m1 to 2020m8
               delta:  1 month

Iteration 1: tolerance = 4.329e-13

GEE population-averaged model                   Number of obs     =      2,000
Group variable:                         id      Number of groups  =        100
Link:                             identity      Obs per group:
Family:                           Gaussian                    min =         20
Correlation:                  exchangeable                    avg =       20.0
                                                              max =         20
                                                Wald chi2(7)      =     373.89
Scale parameter:                  1243.027      Prob > chi2       =     0.0000

                                    (Std. Err. adjusted for clustering on id)
```

_y	Coef.	Robust Std. Err.	z	P>|z|	[95% Conf.	Interval]
_t	1.063626	.4100224	2.59	0.009	.2599965	1.867255
_z	-.2238399	3.415438	-0.07	0.948	-6.917976	6.470296
_z_t	-.4241313	.5905713	-0.72	0.473	-1.58163	.7333672
_x2019m11	-6.642734	3.13572	-2.12	0.034	-12.78863	-.4968352
_x_t2019m11	-1.15273	.6012455	-1.92	0.055	-2.33115	.0256891
_z_x2019m11	27.64472	6.46802	4.27	0.000	14.96763	40.3218
_z_x_t2019m11	6.499996	.8826056	7.36	0.000	4.770121	8.229871
_cons	45.14394	2.581977	17.48	0.000	40.08336	50.20452

```
Comparison of Linear Postintervention Trends: 2019m11

Treated    : _b[_t] + _b[_z_t] + _b[_x_t2019m11] + _b[_z_x_t2019m11]
Controls   : _b[_t] + _b[_x_t2019m11]
Difference : _b[_z_t] +  _b[_z_x_t2019m11]
```

Linear Trend	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
Treated	5.98676	.440477	13.59	0.000	5.123441	6.850079
Controls	-.0891049	.4611882	-0.19	0.847	-.9930172	.8148075
Difference	6.075865	.6377418	9.53	0.000	4.825914	7.325816

As shown in the regression table, the initial mean level difference between the treatment group and control group (`_z`) was not significant ($P = 0.224$, CI = [−6.92, 6.47]), and neither was the difference in the mean baseline slope (`_z_t`) ($P = 0.473$, CI = [−1.58,0.73]), indicating that the groups were balanced on pre-intervention level and trend. This is verified upon visual inspection of figure 2.

There is a statistically significant treatment effect in the starting level of the intervention (`_z_x2019m11`), of 27.64 points ($P < 0.0001$, CI = [14.97, 40.32]) as well as a statistically significant monthly increase in the pre–post trend compared with that of controls (`_z_x_t2019m11`) of 6.50 points per month ($P < 0.0001$, CI = [4.77, 8.23]). Additionally, we see from the `posttrend` output that the treatment group's trend increased monthly in the postintervention period by 5.99 points, while the control group's trend remained flat, with a

statistically significance difference between them of 6.08 points per month ($P < 0.0001$, CI = [4.83, 7.33]).

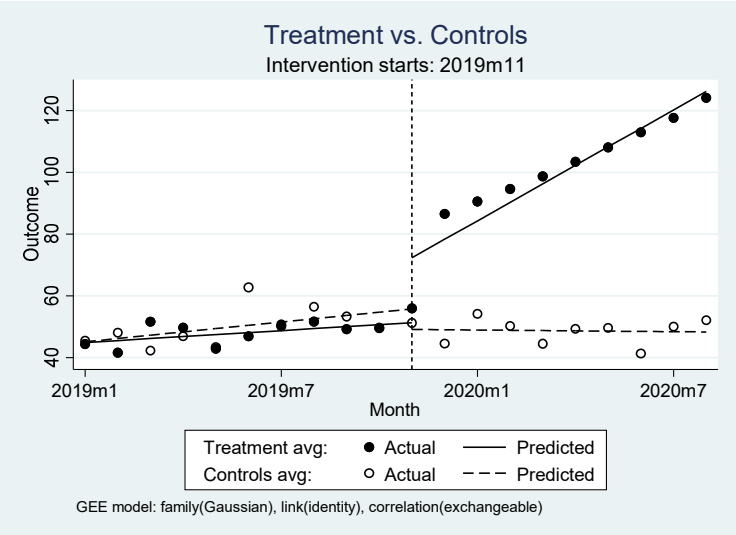

Figure 2. Multiple-group ITSA with default `xtgee` settings

4.3 Multiple treatment periods

`xtitsa` can accommodate design variations in which the effect of multiple treatment periods is of interest. For example, the researcher may be interested in studying the effects of an intervention that is introduced, withdrawn, and reintroduced, or an intervention that is followed by a separate intervention at a later point in time (see Barlow, Hayes, and Nelson [1984] for many other design alternatives). For exposition, in the following example we add an additional intervention to the data, starting in 2019m6. Thus, we re-estimate the single-group ITSA from section 4.1, now with one additional intervention period.

The interpretation of all coefficients up to the first intervention is as before. That is, the first intervention period is compared with the preintervention period. However, the additional coefficients for the second intervention period, `_x2019m11` and `_x_t2019m11`, are now compared with those of the prior (first) intervention period (Linden 2017a). As shown in both the regression table and verified upon visual inspection of figure 3, there is no evidence of a "treatment effect" beginning in 2019m6 while there was a treatment effect after the implementation of the second intervention (which in reality was the true intervention period).

```
. xtitsa y, single treat(z) trperiod(2019m6; 2019m11) vce(robust) posttrend replace fig

       panel variable:  id (strongly balanced)
        time variable:  month, 2019m1 to 2020m8
                delta:  1 month

Iteration 1: tolerance = 2.514e-12

GEE population-averaged model                 Number of obs     =      1,000
Group variable:                         id    Number of groups  =         50
Link:                             identity    Obs per group:
Family:                           Gaussian                  min =         20
Correlation:                  exchangeable                  avg =       20.0
                                                            max =         20
                                              Wald chi2(5)      =     345.18
Scale parameter:                  1631.97     Prob > chi2       =     0.0000

                                  (Std. Err. adjusted for clustering on id)
------------------------------------------------------------------------------
             |               Robust
          _y |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
          _t |   .5160163   1.258605     0.41   0.682    -1.950803    2.982836
      _x2019m6 |   1.049398   5.116332     0.21   0.837    -8.978429    11.07722
   _x_t2019m6 |  -.0837299   2.017461    -0.04   0.967    -4.037882    3.870422
     _x2019m11 |   21.49133    6.99034     3.07   0.002     7.790516    35.19214
  _x_t2019m11 |   5.554474   1.502247     3.70   0.000     2.610124    8.498824
        _cons |   45.03478   2.961353    15.21   0.000     39.23063    50.83893
------------------------------------------------------------------------------

                    Postintervention Linear Trend: 2019m6

Treated: _b[_t]+_b[_x_t2019m6]

------------------------------------------------------------------------------
 Linear Trend |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
--------------+---------------------------------------------------------------
      Treated |   .4322864   1.401609     0.31   0.758    -2.314817     3.17939
------------------------------------------------------------------------------

                    Postintervention Linear Trend: 2019m11

Treated: _b[_t]+_b[_x_t2019m6]+_b[_x_t2019m11]

------------------------------------------------------------------------------
 Linear Trend |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
--------------+---------------------------------------------------------------
      Treated |   5.98676   .4427186    13.52   0.000     5.119048    6.854473
------------------------------------------------------------------------------
```

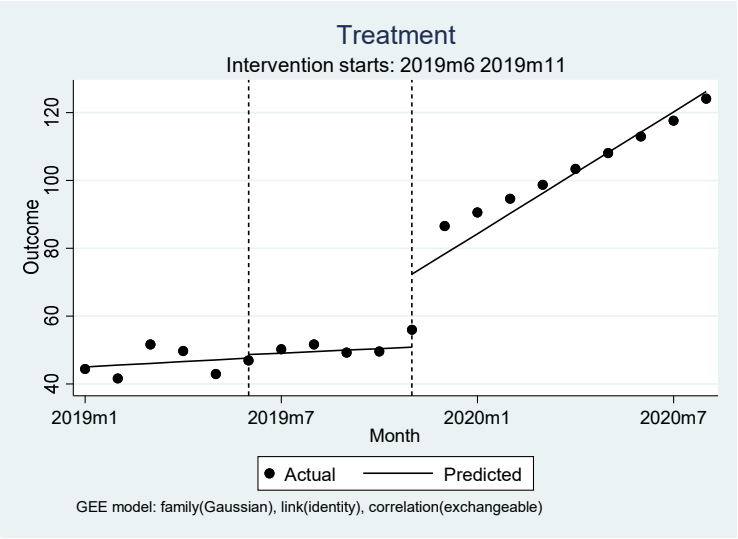

Figure 3. Single-group ITSA with default `xtgee` settings and two intervention periods

5 Discussion

In this article, I have demonstrated the basic implementation of `xtitsa` to estimate treatment effects using the ITSA design when data are available at the individual-level, allowing for the variability within the treatment group to be accounted for during estimation. `xtitsa` can estimate effects for a single treatment group, a multiple-group comparison, and when more than one intervention has been employed sequentially.

More-complex models can easily be estimated with `xtitsa` by including additional covariates to control for confounding, seasonal effects, and the impact of external events. Moreover, as `xtgee` is the underlying model used for estimation, the user can choose from a large array of `family` and `link` functions to fit the distribution of the outcome variable (the dataset and help file accompanying `xtitsa` provide an example of an analysis with an additional fractional outcome). Finally, many post-estimation measures can be computed following `xtitsa`, including those that fulfill the primary goal and those that provide supplementary information about trends (Linden 2017a). While `xtitsa` utilizes the same analytic approach as the `itsa` command (Linden 2015) it is not intended as a replacement for `itsa`. `itsa` should always be used when data are aggregated to a single unit because the underlying models used in the command are designed to accommodate univariate time series data.

6 References

Barlow, D. H., S. C. Hayes, and R. O. Nelson. 1984. *The Scientist Practitioner: Research and Accountability in Clinical and Educational Settings*. New York: Pergamon Press.

Baum, C. F., and M. E. Schaffer. 2013. actest: Stata module to perform Cumby– Huizinga general test for autocorrelation in time series. Statistical Software Components S457668, Department of Economics, Boston College. https://ideas.repec.org/c/boc/bocode/s457668.html.

Campbell, D. T., and J. C. Stanley. 1966. *Experimental and Quasi-experimental Designs for Research*. Chicago, IL: Rand McNally.

Huitema, B. E., and J. W. McKean. 2000. Design specification issues in time-series intervention models. *Educational and Psychological Measurement* 60: 38–58.

Linden, A. 2015. Conducting interrupted time-series analysis for single- and multiple-group comparisons. *Stata Journal* 15: 480-500.

Linden, A. 2017a. A comprehensive set of post-estimation measures to enrich interrupted time series analysis. *Stata Journal* 17: 73-88.

Linden, A. 2017b. Challenges to validity in single-group interrupted time series analysis. *Journal of Evaluation in Clinical Practice* 23: 413–418.

Shadish, W. R., Cook, T. D., and D. T. Campbell. 2002. *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Boston: Houghton Mifflin.

Simonton, D. K. 1977a. Cross-sectional time-series experiments: Some suggested statistical analyses. *Psychological Bulletin* 84: 489–502.

Simonton, D. K. 1977b. Erratum to Simonton. *Psychological Bulletin* 84: 1097.

8 Acknowledgement

I thank Kit Baum for assisting with the `actest` specification.

About the author

Ariel Linden is a health services researcher specializing in the evaluation of health care interventions and policy changes. He is both an independent consultant and a research scientist in the Department of Medicine, at the University of California, San Francisco. Thus far he has written 46 community-contributed packages for Stata.