

tssmooth shwinters — Holt–Winters seasonal smoothing
[Description](#)[Quick start](#)[Menu](#)[Syntax](#)[Options](#)[Remarks and examples](#)[Stored results](#)[Methods and formulas](#)[Acknowledgment](#)[References](#)[Also see](#)

Description

`tssmooth shwinters` performs the seasonal Holt–Winters method on a user-specified expression, which is usually just a variable name, and generates a new variable containing the forecasted series.

Quick start

Create `smooth` using Holt–Winters seasonal smoothing over `y` with `tsset` data

```
tssmooth shwinters smooth=y
```

As above, but forecast 10 periods out of sample

```
tssmooth shwinters smooth=y, forecast(10)
```

As above, but use 111 and 112 as the initial values for the recursion

```
tssmooth shwinters smooth=y, forecast(10) s0(111 112)
```

As above but use 0.5, 0.3, and 0.7 as the smoothing parameters

```
tssmooth shwinters smooth=y, forecast(10) s0(111 112) ///
    parms(.5 .3 .7)
```

As above, but normalize seasonal values

```
tssmooth shwinters smooth=y, forecast(10) s0(111 112) ///
    parms(.5 .3 .7) normalize
```

Note: The above commands can also be used to apply the smoother separately to each panel of a panel dataset when a *panelvar* has been specified using `tsset` or `xtset`.

Menu

Statistics > Time series > Smoothers/univariate forecasters > Holt–Winters seasonal smoothing

Syntax

```
tssmooth shwinters [type] newvar = exp [if] [in] [, options]
```

options Description

Main	
<code>replace</code>	replace <i>newvar</i> if it already exists
<code>parms(#α #β #γ)</code>	use # α , # β , and # γ as smoothing parameters
<code>samp0(#)</code>	use # observations to obtain initial values for recursion
<code>s0(#$_{\text{cons}}$ #$_{\text{lt}}$)</code>	use # $_{\text{cons}}$ and # $_{\text{lt}}$ as initial values for recursion
<code>forecast(#)</code>	use # periods for the out-of-sample forecast
<code>period(#)</code>	use # for period of the seasonality
<code>additive</code>	use additive seasonal Holt–Winters method
Options	
<code>sn0_0(<i>varname</i>)</code>	use initial seasonal values in <i>varname</i>
<code>sn0_v(<i>newvar</i>)</code>	store estimated initial values for seasonal terms in <i>newvar</i>
<code>snt_v(<i>newvar</i>)</code>	store final year's estimated seasonal terms in <i>newvar</i>
<code>normalize</code>	normalize seasonal values
<code>altstarts</code>	use alternative method for computing the starting values
Maximization	
<code>maximize_options</code>	control the maximization process; seldom used
<code>from(#α #β #γ)</code>	use # α , # β , and # γ as starting values for the parameters

You must `tsset` your data before using `tssmooth shwinters`; see [TS] [tsset](#).

exp may contain time-series operators; see [U] [11.4.4 Time-series varlists](#).

Options

Main

`replace` replaces *newvar* if it already exists.

`parms(# α # β # γ)`, $0 \leq \#_{\alpha} \leq 1$, $0 \leq \#_{\beta} \leq 1$, and $0 \leq \#_{\gamma} \leq 1$, specifies the parameters. If `parms()` is not specified, the values are chosen by an iterative process to minimize the in-sample sum-of-squared prediction errors.

If you experience difficulty converging (many iterations and “not concave” messages), try using `from()` to provide better starting values.

`samp0(#)` and `s0(# $_{\text{cons}}$ # $_{\text{lt}}$)` have to do with how the initial values # $_{\text{cons}}$ and # $_{\text{lt}}$ for the recursion are obtained.

`s0(# $_{\text{cons}}$ # $_{\text{lt}}$)` specifies the initial values to be used.

`samp0(#)` specifies that the initial values be obtained using the first # observations of the sample. This calculation is described under [Methods and formulas](#) and depends on whether the `altstart` and `additive` options are also specified.

If neither option is specified, the first half of the sample is used to obtain initial values.

`forecast(#)` specifies the number of periods for the out-of-sample prediction; $0 \leq \# \leq 500$. The default is `forecast(0)`, which is equivalent to not performing an out-of-sample forecast.

`period(#)` specifies the period of the seasonality. If `period()` is not specified, the seasonality is obtained from the `tsset` options `daily`, `weekly`, ..., `yearly`; see [TS] `tsset`. If you did not specify one of those options when you `tsset` the data, you must specify the `period()` option. For instance, if your data are quarterly and you did not specify `tsset`'s quarterly option, you must now specify `period(4)`.

By default, seasonal values are calculated, but you may specify the initial seasonal values to be used via the `sn0_0(varname)` option. The first `period()` observations of `varname` are to contain the initial seasonal values.

`additive` uses the additive seasonal Holt–Winters method instead of the default multiplicative seasonal Holt–Winters method.

Options

`sn0_0(varname)` specifies the initial seasonal values to use. `varname` must contain a complete year's worth of seasonal values, beginning with the first observation in the estimation sample. For example, if you have monthly data, the first 12 observations of `varname` must contain nonmissing data. `sn0_0()` cannot be used with `sn0_v()`.

`sn0_v(newvar)` stores in `newvar` the initial seasonal values after they have been estimated. `sn0_v()` cannot be used with `sn0_0()`.

`snt_v(newvar)` stores in `newvar` the seasonal values for the final year's worth of data.

`normalize` specifies that the seasonal values be normalized. In the multiplicative model, they are normalized to sum to one. In the additive model, the seasonal values are normalized to sum to zero.

`altstarts` uses an alternative method to compute the starting values for the constant, the linear, and the seasonal terms. The default and the alternative methods are described in [Methods and formulas](#). `altstarts` may not be specified with `s0()`.

Maximization

`maximize_options` controls the process for solving for the optimal α , β , and γ when the `parms()` option is not specified.

`maximize_options`: `nodifficult`, `technique(algorithm_spec)`, `iterate(#)`, `[no]log`, `trace`, `gradient`, `showstep`, `hessian`, `showtolerance`, `tolerance(#)`, `ltolerance(#)`, `nrtolerance(#)`, and `nonrntolerance`; see [R] [maximize](#). These options are seldom used.

`from(# α # β # γ)`, $0 < \#_\alpha < 1$, $0 < \#_\beta < 1$, and $0 < \#_\gamma < 1$, specifies starting values from which the optimal values of α , β , and γ will be obtained. If `from()` is not specified, `from(.5 .5 .5)` is used.

Remarks and examples

[stata.com](http://www.stata.com)

Remarks are presented under the following headings:

[Introduction](#)

[Holt–Winters seasonal multiplicative method](#)

[Holt–Winters seasonal additive method](#)

Introduction

The seasonal Holt–Winters methods forecast univariate series that have a seasonal component. If the amplitude of the seasonal component grows with the series, the Holt–Winters multiplicative method should be used. If the amplitude of the seasonal component is not growing with the series, the Holt–Winters additive method should be used. [Abraham and Ledolter \(1983\)](#), [Bowerman, O’Connell, and Koehler \(2005\)](#), and [Montgomery, Johnson, and Gardiner \(1990\)](#) provide good introductions to the Holt–Winters methods in recursive univariate forecasting methods. [Chatfield \(2001, 2004\)](#) provides introductions in the broader context of modern time-series analysis.

Like the other recursive methods in `tssmooth`, `tssmooth shwinters` uses the information stored by `tsset` to detect panel data. When applied to panel data, each series is smoothed separately, and the starting values are computed separately for each panel. If the smoothing parameters are chosen to minimize the in-sample sum-of-squared forecast errors, the optimization is performed separately on each panel.

When there are missing values at the beginning of the series, the sample begins with the first nonmissing observation. Missing values after the first nonmissing observation are filled in with forecasted values.

Holt–Winters seasonal multiplicative method

This method forecasts seasonal time series in which the amplitude of the seasonal component grows with the series. [Chatfield \(2001\)](#) notes that there are some nonlinear state-space models whose optimal prediction equations correspond to the multiplicative Holt–Winters method. This procedure is best applied to data that could be described by

$$x_{t+j} = (\mu_t + \beta j)S_{t+j} + \epsilon_{t+j}$$

where x_t is the series, μ_t is the time-varying mean at time t , β is a parameter, S_t is the seasonal component at time t , and ϵ_t is an idiosyncratic error. See [Methods and formulas](#) for the updating equations.

► Example 1: Forecasting from the multiplicative model

We have quarterly data on turkey sales by a new producer in the 1990s. The data have a strong seasonal component and an upward trend. We use the multiplicative Holt–Winters method to forecast sales for the year 2000. Because we have already `tsset` our data to the quarterly format, we do not need to specify the `period()` option.

```

. use http://www.stata-press.com/data/r14/turksales
. tssmooth shwinters shw1 = sales, forecast(4)
computing optimal weights
Iteration 0:   penalized RSS = -189.34609   (not concave)
Iteration 1:   penalized RSS = -108.68038
Iteration 2:   penalized RSS = -106.88286
Iteration 3:   penalized RSS = -106.15893
Iteration 4:   penalized RSS = -106.14094
Iteration 5:   penalized RSS = -106.14093
Iteration 6:   penalized RSS = -106.14093
Optimal weights:
                alpha = 0.1310
                beta  = 0.1428
                gamma = 0.2999
penalized sum-of-squared residuals = 106.1409
sum-of-squared residuals = 106.1409
root mean squared error = 1.628964

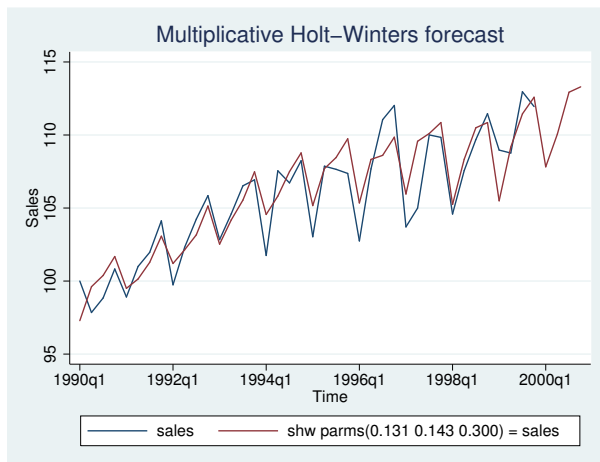
```

The graph below describes the fit and the forecast that was obtained.

```

. line sales shw1 t, title("Multiplicative Holt–Winters forecast")
> xtitle(Time) ytitle(Sales)

```



4

Holt–Winters seasonal additive method

This method is similar to the previous one, but the seasonal effect is assumed to be additive rather than multiplicative. This method forecasts series that can be described by the equation

$$x_{t+j} = (\mu_t + \beta j) + S_{t+j} + \epsilon_{t+j}$$

See [Methods and formulas](#) for the updating equations.

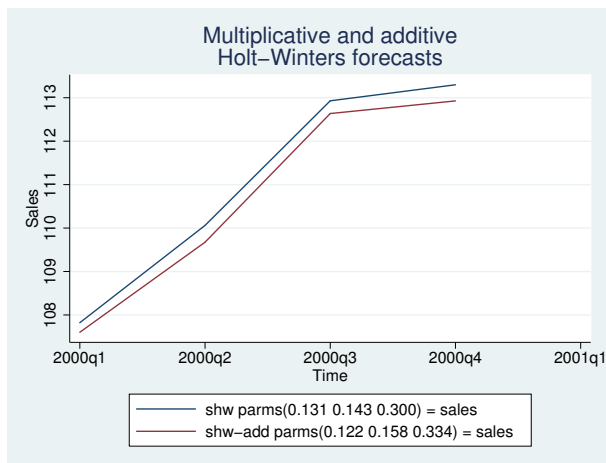
► Example 2: Forecasting from the additive model

In this example, we fit the data from the previous example to the additive model to forecast sales in the coming year. We use the `snt_v()` option to save the last year's seasonal terms in the new variable `seas`.

```
. tssmooth shwinters shwa = sales, forecast(4) snt_v(seas) normalize additive
computing optimal weights
Iteration 0:   penalized RSS = -190.90242   (not concave)
Iteration 1:   penalized RSS = -108.8357
Iteration 2:   penalized RSS = -108.25359
Iteration 3:   penalized RSS = -107.68187
Iteration 4:   penalized RSS = -107.66444
Iteration 5:   penalized RSS = -107.66442
Iteration 6:   penalized RSS = -107.66442
Optimal weights:
                alpha = 0.1219
                beta  = 0.1580
                gamma = 0.3340
penalized sum-of-squared residuals = 107.6644
sum-of-squared residuals = 107.6644
root mean squared error = 1.640613
```

The output reveals that the multiplicative model has a better in-sample fit, and the graph below shows that the forecast from the multiplicative model is higher than that of the additive model.

```
. line shw1 shwa t if t>=tq(2000q1), title("Multiplicative and additive"
> "Holt-Winters forecasts") xtitle("Time") ytitle("Sales") legend(cols(1))
```



To check whether the estimated seasonal components are intuitively sound, we list the last year's seasonal components.

```
. list t seas if seas < .
```

	t	seas
37.	1999q1	-2.7533393
38.	1999q2	-.91752573
39.	1999q3	1.8082417
40.	1999q4	1.8626233

The output indicates that the signs of the estimated seasonal components agree with our intuition. ◀

Stored results

`tssmooth shwinters` stores the following in `r()`:

Scalars

<code>r(N)</code>	number of observations	<code>r(N-pre)</code>	number of seasons used
<code>r(alpha)</code>	α smoothing parameter		in calculating starting values
<code>r(beta)</code>	β smoothing parameter	<code>r(s2_0)</code>	initial value for linear term
<code>r(gamma)</code>	γ smoothing parameter	<code>r(s1_0)</code>	initial value for constant term
<code>r(prss)</code>	penalized sum-of-squared errors	<code>r(linear)</code>	final value of linear term
<code>r(rss)</code>	sum-of-squared errors	<code>r(constant)</code>	final value of constant term
<code>r(rmse)</code>	root mean squared error	<code>r(period)</code>	period, if filter is seasonal

Macros

<code>r(method)</code>	<code>shwinters, additive or shwinters, multiplicative</code>	<code>r(exp)</code>	expression specified
		<code>r(timevar)</code>	time variable specified in <code>tsset</code>
<code>r(normalize)</code>	<code>normalize, if specified</code>	<code>r(panelvar)</code>	panel variable specified in <code>tsset</code>

Methods and formulas

A truncated description of the specified seasonal Holt–Winters filter labels the new variable. See [\[D\] label](#) for more information on labels.

An untruncated description of the specified seasonal Holt–Winters filter is saved in the characteristic named `tssmooth` for the new variable. See [\[P\] char](#) for more information on characteristics.

When the `parms()` option is not specified, the smoothing parameters are chosen to minimize the in-sample sum of penalized squared-forecast errors. Sometimes, one or more of the three optimal parameters lies on the boundary $[0, 1]$. To keep the estimates inside $[0, 1]$, `tssmooth shwinters` parameterizes the objective function in terms of their inverse logits, that is, $\exp(\alpha)/\{1 + \exp(\alpha)\}$, $\exp(\beta)/\{1 + \exp(\beta)\}$, and $\exp(\gamma)/\{1 + \exp(\gamma)\}$. When one of these parameters is actually on the boundary, this can complicate the optimization. For this reason, `tssmooth shwinters` optimizes a penalized sum-of-squared forecast errors. Let $\hat{x}_t(\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ be the forecast for the series x_t given the choices of $\tilde{\alpha}$, $\tilde{\beta}$, and $\tilde{\gamma}$. Then the in-sample penalized sum-of-squared prediction errors is

$$P = \sum_{t=1}^T \left[\{x_t - \hat{x}_t(\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})\}^2 + I_{|f(\tilde{\alpha})| > 12} (|f(\tilde{\alpha})| - 12)^2 + I_{|f(\tilde{\beta})| > 12} (|f(\tilde{\beta})| - 12)^2 + I_{|f(\tilde{\gamma})| > 12} (|f(\tilde{\gamma})| - 12)^2 \right]$$

where $f(x) = \ln\left(\frac{x}{1-x}\right)$. The penalty term is zero unless one of the parameters is close to the boundary. When one of the parameters is close to the boundary, the penalty term will help to obtain convergence.

Holt–Winters seasonal multiplicative procedure

As with the other recursive methods in `tssmooth`, there are three aspects to implementing the Holt–Winters seasonal multiplicative procedure: the forecasting equation, the initial values, and the updating equations. Unlike in the other methods, the data are now assumed to be seasonal with period L .

Given the estimates $a(t)$, $b(t)$, and $s(t + \tau - L)$, a τ step-ahead point forecast of x_t , denoted by $\hat{y}_{t+\tau}$, is

$$\hat{y}_{t+\tau} = \{a(t) + b(t)\tau\} s(t + \tau - L)$$

Given the smoothing parameters α , β , and γ , the updating equations are

$$a(t) = \alpha \frac{x_t}{s(t-L)} + (1 - \alpha) \{a(t-1) + b(t-1)\}$$

$$b(t) = \beta \{a(t) - a(t-1)\} + (1 - \beta) b(t-1)$$

and

$$s(t) = \gamma \left\{ \frac{x_t}{a(t)} \right\} + (1 - \gamma) s(t-L)$$

To restrict the seasonal terms to sum to 1 over each year, specify the `normalize` option.

The updating equations require the $L + 2$ initial values $a(0)$, $b(0)$, $s(1 - L)$, $s(2 - L)$, \dots , $s(0)$. Two methods calculate the initial values with the first m years, each of which contains L seasons. By default, m is set to the number of seasons in half the sample.

The initial value of the trend component, $b(0)$, can be estimated by

$$b(0) = \frac{\bar{x}_m - \bar{x}_1}{(m-1)L}$$

where \bar{x}_m is the average level of x_t in year m and \bar{x}_1 is the average level of x_t in the first year.

The initial value for the linear term, $a(0)$, is then calculated as

$$a(0) = \bar{x}_1 - \frac{L}{2} b(0)$$

To calculate the initial values for the seasons $1, 2, \dots, L$, we first calculate the deviation-adjusted values,

$$S(t) = \frac{x_t}{\bar{x}_i - \left\{ \frac{(L+1)}{2} - j \right\} b(0)}$$

where i is the year that corresponds to time t , j is the season that corresponds to time t , and \bar{x}_i is the average level of x_t in year i .

Next, for each season $l = 1, 2, \dots, L$, we define \bar{s}_l as the average S_t over the years. That is,

$$\bar{s}_l = \frac{1}{m} \sum_{k=0}^{m-1} S_{l+kL} \quad \text{for } l = 1, 2, \dots, L$$

Then the initial seasonal estimates are

$$s_{0l} = \bar{s}_l \left(\frac{L}{\sum_{l=1}^L \bar{s}_l} \right) \quad \text{for } l = 1, 2, \dots, L$$

and these values are used to fill in $s(1 - L), \dots, s(0)$.

If the `altstarts` option is specified, the starting values are computed based on a regression with seasonal indicator variables. Specifically, the series x_t is regressed on a time variable normalized to equal one in the first period in the sample and on a constant. Then $b(0)$ is set to the estimated coefficient on the time variable, and $a(0)$ is set to the estimated constant term. To calculate the seasonal starting values, x_t is regressed on a set of L seasonal dummy variables. The l th seasonal starting value is set to $(\frac{1}{\mu})\hat{\beta}_l$, where μ is the mean of x_t and $\hat{\beta}_l$ is the estimated coefficient on the l th seasonal dummy variable. The sample used in both regressions and the mean computation is restricted to include the first `samp0()` years. By default, `samp0()` includes half the data.

□ Technical note

If there are missing values in the first few years, a small value of m can cause the starting value methods for seasonal term to fail. Here you should either specify a larger value of m by using `samp0()` or directly specify the seasonal starting values by using the `snt0_0()` option. □

Holt–Winters seasonal additive procedure

This procedure is similar to the previous one, except that the data are assumed to be described by

$$x_t = (\beta_0 + \beta_1 t) + s_t + \epsilon_t$$

As in the multiplicative case, there are three smoothing parameters, α , β , and γ , which can either be set or chosen to minimize the in-sample sum-of-squared forecast errors.

The updating equations are

$$a(t) = \alpha \{x_t - s(t-L)\} + (1 - \alpha) \{a(t-1) + b(t-1)\}$$

$$b(t) = \beta \{a(t) - a(t-1)\} + (1 - \beta)b(t-1)$$

and

$$s(t) = \gamma \{x_t - a(t)\} + (1 - \gamma)s(t-L)$$

To restrict the seasonal terms to sum to 0 over each year, specify the `normalize` option.

A τ -step-ahead forecast, denoted by $\hat{y}_{t+\tau}$, is given by

$$\hat{x}_{t+\tau} = a(t) + b(t)\tau + s(t + \tau - L)$$

As in the multiplicative case, there are two methods for setting the initial values.

The default method is to obtain the initial values for $a(0)$, $b(0)$, $s(1-L)$, \dots , $s(0)$ from the regression

$$x_t = a(0) + b(0)t + \beta_{s,1-L}D_1 + \beta_{s,2-L}D_2 + \dots + \beta_{s,0}D_L + e_t$$

where the D_1, \dots, D_L are dummy variables with

$$D_i = \begin{cases} 1 & \text{if } t \text{ corresponds to season } i \\ 0 & \text{otherwise} \end{cases}$$

When `altstarts` is specified, an alternative method is used that regresses the x_t series on a time variable that has been normalized to equal one in the first period in the sample and on a constant term. $b(0)$ is set to the estimated coefficient on the time variable, and $a(0)$ is set to the estimated constant term. Then the demeaned series $\tilde{x}_t = x_t - \mu$ is created, where μ is the mean of the x_t . The \tilde{x}_t are regressed on L seasonal dummy variables. The l th seasonal starting value is then set to β_l , where β_l is the estimated coefficient on the l th seasonal dummy variable. The sample in both the regression and the mean calculation is restricted to include the first `samp0` years, where, by default, `samp0()` includes half the data.

Acknowledgment

We thank Nicholas J. Cox of the Department of Geography at Durham University, UK, and coeditor of the *Stata Journal* and author of *Speaking Stata Graphics* for his helpful comments.

References

- Abraham, B., and J. Ledolter. 1983. *Statistical Methods for Forecasting*. New York: Wiley.
- Bowerman, B. L., R. T. O'Connell, and A. B. Koehler. 2005. *Forecasting, Time Series, and Regression: An Applied Approach*. 4th ed. Pacific Grove, CA: Brooks/Cole.
- Chatfield, C. 2001. *Time-Series Forecasting*. London: Chapman & Hall/CRC.
- . 2004. *The Analysis of Time Series: An Introduction*. 6th ed. Boca Raton, FL: Chapman & Hall/CRC.
- Chatfield, C., and M. Yar. 1988. Holt-Winters forecasting: Some practical issues. *Statistician* 37: 129–140.
- Holt, C. C. 2004. Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting* 20: 5–10.
- Montgomery, D. C., L. A. Johnson, and J. S. Gardiner. 1990. *Forecasting and Time Series Analysis*. 2nd ed. New York: McGraw-Hill.
- Winters, P. R. 1960. Forecasting sales by exponentially weighted moving averages. *Management Science* 6: 324–342.

Also see

[TS] **tsset** — Declare data to be time-series data

[TS] **tssmooth** — Smooth and forecast univariate time-series data