tssmooth dexponential --- Double-exponential smoothing

Description Options References

Quick start Remarks and examples Also see Menu Stored results Syntax Methods and formulas

Description

tssmooth dexponential models the trend of a variable whose difference between changes from the previous values is serially correlated. More precisely, it models a variable whose second difference follows a low-order, moving-average process.

Quick start

Create smooth using a double-exponential smoother over y with tsset data tssmooth dexponential smooth=y

Same as above, but forecast 10 periods out of sample

tssmooth dexponential smooth=y, forecast(10)

Same as above, but use 111 and 112 as the initial values for the recursion tssmooth dexponential smooth=y, forecast(10) s0(111 112)

the second desponential smooth y, forecast(10) bo(111

Same as above, but use 0.5 as the smoothing parameter

tssmooth dexponential smooth=y, forecast(10) s0(111 112) parms(.5)

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

 $Statistics > {\sf Time\ series} > {\sf Smoothers/univariate\ forecasters} > {\sf Double-exponential\ smoothing}$

Syntax

tssmooth dexponential [type] newvar = exp [if] [in] [, options]

options	Description	
Main		
replace	replace newvar if it already exists	
$parms(\#_{\alpha})$	use $\#_{\alpha}$ as smoothing parameter	
<u>sa</u> mp0(#)	use # observations to obtain initial values for recursions	
$s0(\#_1 \#_2)$	use $\#_1$ and $\#_2$ as initial values for recursions	
<u>f</u> orecast(#)	use # periods for the out-of-sample forecast	

You must tsset your data before using tssmooth dexponential; see [TS] tsset. exp may contain time-series operators; see [U] 11.4.4 Time-series varlists. collect is allowed; see [U] 11.1.10 Prefix commands.

Options

🛛 Main 🛛

replace replaces *newvar* if it already exists.

parms $(\#_{\alpha})$ specifies the parameter α for the double-exponential smoothers; $0 < \#_{\alpha} < 1$. If parms $(\#_{\alpha})$ is not specified, the smoothing parameter is chosen to minimize the in-sample sum-of-squared forecast errors.

samp0(#) and s0($\#_1$ $\#_2$) are mutually exclusive ways of specifying the initial values for the recursion.

By default, initial values are obtained by fitting a linear regression with a time trend, using the first half of the observations in the dataset; see *Remarks and examples*.

samp0(#) specifies that the first # be used in that regression.

 $s0(\#_1 \#_2)$ specifies that $\#_1 \#_2$ be used as initial values.

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

Remarks and examples

The double-exponential smoothing procedure is designed for series that can be locally approximated as

$$\hat{x}_t = m_t + b_t t$$

where \hat{x}_t is the smoothed or predicted value of the series x, and the terms m_t and b_t change over time. Abraham and Ledolter (1983), Bowerman, O'Connell, and Koehler (2005), and Montgomery, Johnson, and Gardiner (1990) all provide good introductions to double-exponential smoothing. Chatfield (2001, 2004) provides helpful discussions of how double-exponential smoothing relates to modern time-series methods.

The double-exponential method has been used both as a smoother and as a prediction method. [TS] **tss-mooth exponential** shows that the single-exponential smoothed series is given by

$$S_t = \alpha x_t + (1 - \alpha)S_{t-1}$$

where α is the smoothing constant and x_t is the original series. The double-exponential smoother is obtained by smoothing the smoothed series,

$$S_t^{[2]} = \alpha S_t + (1-\alpha) S_{t-1}^{[2]}$$

Values of S_0 and $S_0^{[2]}$ are necessary to begin the process. Per Montgomery, Johnson, and Gardiner (1990), the default method is to obtain S_0 and $S_0^{[2]}$ from a regression of the first $N_{\rm pre}$ values of x_t on $\tilde{t} = (1, \ldots, N_{\rm pre} - t_0)'$. By default, $N_{\rm pre}$ is equal to one-half the number of observations in the sample. $N_{\rm pre}$ can be specified using the samp0() option.

The values of S_0 and $S_0^{[2]}$ can also be specified using the $\mathfrak{s0}$ () option.

Example 1: Smoothing a locally trending series

Suppose that we had some data on the monthly sales of a book and that we wanted to smooth this series. The graph below illustrates that this series is locally trending over time, so we would not want to use single-exponential smoothing.



The following example illustrates that double-exponential smoothing is simply smoothing the smoothed series. Because the starting values are treated as time-zero values, we actually lose 2 observations when smoothing the smoothed series.

```
. use https://www.stata-press.com/data/r19/sales2
. tssmooth exponential double sm1=sales, p(.7) s0(1031)
exponential coefficient
                                  0.7000
                         =
sum-of-squared residuals =
                                   13923
root mean squared error =
                                  13.192
. tssmooth exponential double sm2=sm1, p(.7) s0(1031)
                                  0.7000
exponential coefficient
                         =
sum-of-squared residuals =
                                  7698.6
root mean squared error
                                  9.8098
                        =
. tssmooth dexponential double sm2b=sales, p(.7) s0(1031 1031)
double-exponential coefficient
                                 =
                                          0.7000
sum-of-squared residuals
                                 =
                                          3724.4
root mean squared error
                                 =
                                          6.8231
. generate double sm2c = f2.sm2
(2 missing values generated)
. list sm2b sm2c in 1/10
                         sm2c
            sm2b
  1.
            1031
                         1031
  2
       1028.3834
                   1028.3834
  з.
       1030.6306
                   1030.6306
       1017.8182
                   1017.8182
  4.
  5.
        1022.938
                    1022.938
  6.
       1026.0752
                   1026.0752
  7.
       1041.8587
                   1041.8587
  8
       1042.8341
                   1042.8341
  9.
       1035.9571
                   1035.9571
 10.
       1030.6651
                   1030.6651
```

4

The double-exponential method can also be viewed as a forecasting mechanism. The exponential forecast method is a constrained version of the Holt–Winters method implemented in [TS] **tssmooth hwinters** (as discussed by Gardner [1985] and Chatfield [2001]). Chatfield (2001) also notes that the double-exponential method arises when the underlying model is an ARIMA(0,2,2) with equal roots.

This method produces predictions \hat{x}_t for $t = t_1, \ldots, T + \texttt{forecast}()$. These predictions are obtained as a function of the smoothed series and the smoothed-smoothed series. For $t \in [t_0, T]$,

$$\hat{x}_t = \Big(2 + \frac{\alpha}{1-\alpha}\Big)S_t - \Big(1 + \frac{\alpha}{1-\alpha}\Big)S_t^{[2]}$$

where S_t and $S_t^{[2]}$ are as given above.

The out-of-sample predictions are obtained as a function of the constant term, the linear term of the smoothed series at the last observation in the sample, and time. The constant term is $a_T = 2S_T - S_T^{[2]}$, and the linear term is $b_T = \frac{\alpha}{1-\alpha}(S_T - S_T^{[2]})$. The τ th-step-ahead out-of-sample prediction is given by

$$\hat{x}_t = a_t + \tau b_T$$

Example 2: Forecasting a locally trending series

Specifying the forecast option puts the double-exponential forecast into the new variable instead of the double-exponential smoothed series. The code given below uses the smoothed series sm1 and sm2 that were generated above to illustrate how the double-exponential forecasts are computed.

```
. tssmooth dexponential double f1=sales, p(.7) s0(1031 1031) forecast(4)
double-exponential coefficient = 0.7000
sum-of-squared residuals = 20737
root mean squared error = 16.1
. generate double xhat = (2 + .7/.3) * sm1 - (1 + .7/.3)* f.sm2
(5 missing values generated)
. list xhat f1 in 1/10
```

	xhat	f1
1.	1031	1031
2.	1031	1031
3.	1023.524	1023.524
4.	1034.8039	1034.8039
5.	994.0237	994.0237
6.	1032.4463	1032.4463
7.	1031.9015	1031.9015
8.	1071.1709	1071.1709
9.	1044.6454	1044.6454
10.	1023.1855	1023.1855

4

Example 3: Choosing an optimal parameter to forecast

Generally, when you are forecasting, you do not know the smoothing parameter. tssmooth dexponential computes the double-exponential forecasts of a series and obtains the optimal smoothing parameter by finding the smoothing parameter that minimizes the in-sample sum-of-squared forecast errors.

. tssmooth dexponential f2=sales, fored	cast(4)	
computing optimal double-exponential co	peffici	lent (0,1)
optimal double-exponential coefficient	=	0.3631
sum-of-squared residuals	=	16075.805
root mean squared error	=	14.175598

The following graph describes the fit that we obtained by applying the double-exponential forecast method to our sales data. The out-of-sample dynamic predictions are not constant, as in the single-exponential case.

. tsline f2 sales, title("Double exponential forecast with optimal alpha")



4

tssmooth dexponential automatically detects panel data from the information provided when the dataset was tsset. The starting values are chosen separately for each series. If the smoothing parameter is chosen to minimize the sum-of-squared prediction errors, the optimization is performed separately on each panel. The stored results contain the results from the last panel. Missing values at the beginning of the sample are excluded from the sample. After at least one value has been found, missing values are filled in using the one-step-ahead predictions from the previous period.

Stored results

tssmooth dexponential stores the following in r():

Scalars

	r(N)	number of observations
	r(alpha)	α smoothing parameter
	r(rss)	sum-of-squared errors
	r(rmse)	root mean squared error
	r(N_pre)	number of observations used in calculating starting values, if starting values calculated
	r(s2_0)	initial value for linear term, i.e., $S_0^{[2]}$
	r(s1_0)	initial value for constant term, i.e., S_0
	r(linear)	final value of linear term
	r(constant)	final value of constant term
	r(period)	period, if filter is seasonal
Mac	cros	
	r(method)	smoothing method
	r(exp)	expression specified
	r(timevar)	time variable specified in tsset
	r(panelvar)	panel variable specified in tsset

Methods and formulas

A truncated description of the specified double-exponential filter is used to label the new variable. See [D] **label** for more information on labels.

An untruncated description of the specified double-exponential filter is saved in the characteristic tssmooth for the new variable. See [P] char for more information on characteristics.

The updating equations for the smoothing and forecasting versions are as given previously.

The starting values for both the smoothing and forecasting versions of double-exponential are obtained using the same method, which begins with the model

$$x_t = \beta_0 + \beta_1 t$$

where x_t is the series to be smoothed and t is a time variable that has been normalized to equal 1 in the first period included in the sample. The regression coefficient estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ are obtained via OLS. The sample is determined by the samp0() option. By default, samp0() includes the first half of the observations. Given the estimates $\hat{\beta}_0$ and $\hat{\beta}_1$, the starting values are

$$\begin{split} S_0 &= \hat{\beta}_0 - \{(1-\alpha)/\alpha\} \hat{\beta}_1 \\ S_0^{[2]} &= \hat{\beta}_0 - 2\{(1-\alpha)/\alpha\} \hat{\beta}_1 \end{split}$$

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 and Hall/CRC.

_____. 2004. The Analysis of Time Series: An Introduction. 6th ed. Boca Raton, FL: Chapman and Hall/CRC.

- Chatfield, C., and M. Yar. 1988. Holt-Winters forecasting: Some practical issues. *Statistician* 37: 129–140. https://doi. org/10.2307/2348687.
- Gardner, E. S., Jr. 1985. Exponential smoothing: The state of the art. *Journal of Forecasting* 4: 1–28. https://doi.org/10. 1002/for.3980040103.
- Holt, C. C. 2004. Forecasting seasonals and trends by exponentially weighted moving averages. International Journal of Forecasting 20: 5–10. https://doi.org/10.1016/j.ijforecast.2003.09.015.
- 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. https://doi.org/10.1287/mnsc.6.3.324.

Also see

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

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

Stata, Stata Press, Mata, NetCourse, and NetCourseNow are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow is a trademark of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2025 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on citing Stata documentation.