

tsfilter bk — Baxter–King time-series filter

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

Description

`tsfilter bk` uses the [Baxter and King \(1999\)](#) band-pass filter to separate a time series into trend and cyclical components. The trend component may contain a deterministic or a stochastic trend. The stationary cyclical component is driven by stochastic cycles at the specified periods.

See [\[TS\] `tsfilter`](#) for an introduction to the methods implemented in `tsfilter bk`.

Quick start

Use the Baxter–King filter for `y` and obtain cyclical component `ct` using `tsset` data

```
tsfilter bk ct=y
```

As above, but filter out stochastic cycles with periods smaller than 5 and those with periods larger than 20

```
tsfilter bk ct=y, minperiod(5) maxperiod(20)
```

As above, and save the trend component in the variable `trendvar`

```
tsfilter bk ct=y, minperiod(5) maxperiod(20) trend(trendvar)
```

Save gain in `gain` and angular frequency in `angle`

```
tsfilter bk ct=y, gain(gain angle)
```

Use the Baxter–King filter for variables `y1`, `y2`, and `y3` to obtain cyclical components with prefix `cycl`

```
tsfilter bk cycl*=y1 y2 y3
```

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

Menu

Statistics > Time series > Filters for cyclical components > Baxter-King

Syntax

Filter one variable

```
tsfilter bk [type] newvar = varname [if] [in] [, options]
```

Filter multiple variables, unique names

```
tsfilter bk [type] newvarlist = varlist [if] [in] [, options]
```

Filter multiple variables, common name *stub*

```
tsfilter bk [type] stub* = varlist [if] [in] [, options]
```

options

Description

Main

minperiod(*#*)

filter out stochastic cycles at periods smaller than *#*

maxperiod(*#*)

filter out stochastic cycles at periods larger than *#*

smaorder(*#*)

number of observations in each direction that contribute to each filtered value

stationary

use calculations for a stationary time series

Trend

trend(*newvar* | *newvarlist* | *stub**) save the trend component(s) in new variable(s)

Gain

gain(*gainvar* *anglevar*) save the gain and angular frequency

You must `tsset` or `xtset` your data before using `tsfilter`; see [\[TS\] tsset](#) and [\[XT\] xtset](#).
varname and *varlist* may contain time-series operators; see [\[U\] 11.4.4 Time-series varlists](#).

Options

Main

`minperiod(#)` filters out stochastic cycles at periods smaller than *#*, where *#* must be at least 2 and less than `maxperiod()`. By default, if the units of the time variable are set to daily, weekly, monthly, quarterly, or half-yearly, then *#* is set to the number of periods equivalent to 1.5 years; yearly data use `minperiod(2)`; otherwise, the default value is `minperiod(6)`.

`maxperiod(#)` filters out stochastic cycles at periods larger than *#*, where *#* must be greater than `minperiod()`. By default, if the units of the time variable are set to daily, weekly, monthly, quarterly, half-yearly, or yearly, then *#* is set to the number of periods equivalent to 8 years; otherwise, the default value is `maxperiod(32)`.

`smaorder(#)` sets the order of the symmetric moving average, denoted by *q*. The order is an integer that specifies the number of observations in each direction used in calculating the symmetric moving average estimate of the cyclical component. This number must be an integer greater than zero and less than $(T - 1)/2$. The estimate for the cyclical component for the *t*th observation, y_t , is based upon the $2q + 1$ values $y_{t-q}, y_{t-q+1}, \dots, y_t, y_{t+1}, \dots, y_{t+q}$. By default, if the units of the time variable are set to daily, weekly, monthly, quarterly, half-yearly, or yearly, then *#* is set to the number of periods equivalent to 3 years; otherwise, the default value is `smaorder(12)`.

`stationary` modifies the filter calculations to those appropriate for a stationary series. By default, the series is assumed nonstationary.

Trend

`trend(newvar | newvarlist | stub*)` saves the trend component(s) in the new variable(s) specified by `newvar`, `newvarlist`, or `stub*`.

Gain

`gain(gainvar anglevar)` saves the gain in `gainvar` and its associated angular frequency in `anglevar`. Gains are calculated at the N angular frequencies that uniformly partition the interval $(0, \pi]$, where N is the sample size.

Remarks and examples

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

We assume that you have already read [TS] [tsfilter](#), which provides an introduction to filtering and the methods implemented in `tsfilter bk`, more examples using `tsfilter bk`, and a comparison of the four filters implemented by `tsfilter`. In particular, an understanding of gain functions as presented in [TS] [tsfilter](#) is required to understand these remarks.

`tsfilter bk` uses the Baxter–King (BK) band-pass filter to separate a time-series y_t into trend and cyclical components:

$$y_t = \tau_t + c_t$$

where τ_t is the trend component and c_t is the cyclical component. τ_t may be nonstationary; it may contain a deterministic or a stochastic trend, as discussed below.

The primary objective is to estimate c_t , a stationary cyclical component that is driven by stochastic cycles within a specified range of periods. The trend component τ_t is calculated by the difference $\tau_t = y_t - c_t$.

Although the BK band-pass filter implemented in `tsfilter bk` has been widely applied by macroeconomists, it is a general time-series method and may be of interest to other researchers.

Symmetric moving-average (SMA) filters with coefficients that sum to zero remove stochastic and deterministic trends of first and second order; see [Fuller \(1996\)](#), [Baxter and King \(1995\)](#), and [Baxter and King \(1999\)](#).

For an infinitely long series, there is an ideal band-pass filter for which the gain function is 1 for $\omega \in [\omega_0, \omega_1]$ and 0 for all other frequencies; see [TS] [tsfilter](#) for an introduction to gain functions. It just so happens that this ideal band-pass filter is an SMA filter with coefficients that sum to zero. [Baxter and King \(1999\)](#) derive the coefficients of this ideal band-pass filter and then define the BK filter to be the SMA filter with $2q + 1$ terms that are as close as possible to those of the ideal filter. There is a trade-off in choosing q : larger values of q cause the gain of the BK filter to be closer to the gain of the ideal filter, but they also increase the number of missing observations in the filtered series.

The `smaorder()` option specifies q . The default value of `smaorder()` is the number of periods equivalent to 3 years, following the [Baxter and King \(1999\)](#) recommendation.

Although the mathematics of the frequency-domain approach to time-series analysis is in terms of stochastic cycles at frequencies $\omega \in [-\pi, \pi]$, applied work is generally in terms of periods p , where $p = 2\pi/\omega$. So `tsfilter bk` has the `minperiod()` and `maxperiod()` options to specify the desired range of stochastic cycles.

Among economists, the BK filter is commonly used for investigating business cycles. Burns and Mitchell (1946) defined business cycles as stochastic cycles in business data corresponding to periods between 1.5 and 8 years. The default values for `minperiod()` and `maxperiod()` are the Burns–Mitchell values of 1.5 and 8 years, scaled to the frequency of the dataset. The calculations of the default values assume that the time variable is formatted as daily, weekly, monthly, quarterly, half-yearly, or yearly; see [D] [format](#).

For each variable, the band-pass BK filter estimate of c_t is put in the corresponding new variable, and when the `trend()` option is specified, the estimate of τ_t is put in the corresponding new variable.


`tsfilter bk` automatically detects panel data from the information provided when the dataset was `tsset` or `xtset`. All calculations are done separately on each panel. Missing values at the beginning and end of the sample are excluded from the sample. The sample may not contain gaps.

Baxter and King (1999) derived their method for nonstationary time series, but they noted that a small modification makes it applicable to stationary time series. Imposing the condition that the filter coefficients sum to zero is what makes their method applicable to nonstationary time series; dropping this condition yields a filter for stationary time series. Specifying the `stationary` option causes `tsfilter bk` to use coefficients calculated without the constraint that they sum to zero.

▷ Example 1: Estimating a business-cycle component

In this and the subsequent examples, we use `tsfilter bk` to estimate the business-cycle component of the natural log of real gross domestic product (GDP) of the United States. Our sample of quarterly data goes from 1952q1 to 2010q4. Below we read in and plot the data.


```
. use http://www.stata-press.com/data/r15/gdp2
(Federal Reserve Economic Data, St. Louis Fed)
. tsline gdp_ln
```



The series is nonstationary and is thus a candidate for the BK filter.

Below we use `tsfilter bk` to filter `gdp_ln`, and we use `pergram` (see [TS] [pergram](#)) to compute and to plot the periodogram of the estimated cyclical component.

```
. tsfilter bk gdp_bk = gdp_ln
. pergram gdp_bk, xline(.03125 .16667)
```



Because our sample is of quarterly data, `tsfilter bk` used the default values of `minperiod(6)`, `maxperiod(32)`, and `smaorder(12)`. The minimum and maximum periods are the [Burns and Mitchell \(1946\)](#) business-cycle periods for quarterly data. The default of `smaorder(12)` was recommended by [Baxter and King \(1999\)](#) for quarterly data.

In the periodogram, we added vertical lines at the natural frequencies corresponding to the conventional [Burns and Mitchell \(1946\)](#) values for business-cycle components. `pergram` displays the results in natural frequencies, which are the standard frequencies divided by 2π . We use the `xline()` option to draw vertical lines at the lower natural-frequency cutoff ($1/32 = 0.03125$) and the upper natural-frequency cutoff ($1/6 \approx 0.16667$).

If the filter completely removed the stochastic cycles at the unwanted frequencies, the periodogram would be a flat line at the minimum value of -6 outside the range identified by the vertical lines.

The periodogram reveals that the default value of `smaorder(12)` did not do a good job of filtering out the high-periodicity stochastic cycles, because there are too many points above -6.00 to the left of the left-hand vertical line. It also reveals that the filter did not remove enough low-periodicity stochastic cycles, because there are too many points above -6.00 to the right of the right-hand vertical line.

We address these problems in the next example.

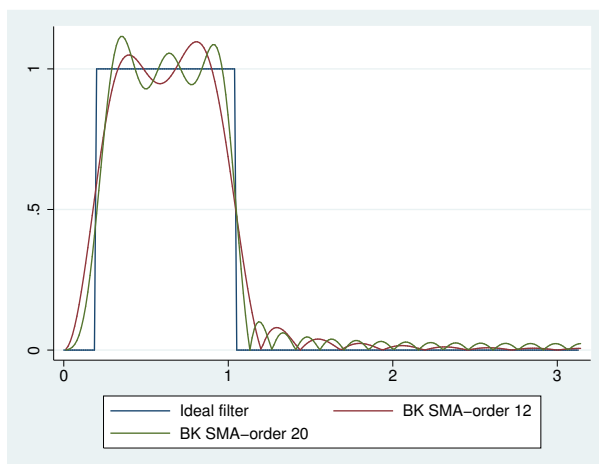
◀

► Example 2: Changing the order of the filter

In this example, we change the symmetric moving average of the filter via the `smaorder()` option so that it will remove more of the unwanted stochastic cycles. As mentioned, larger values of q cause the gain of the BK filter to be closer to the gain of the ideal filter, but larger values also increase the number of missing observations in the filtered series.

In the output below, we estimate the business-cycle component and compute the gain functions when the SMA-order of the filter is 12 and when it is 20. We also generate `ideal`, the gain function of the ideal band-pass filter at the frequencies `f`. Then we plot the gain functions from all three filters.

```
. tsfilter bk gdp_bk12 = gdp_ln, gain(g12 a12)
. label variable g12 "BK SMA-order 12"
. tsfilter bk gdp_bk20 = gdp_ln, gain(g20 a20) smaorder(20)
. label variable g20 "BK SMA-order 20"
. generate f = _pi*(_n-1)/_N
. generate ideal = cond(f<_pi/16, 0, cond(f<_pi/3, 1,0))
. label variable ideal "Ideal filter"
. twoway line ideal f || line g12 a12 || line g20 a20
```




As discussed in [TS] [tsfilter](#), the gain function of the ideal filter is a square wave with a value of 0 at the frequencies corresponding to unwanted frequencies and a value of 1 at the desired frequencies. The vertical lines in the gain function of the ideal filter occur at the frequencies $\pi/16$, corresponding to 32 periods, and at $\pi/3$, corresponding to 6 periods. (Given that $p = 2\pi/\omega$, where p is the period corresponding to the frequency ω , the frequency is given by $2\pi/p$.)

The differences between the gain function of the filter with SMA-order 12 and the gain function of the ideal band-pass filter is the root of the issues mentioned at the end of [example 1](#). The filter with SMA-order 20 is closer to the gain function of the ideal band-pass filter at the cost of 16 more missing values in the filtered series.

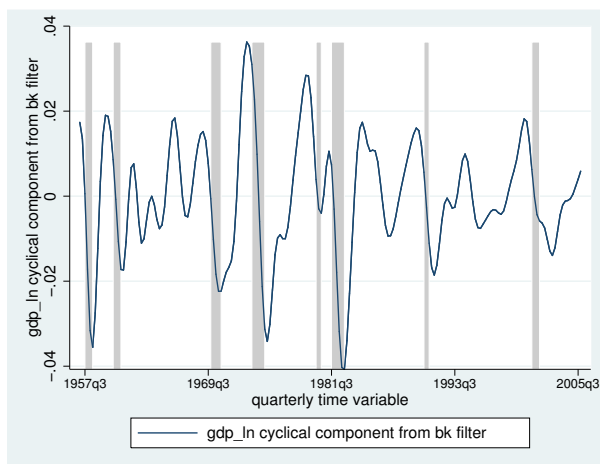
Below we compute and graph the periodogram of the series filtered with SMA-order 20.

```
. pergram gdp_bk20, xline(.03125 .16667)
```



The above periodogram indicates that the filter of SMA-order 20 removed more of the stochastic cycles at the unwanted periodicities than did the filter of SMA-order 12. Whether removing the stochastic cycles at the unwanted periodicities is worth losing more observations in the filtered series is a judgment call.

Below we plot the estimated business-cycle component with recessions identified by the shaded areas.



4

Stored results

`tsfilter bk` stores the following in `r()`:

Scalars

`r(smaorder)` order of the symmetric moving average
`r(minperiod)` minimum period of stochastic cycles
`r(maxperiod)` maximum period of stochastic cycles

Macros

`r(varlist)` original time-series variables
`r(filterlist)` variables containing estimates of the cyclical components
`r(trendlist)` variables containing estimates of the trend components, if `trend()` was specified
`r(method)` Baxter–King
`r(stationary)` yes or no, indicating whether the calculations assumed the series was or was not stationary
`r(unit)` units of time variable set using `tsset` or `xtset`

Matrices

`r(filter)` $(q+1) \times 1$ matrix of filter weights, where q is the order of the symmetric moving average

Methods and formulas

Baxter and King (1999) showed that there is an infinite-order SMA filter with coefficients that sum to zero that can extract the specified components from a nonstationary time series. The components are specified in terms of the minimum and maximum periods of the stochastic cycles that drive these components in the frequency domain. This ideal filter is not feasible, because the constraints imposed on the filter can only be satisfied using an infinite number of coefficients, so Baxter and King (1999) derived a finite approximation to this ideal filter.

The infinite-order, ideal band-pass filter obtains the cyclical component with the calculation

$$c_t = \sum_{j=-\infty}^{\infty} b_j y_{t-j}$$

Letting p_l and p_h be the minimum and maximum periods of the stochastic cycles of interest, the weights b_j in this calculation are given by

$$b_j = \begin{cases} \pi^{-1}(\omega_h - \omega_l) & \text{if } j = 0 \\ (j\pi)^{-1}\{\sin(j\omega_h) - \sin(j\omega_l)\} & \text{if } j \neq 0 \end{cases}$$

where $\omega_l = 2\pi/p_l$ and $\omega_h = 2\pi/p_h$ are the lower and higher cutoff frequencies, respectively.

For the default case of nonstationary time series with finite length, the ideal band-pass filter cannot be used without modification. [Baxter and King \(1999\)](#) derived modified weights for a finite order SMA filter with coefficients that sum to zero.

As a result, [Baxter and King \(1999\)](#) estimate c_t by

$$c_t = \sum_{j=-q}^{+q} \hat{b}_j y_{t-j}$$

The coefficients \hat{b}_j in this calculation are equal to $\hat{b}_j = b_j - \bar{b}_q$, where $\hat{b}_{-j} = \hat{b}_j$ and \bar{b}_q is the mean of the ideal coefficients truncated at $\pm q$:

$$\bar{b}_q = (2q + 1)^{-1} \sum_{j=-q}^q b_j$$

Note that $\sum_{j=-q}^{+q} \hat{b}_j = 0$ and that the first and last q values of the cyclical component cannot be estimated using this filter.

If the `stationary` option is set, the BK filter sets the coefficients to the ideal coefficients, that is, $\hat{b}_j = b_j$. For these weights, $\hat{b}_j = \hat{b}_{-j}$, and although $\sum_{j=-\infty}^{\infty} \hat{b}_j = 0$, for small q , $\sum_{j=-q}^q \hat{b}_j \neq 0$.

References

- Baxter, M., and R. G. King. 1995. Measuring business cycles approximate band-pass filters for economic time series. NBER Working Paper No. 5022, National Bureau of Economic Research. <http://www.nber.org/papers/w5022>.
- . 1999. Measuring business cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics* 81: 575–593.
- Burns, A. F., and W. C. Mitchell. 1946. *Measuring Business Cycles*. New York: National Bureau of Economic Research.
- Fuller, W. A. 1996. *Introduction to Statistical Time Series*. 2nd ed. New York: Wiley.
- Pollock, D. S. G. 1999. *A Handbook of Time-Series Analysis, Signal Processing and Dynamics*. London: Academic Press.
- . 2006. Econometric methods of signal extraction. *Computational Statistics & Data Analysis* 50: 2268–2292.

Also see

- [TS] **tsset** — Declare data to be time-series data
- [XT] **xtset** — Declare data to be panel data
- [TS] **tsfilter** — Filter a time-series, keeping only selected periodicities
- [D] **format** — Set variables' output format
- [TS] **tssmooth** — Smooth and forecast univariate time-series data