

tsfilter cf — Christiano–Fitzgerald time-series filter

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

Description

`tsfilter cf` uses the [Christiano and Fitzgerald \(2003\)](#) 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 cf`.

Quick start

Use the Christiano–Fitzgerald filter for `y` to obtain cyclical component `ct` using `tsset` data

```
tsfilter cf ct=y
```

As above, but filter for periods smaller than 5 and larger than 20

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

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

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

As above, and save gain and angular frequency with the names `gain` and `angle`

```
tsfilter cf ct=y, minperiod(5) maxperiod(20) trend(trendvar) ///
gain(gain angle)
```

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

```
tsfilter cf 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 > Christiano–Fitzgerald

Syntax

Filter one variable

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

Filter multiple variables, unique names

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

Filter multiple variables, common name *stub*

```
tsfilter cf [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

drift

remove drift from the 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*. By default, `smaorder()` is not set, which invokes the asymmetric calculations for the Christiano–Fitzgerald filter. 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 of the cyclical component for the *t*th observation, y_t , is based upon the $2q + 1$ values $y_{t-q}, y_{t-q+1}, \dots, y_t, y_{t+1}, \dots, y_{t+q}$.

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

drift removes drift using the approach described in [Christiano and Fitzgerald \(2003\)](#). By default, drift is not removed.

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 cf`, more examples using `tsfilter cf`, 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 cf` uses the Christiano–Fitzgerald (CF) 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 at a specified range of periods. The trend component τ_t is calculated by the difference $\tau_t = y_t - c_t$.

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

As discussed by [Christiano and Fitzgerald \(2003\)](#) and in [\[TS\] tsfilter](#), if one had an infinitely long series, one could apply an ideal band-pass filter that perfectly separates out cyclical components driven by stochastic cycles at the specified periodicities. In finite samples, it is not possible to exactly satisfy the conditions that a filter must fulfill to perfectly separate out the specified stochastic cycles; the expansive filter literature reflects the trade-offs involved in choosing a finite-length filter to separate out the specified stochastic cycles.

[Christiano and Fitzgerald \(2003\)](#) derive a finite-length CF band-pass filter that minimizes the mean squared error between the filtered series and the series filtered by an ideal band-pass filter that perfectly separates out components driven by stochastic cycles at the specified periodicities. [Christiano and Fitzgerald \(2003\)](#) place two important restrictions on the mean squared error problem that their filter solves. First, the CF filter is restricted to be a linear filter. Second, y_t is assumed to be a random-walk process; in other words, $y_t = y_{t-1} + \epsilon_t$, where ϵ_t is independent and identically distributed with mean zero and finite variance. The CF filter is the best linear predictor of the series filtered by the ideal band-pass filter when y_t is a random walk.

Christiano and Fitzgerald (2003) make four points in support of the random-walk assumption. First, the mean squared error problem solved by their filter requires that the process for y_t be specified. Second, they provide a method for removing drift so that their filter handles cases in which y_t is a random walk with drift. Third, many economic time series are well approximated by a random-walk-plus-drift process. (We add that many time series encountered in applied statistics are well approximated by a random-walk-plus-drift process.) Fourth, they provide simulation evidence that their filter performs well when the process generating y_t is not a random-walk-plus-drift process but is close to being a random-walk-plus-drift process.

Comparing the CF filter with the Baxter–King (BK) filter provides some intuition and explains the `smaorder()` option in `tsfilter cf`. As discussed in [TS] `tsfilter` and Baxter and King (1999), symmetric moving-average (SMA) filters with coefficients that sum to zero can extract the components driven by stochastic cycles at specified periodicities when the series to be filtered has a deterministic or stochastic trend of order 1 or 2.

The coefficients of the finite-length BK filter are as close as possible to the coefficients of an ideal SMA band-pass filter under the constraints that the BK coefficients are symmetric and sum to zero. The coefficients of the CF filter are not symmetric nor do they sum to zero, but the CF filter was designed to filter out the specified periodicities when y_t has a first-order stochastic trend.

To be robust to second-order trends, Christiano and Fitzgerald (2003) derive a constrained version of the CF filter. The coefficients of the constrained filter are constrained to be symmetric and to sum to zero. Subject to these constraints, the coefficients of the constrained CF filter minimize the mean squared error between the filtered series and the series filtered by an ideal band-pass filter that perfectly separates out the components. Christiano and Fitzgerald (2003) note that the higher-order detrending properties of this constrained filter come at the cost of lost efficiency. If the constraints are binding, the constrained filter cannot predict the series filtered by the ideal filter as well as the unconstrained filter can.

Specifying the `smaorder()` option causes `tsfilter cf` to compute the SMA-constrained CF filter.

The choice between the BK and the CF filters is one between robustness and efficiency. The BK filter handles a broader class of stochastic processes than does the CF filter, but the CF filter produces a better estimate of c_t if y_t is close to a random-walk process or a random-walk-plus-drift process.

Among economists, the CF 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`.

When y_t is assumed to be a random-walk-plus-drift process instead of a random-walk process, specify the `drift` option, which removes the linear drift in the series before applying the filter. Drift is removed by transforming the original series to a new series by using the calculation

$$z_t = y_t - \frac{(t-1)(y_T - y_1)}{T-1}$$

The cyclical component c_t is calculated from drift-adjusted series z_t . The trend component τ_t is calculated by $\tau_t = y_t - c_t$.

By default, the CF filter assumes the series is nonstationary. If the series is stationary, the `stationary` option is used to change the calculations to those appropriate for a stationary series.


For each variable, the CF 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 cf` 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.

► Example 1: Estimating a business-cycle component

In this and the subsequent examples, we use `tsfilter cf` 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 looks like it might be generated by a random-walk-plus-drift process and is thus a candidate for the CF filter.

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

```
. tsfilter cf gdp_cf = gdp_ln
. pergram gdp_cf, xline(.03125 .16667)
```




Because our sample is of quarterly data, `tsfilter cf` used the default values of `minperiod(6)` and `maxperiod(32)`. The minimum and maximum periods are the [Burns and Mitchell \(1946\)](#) business-cycle periods 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 CF did a reasonable job of filtering out the unwanted stochastic cycles.

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



4

Stored results

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

Scalars

`r(smaorder)` order of the symmetric moving average, if specified
`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)` Christiano–Fitzgerald
`r(symmetric)` yes or no, indicating whether the symmetric version of the filter was or was not used
`r(drift)` yes or no, indicating whether drift was or was not removed before filtering
`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 weights $(\hat{b}_0, \hat{b}_1, \dots, \hat{b}_q)'$, where q is the order of the symmetric moving average, and the weights are the Christiano–Fitzgerald coefficients; only returned when `smaorder()` is used to set q

Methods and formulas

For an infinitely long series, there is an ideal band-pass filter that extracts the cyclical component by using the calculation

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

If p_l and p_h are the minimum and maximum periods of the stochastic cycles of interest, the weights b_j in the ideal band-pass filter 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.

Because our time series has finite length, the ideal band-pass filter cannot be computed exactly. [Christiano and Fitzgerald \(2003\)](#) derive the finite-length CF band-pass filter that minimizes the mean squared error between the filtered series and the series filtered by an ideal band-pass filter that perfectly separates out the components. This filter is not symmetric nor do the coefficients sum to zero. The formula for calculating the value of cyclical component c_t for $t = 2, 3, \dots, T-1$ using the asymmetric version of the CF filter can be expressed as

$$c_t = b_0 y_t + \sum_{j=1}^{T-t-1} b_j y_{t+j} + \tilde{b}_{T-t} y_T + \sum_{j=1}^{t-2} b_j y_{t-j} + \tilde{b}_{t-1} y_1$$

where b_0, b_1, \dots are the weights used by the ideal band-pass filter. \tilde{b}_{T-t} and \tilde{b}_{t-1} are linear functions of the ideal weights used in this calculation. The CF filter uses two different calculations for \tilde{b}_t depending upon whether the series is assumed to be stationary or nonstationary.

For the default nonstationary case with $1 < t < T$, [Christiano and Fitzgerald \(2003\)](#) set \tilde{b}_{T-t} and \tilde{b}_{t-1} to

$$\tilde{b}_{T-t} = -\frac{1}{2}b_0 - \sum_{j=1}^{T-t-1} b_j \quad \text{and} \quad \tilde{b}_{t-1} = -\frac{1}{2}b_0 - \sum_{j=1}^{t-2} b_j$$

which forces the weights to sum to zero.

For the nonstationary case, when $t = 1$ or $t = T$, the two endpoints (c_1 and c_T) use only one modified weight, \tilde{b}_{T-1} :

$$c_1 = \frac{1}{2}b_0 y_1 + \sum_{j=1}^{T-2} b_j y_{j+1} + \tilde{b}_{T-1} y_T \quad \text{and} \quad c_T = \frac{1}{2}b_0 y_T + \sum_{j=1}^{T-2} b_j y_{T-j} + \tilde{b}_{T-1} y_1$$

When the `stationary` option is used to invoke the stationary calculations, all weights are set to the ideal filter weight, that is, $\tilde{b}_j = b_j$.

If the `smaorder()` option is set, the symmetric version of the CF filter is used. This option specifies the length of the symmetric moving average denoted by q . The symmetric calculations for c_t are similar to those used by the BK filter:

$$c_t = \hat{b}_q \{L^{-q}(y_t) + L^q(y_t)\} + \sum_{j=-q+1}^{q-1} b_j L^j(y_t)$$

where, for the default nonstationary calculations, $\hat{b}_q = -(1/2)b_0 - \sum_{j=1}^{q-1} b_j$. If the `smaorder()` and `stationary` options are set, then \hat{b}_q is set equal to the ideal weight b_q .

References

- Baxter, M., and R. G. King. 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.
- Christiano, L. J., and T. J. Fitzgerald. 2003. The band pass filter. *International Economic Review* 44: 435–465.
- 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