

tsfilter hp — Hodrick–Prescott time-series filter

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

Description

`tsfilter hp` uses the Hodrick–Prescott high-pass filter to separate a time series into trend and cyclical components. The trend component may contain a deterministic or a stochastic trend. The smoothing parameter determines the periods of the stochastic cycles that drive the stationary cyclical component.

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

Quick start

Use the Hodrick–Prescott filter for `y` to obtain cyclical component `ct` using `tsset` data

```
tsfilter hp ct=y
```

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

```
tsfilter hp ct=y, trend(trendvar)
```

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

```
tsfilter hp ct=y, trend(trendvar) gain(gain angle)
```

As above, but set the Hodrick–Prescott smoothing parameter to be 1700

```
tsfilter hp ct=y, trend(trendvar) gain(gain angle) smooth(1700)
```

Use Hodrick–Prescott filter for variables `y1`, `y2`, and `y3` to obtain cyclical components with prefix `cycl`

```
tsfilter hp 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 > Hodrick-Prescott

Syntax

Filter one variable

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

Filter multiple variables, unique names

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

Filter multiple variables, common name *stub*

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

options

Description

Main

smooth(#) smoothing parameter for the Hodrick–Prescott filter

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

smooth(#) sets the smoothing parameter for the Hodrick–Prescott filter. By default, if the units of the time variable are set to daily, weekly, monthly, quarterly, half-yearly, or yearly, then the Ravn–Uhlig rule is used to set the smoothing parameter; otherwise, the default value is **smooth**(1600). The Ravn–Uhlig rule sets # to $1600p_q^4$, where p_q is the number of periods per quarter. The smoothing parameter must be greater than 0.

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

We assume that you have already read [TS] [tsfilter](#), which provides an introduction to filtering and the methods implemented in `tsfilter hp`, more examples using `tsfilter hp`, 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 hp` uses the Hodrick–Prescott (HP) high-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 range of periods. The trend component τ_t is calculated by the difference $\tau_t = y_t - c_t$.

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

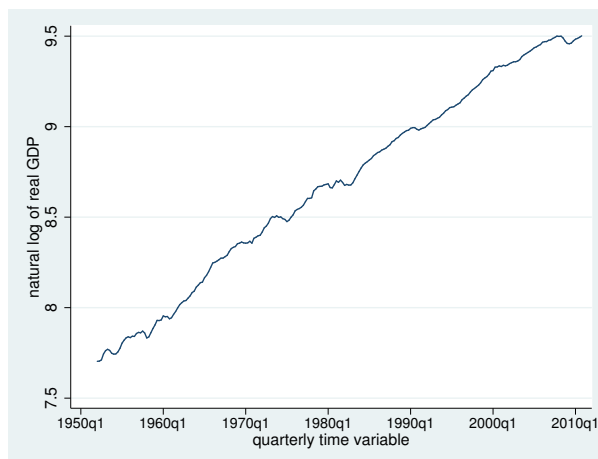
[Hodrick and Prescott \(1997\)](#) motivated the HP filter as a trend-removal technique that could be applied to data that came from a wide class of data-generating processes. In their view, the technique specified a trend in the data and the data was filtered by removing the trend. The smoothness of the trend depends on a parameter λ . The trend becomes smoother as $\lambda \rightarrow \infty$, and [Hodrick and Prescott \(1997\)](#) recommended setting λ to 1,600 for quarterly data.

[King and Rebelo \(1993\)](#) showed that removing a trend estimated by the HP filter is equivalent to a high-pass filter. They derived the gain function of this high-pass filter and showed that the filter would make integrated processes of order 4 or less stationary, making the HP filter comparable to the other filters implemented in `tsfilter`.

► Example 1: Estimating a business-cycle component

In this and the subsequent examples, we use `tsfilter hp` 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/r14/gdp2
(Federal Reserve Economic Data, St. Louis Fed)
. tsline gdp_ln
```



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

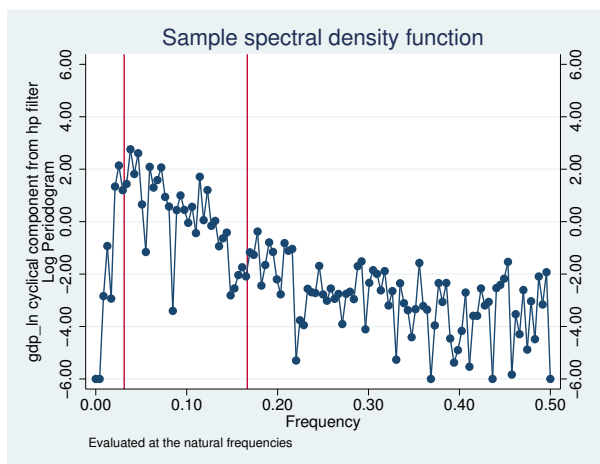
Below we use `tsfilter hp` to filter `gdp_ln`, and we use `pergram` (see [TS] [pergram](#)) to compute and to plot the periodogram of the estimated cyclical component.

```
. tsfilter hp gdp_hp = gdp_ln
. pergram gdp_hp, xline(.03125 .16667)
```

Because our sample is of quarterly data, `tsfilter hp` used the default value for the smoothing parameter of 1,600.

In the periodogram, we added vertical lines at the natural frequencies corresponding to the conventional [Burns and Mitchell \(1946\)](#) values for business-cycle components of 32 periods and 6 periods. `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 a high-periodicity issue and a low-periodicity issue. The points above -6.00 to the left of the left-hand vertical line in the periodogram reveal that the filter did not do a good job of filtering out the high-periodicity stochastic cycles with the default value smoothing parameter of 1,600. That there is no tendency of the points to the right of the right-hand vertical line to be smoothed toward -6.00 reveals that the HP filter did not remove any of the low-periodicity stochastic cycles. This result is not surprising, because the HP filter is a high-pass filter.

In the next example, we address the high-periodicity issue. See [TS] [tsfilter](#) and [TS] [tsfilter bw](#) for how to turn a high-pass filter into a band-pass filter.

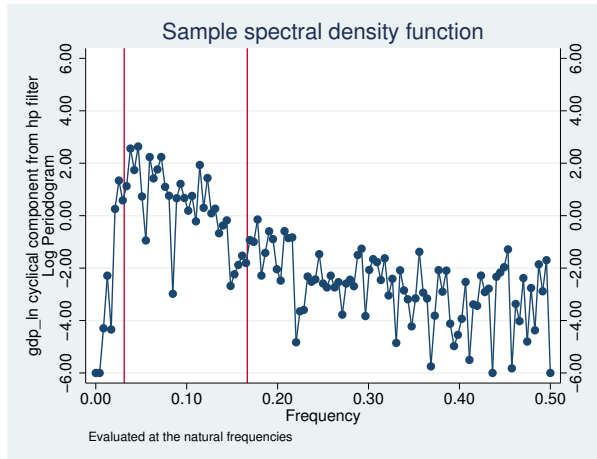
◀

► Example 2: Choosing the filter parameters

In the filter literature, filter parameters are set as functions of the cutoff frequency; see [Pollock \(2000, 324\)](#), for instance. This method finds the filter parameter that sets the gain of the filter equal to $1/2$ at the cutoff frequency. In a [technical note](#) in [TS] [tsfilter](#), we showed that applying this method to selecting λ at the cutoff frequency of 32 periods suggests setting $\lambda \approx 677.13$. In the output below, we

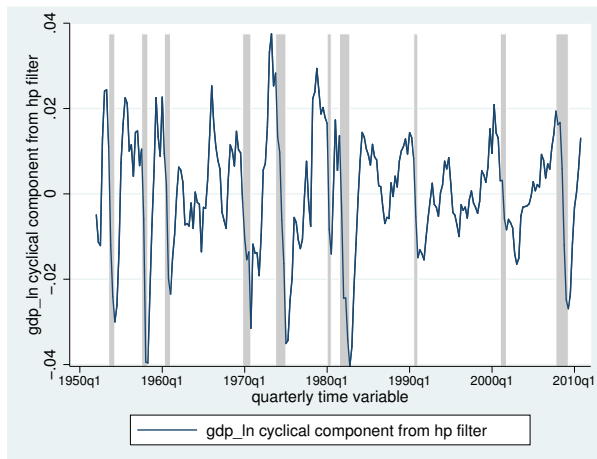
estimate the business-cycle component using this value for the smoothing parameter, and we compute and plot the periodogram of the estimated business-cycle component.

```
. tsfilter hp gdp_hp677 = gdp_ln, smooth(677.13)
. pergram gdp_hp677, xline(.03125 .16667)
```



A comparison of the two periodograms reveals that setting the smoothing parameter to 677.13 removes more of the high-periodicity stochastic cycles than does the default 1,600. In [TS] `tsfilter`, we found that the HP filter was not as good at removing the high-periodicity stochastic cycles as was the Christiano–Fitzgerald filter implemented in `tsfilter cf` or as was the Butterworth filter implemented in `tsfilter bw`.

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



4

`tsfilter hp` 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.

Stored results

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

Scalars

`r(smooth)` smoothing parameter λ

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)` Hodrick–Prescott

`r(unit)` units of time variable set using `tsset` or `xtset`

Methods and formulas

Formally, the filter is defined as the solution to the following optimization problem for τ_t

$$\min_{\tau_t} \left[\sum_{t=1}^T (y_t - \tau_t)^2 + \lambda \sum_{t=2}^{T-1} \{(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})\}^2 \right]$$

where the smoothing parameter λ is set fixed to a value.

If $\lambda = 0$, the solution degenerates to $\tau_t = y_t$, in which case the filter excludes all frequencies, that is, $c_t = 0$. On the other extreme, as $\lambda \rightarrow \infty$, the solution approaches the least-squares fit to the line $\tau_t = \beta_0 + \beta_1 t$; see [Hodrick and Prescott \(1997\)](#) for a discussion.

For a fixed λ , it can be shown that the cyclical component $\mathbf{c}' = (c_1, c_2, \dots, c_T)$ is calculated by

$$\mathbf{c} = (\mathbf{I}_T - \mathbf{M}^{-1})\mathbf{y}$$

where \mathbf{y} is the column vector $\mathbf{y}' = (y_1, y_2, \dots, y_T)$, \mathbf{I}_T is the $T \times T$ identity matrix, and \mathbf{M} is the $T \times T$ matrix:

$$\mathbf{M} = \begin{pmatrix} (1 + \lambda) & -2\lambda & \lambda & 0 & 0 & 0 & \dots & 0 \\ -2\lambda & (1 + 5\lambda) & -4\lambda & \lambda & 0 & 0 & \dots & 0 \\ \lambda & -4\lambda & (1 + 6\lambda) & -4\lambda & \lambda & 0 & \dots & 0 \\ 0 & \lambda & -4\lambda & (1 + 6\lambda) & -4\lambda & \lambda & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \dots & \vdots \\ \vdots & 0 & \ddots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & \lambda & -4\lambda & (1 + 6\lambda) & -4\lambda & \lambda & 0 \\ 0 & \dots & 0 & \lambda & -4\lambda & (1 + 6\lambda) & -4\lambda & \lambda \\ 0 & \dots & 0 & 0 & \lambda & -4\lambda & (1 + 5\lambda) & -2\lambda \\ 0 & \dots & 0 & 0 & 0 & \lambda & -2\lambda & (1 + \lambda) \end{pmatrix}$$

The gain of the HP filter is given by (see [King and Rebelo \[1993\]](#), [Maravall and del Rio \[2007\]](#), or [Harvey and Trimbur \[2008\]](#))

$$\psi(\omega) = \frac{4\lambda\{1 - \cos(\omega)\}^2}{1 + 4\lambda\{1 - \cos(\omega)\}^2}$$

As discussed in [TS] **tsfilter**, there are two approaches to selecting λ . One method, based on the heuristic argument of Hodrick and Prescott (1997), is used to compute the default values for λ . The method sets λ to 1,600 for quarterly data and to the rescaled values worked out by Ravn and Uhlig (2002). The rescaled default values for λ are 6.25 for yearly data, 100 for half-yearly data, 129,600 for monthly data, 1600×12^4 for weekly data, and $1600 \times (365/4)^4$ for daily data.

The second method for selecting λ uses the recommendations of Pollock (2000, 324), who uses the gain function of the filter to identify a value for λ .

Additional literature critiques the HP filter by pointing out that the HP filter corresponds to a specific model. Harvey and Trimbur (2008) show that the cyclical component estimated by the HP filter is equivalent to one estimated by a particular unobserved-components model. Harvey and Jaeger (1993), Gómez (1999), Pollock (2000), and Gómez (2001) also show this result and provide interesting comparisons of estimating c_t by filtering and model-based methods.

References

- Burns, A. F., and W. C. Mitchell. 1946. *Measuring Business Cycles*. New York: National Bureau of Economic Research.
- Gómez, V. 1999. Three equivalent methods for filtering finite nonstationary time series. *Journal of Business and Economic Statistics* 17: 109–116.
- . 2001. The use of Butterworth filters for trend and cycle estimation in economic time series. *Journal of Business and Economic Statistics* 19: 365–373.
- Harvey, A. C., and A. Jaeger. 1993. Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* 8: 231–247.
- Harvey, A. C., and T. M. Trimbur. 2008. Trend estimation and the Hodrick–Prescott filter. *Journal of the Japanese Statistical Society* 38: 41–49.
- Hodrick, R. J., and E. C. Prescott. 1997. Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit, and Banking* 29: 1–16.
- King, R. G., and S. T. Rebelo. 1993. Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17: 207–231.
- Leser, C. E. V. 1961. A simple method of trend construction. *Journal of the Royal Statistical Society, Series B* 23: 91–107.
- Maravall, A., and A. del Rio. 2007. Temporal aggregation, systematic sampling, and the Hodrick–Prescott filter. Working Paper No. 0728, Banco de España. <http://www.bde.es/webbde/Secciones/Publicaciones/PublicacionesSeriadadas/DocumentosTrabajo/07/Fic/dt0728e.pdf>.
- Pollock, D. S. G. 1999. *A Handbook of Time-Series Analysis, Signal Processing and Dynamics*. London: Academic Press.
- . 2000. Trend estimation and de-trending via rational square-wave filters. *Journal of Econometrics* 99: 317–334.
- . 2006. Econometric methods of signal extraction. *Computational Statistics & Data Analysis* 50: 2268–2292.
- Ravn, M. O., and H. Uhlig. 2002. On adjusting the Hodrick–Prescott filter for the frequency of observations. *Review of Economics and Statistics* 84: 371–376.

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