Getting started with mhpfilter

Muhammad Yaseen

2026-11-05

Table of contents

1 Abstract **1**

2 The data **2**

3 Three filters **2**
3.1 The HP filter 2
3.2 The modified HP filter 3
3.3 The fully modified HP filter 4

4 Comparing the filters **6**

5 Many series **7**

6 Reading the results **9**

7 Where next **10**

8 References **10**

1 Abstract

mhpfilter splits a macroeconomic time series into a trend and a cycle. This vignette takes the quarterly real GDP of the United States from the raw series to the three filters of the package, their comparison, many series at once and a plain-language reading of the results. The later vignettes give the theory (vignette 2), the two papers behind the package (vignettes 3 and 4), the simulation study (vignette 5), the application (vignette 6) and the checks of the code (vignette 7).

2 The data

`us_macro` holds quarterly real GDP, consumption and investment of the United States from 1959 to 2009 (FRED). As a `ts` the series keeps its dates, and every result of the package keeps them too.

```
data(us_macro)
head(us_macro)
```

```
  year quarter      gdp consumption investment
1 1959       1 2710.349      1707.4    286.898
2 1959       2 2778.801      1733.7    310.859
3 1959       3 2775.488      1751.8    289.226
4 1959       4 2785.204      1753.7    299.356
5 1960       1 2847.699      1770.5    331.722
6 1960       2 2834.390      1792.9    298.152
```

```
gdp <- ts(data = us_macro$gdp, start = c(1959, 1), frequency = 4)
```

3 Three filters

3.1 The HP filter

The filter of Hodrick and Prescott (1997) fits a smooth trend $\mathbf{g}$ to the series $\mathbf{x}$ by minimising

$$\sum_{t=1}^{T}(x_t - g_t)^2 + \lambda \sum_{t=2}^{T-1}\left[(g_{t+1} - g_t) - (g_t - g_{t-1})\right]^2,$$

and the cycle is $\mathbf{c} = \mathbf{x} - \mathbf{g}$. The smoothing parameter λ is conventionally 1600 for quarterly data. On the logarithm of a series the cycle is close to the percentage deviation from the trend.

```
hp <-
  hp_filter(
      x         = log(gdp)
    , lambda    = 1600
    , as_dt     = TRUE
    , log_scale = FALSE
    )
hp
```

With `as_dt = TRUE`, the default kept from version 0.1.0, the result is a `data.table`; `as_dt = FALSE` gives a list with methods for printing and plotting.

3.2 The modified HP filter

The modified HP filter of McDermott (1997), assessed by Choudhary, Hanif and Iqbal (2014), lets the data choose λ: it tries every $\lambda = 1, 2, ...$ and keeps the one that minimises a generalised cross-validation criterion.

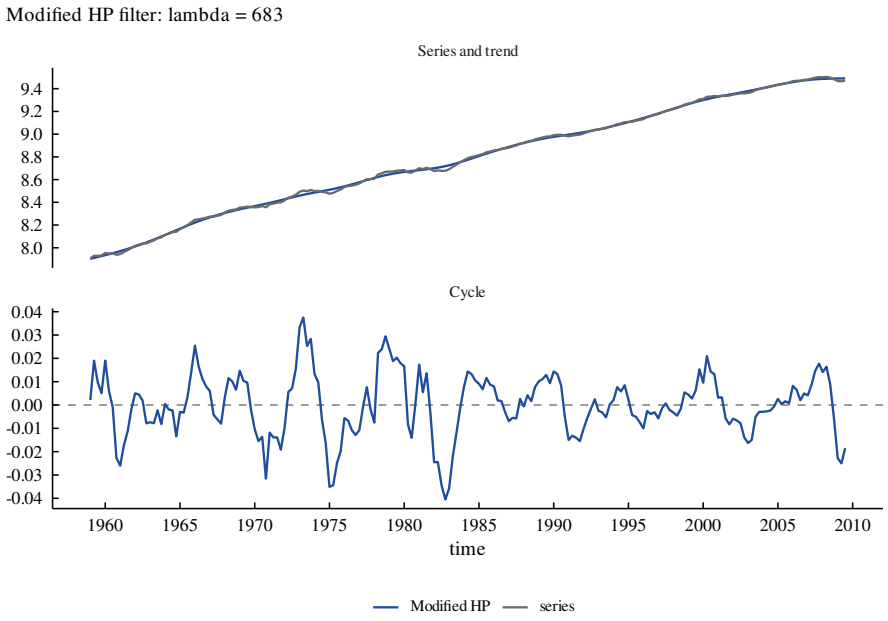
```
mhp <-
  mhp_filter(
      x          = log(gdp)
    , max_lambda = 100000
    , as_dt      = FALSE
    , log_scale  = FALSE
    , criterion  = "approximate"
    )
mhp
```

```
Modified HP Filter Result
==============================
Observations: 203
Optimal lambda: 683
GCV: 2.82e-04

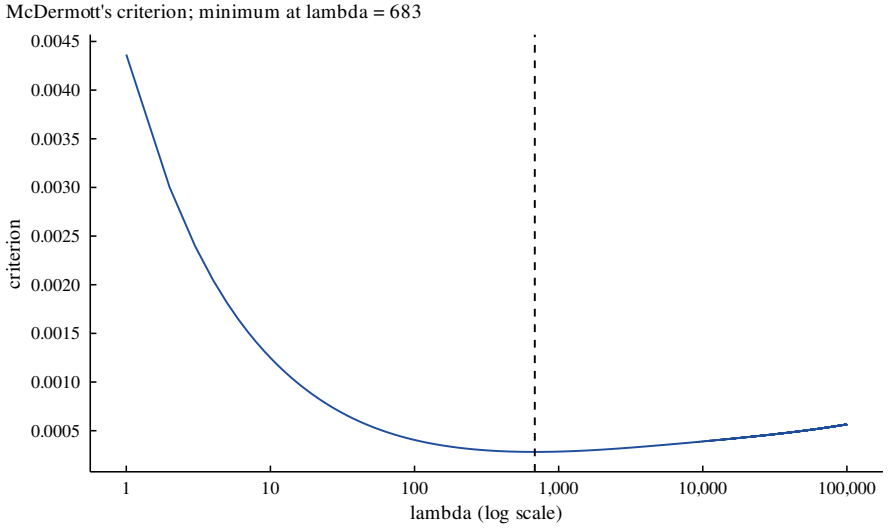
Cycle Statistics:
  Mean: -1.634e-13
  SD:   0.01333
  AR1:  0.828
  Range: 0.07802

Sources: Hodrick and Prescott (1997); McDermott (1997); Coe and McDermott
(1997); Choudhary et al. (2014); Choudhary et al. (2014, MATLAB code). Full
references: mhp_cite(x).
```

```
plot(x = mhp)
```




```
plot(x = mhp, which = "gcv")
```



3.3 The fully modified HP filter

The HP filter has too little information at the ends of a series, so its trend follows the last observations too closely and the end of the cycle is biased towards zero,

just where a policy maker needs it most. The fully modified HP filter of Hanif, Iqbal and Choudhary (2017) raises λ at the first and last observations, with weights chosen to make the filter at the ends behave like the filter in the middle, and chooses λ from the data as the modified HP filter does. The published code filters the logarithm of the series and returns the trend and the cycle in levels; `log_scale = TRUE` does the same.

```
fm <-
  fmhp_filter(
      x               = gdp
    , max_lambda      = 3000
    , k_linear        = 15
    , k_quadratic     = NULL
    , alpha_linear    = seq(from = 0, to = 2000, by = 10)
    , alpha_quadratic = seq(from = 0, to = 2000, by = 5)
    , schemes         = c("linear", "quadratic")
    , form            = "rows"
    , lambda          = NULL
    , log_scale       = TRUE
    , as_dt           = FALSE
    )
fm
```

```
Fully modified HP Filter Result
===============================
Observations: 203
Fully modified lambda: 1448  (modified HP lambda: 683 )
End-point weights: quadratic, alpha = 15 on 22 end values; loss = 0.6163
GCV: 3.336e-04
Filtered on the log scale; cycle = series - trend in levels

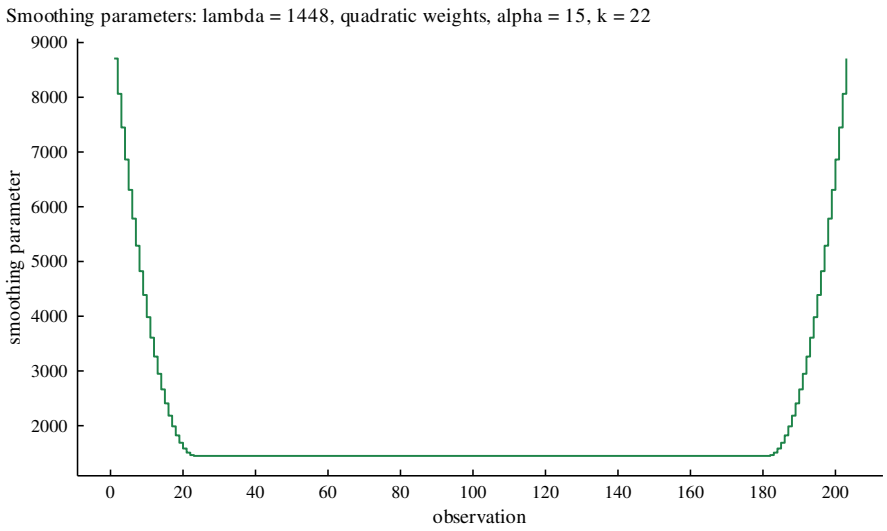
Cycle Statistics:
  Mean: -2.603
  SD:   120.4
  AR1:  0.9522
  Range: 900

Sources: Hanif et al. (2017); Hanif (2017); Bloechl (2014); McDermott (1997);
Hodrick and Prescott (1997). Full references: mhp_cite(x).
```

```
fm$schemes
```

```
      scheme  k alpha      loss lambda          gcv
1:    linear 15   240 0.8548936   1169 0.0003168500
2: quadratic 22    15 0.6163442   1448 0.0003335623
```

```
plot(x = fm, which = "lambda")
```



The smoothing parameter rises from 1448 in the middle to 8708 at the first and last observation (quadratic weights on 22 values at each end).

4 Comparing the filters

`fmhp_compare()` filters the series with the three filters and compares the cycles and their end-point loss, the measure of end-point bias used by Hanif et al. (2017): zero for a filter that treats the ends like the middle.

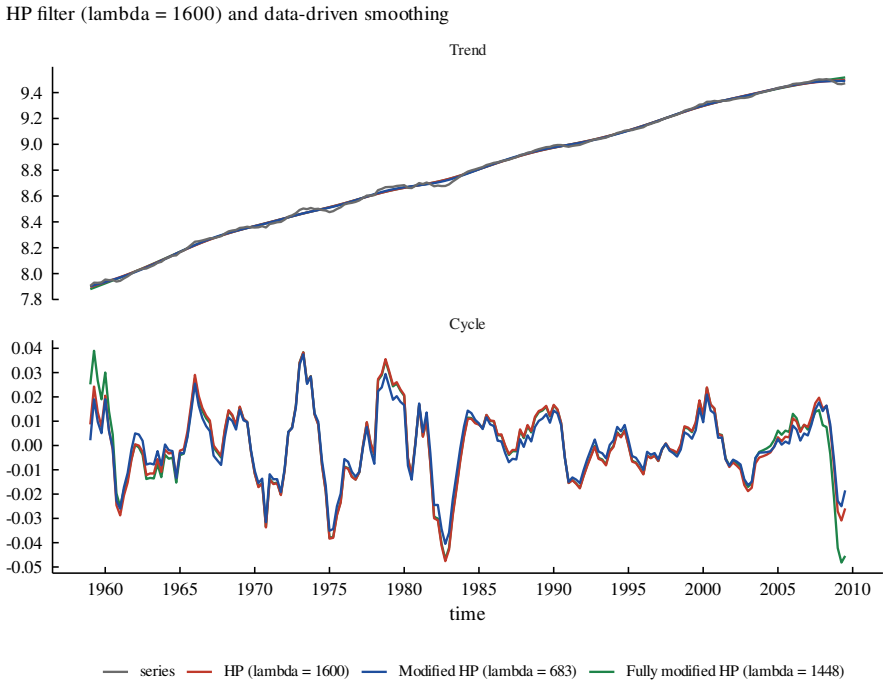
```
cmp <-
  fmhp_compare(
      x          = gdp
    , frequency  = "quarterly"
    , max_lambda = 3000
    , log_scale  = TRUE
    , end_share  = 0.1
    )
cmp
```

```
              method lambda   cycle_sd    cycle_mean       ar1 cycle_range
1:                HP   1600 0.01543904  4.942713e-12 0.8668207  0.08590516
2:       Modified HP    683 0.01333176 -1.634336e-13 0.8280266  0.07801575
3: Fully modified HP   1448 0.01618038 -4.145448e-05 0.8815215  0.08713978
            gcv      loss    sd_ends  sd_middle
1:           NA 1.7617778 0.01458541 0.01568213
2: 0.0002819926 1.7652928 0.01202396 0.01366805
```

```
3: 0.0003335623 0.7898641 0.01930994 0.01538061
```

The fully modified HP filter cuts the loss of the HP filter by more than half. Its cycle is also more variable at the ends than in the middle, unlike the cycles of the other two filters: the end-point bias of the HP filter shows as cycles that are too small at the ends.

```
plot_comparison(
    x          = log(gdp)
  , frequency  = "quarterly"
  , max_lambda = 100000
  , show_cycle = TRUE
  , fmhp       = TRUE
  )
```



5 Many series

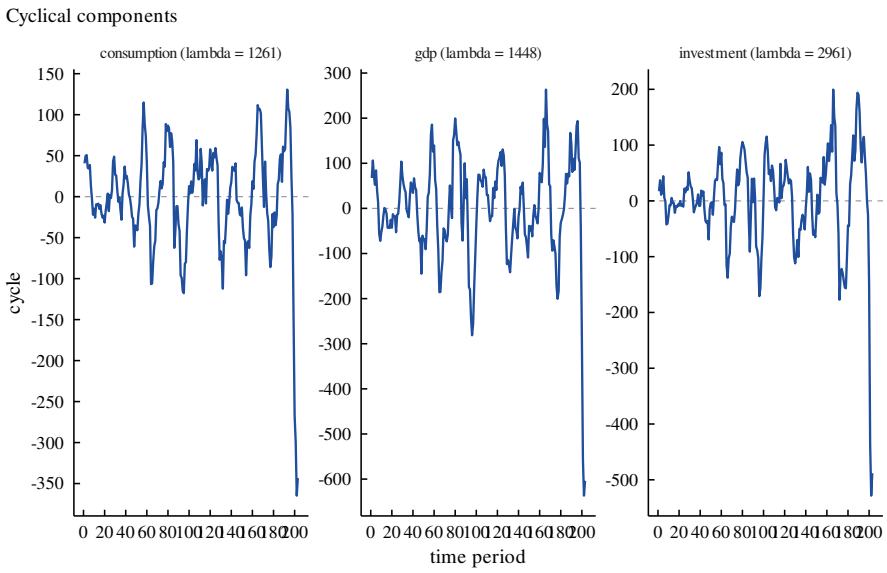
`mhp_batch()` and `fmhp_batch()` filter every column of a matrix; series of different lengths can share it, padded with missing values at the start or the end.

```
X <- as.matrix(us_macro[, c("gdp", "consumption", "investment")])
fb <-
  fmhp_batch(
```

```
      X               = X
    , max_lambda      = 3000
    , k_linear        = 15
    , k_quadratic     = NULL
    , alpha_linear    = seq(from = 0, to = 2000, by = 10)
    , alpha_quadratic = seq(from = 0, to = 2000, by = 5)
    , schemes         = c("linear", "quadratic")
    , form            = "rows"
    , lambda          = NULL
    , log_scale       = TRUE
    )
attr(x = fb, which = "lambdas")
```

```
        series lambda          gcv lambda_mhp    scheme alpha  k      loss
1:         gdp   1448 0.0003335623        683 quadratic    15 22 0.6163442
2: consumption   1261 0.0002248243        534 quadratic    15 22 0.5850792
3:  investment   2961 0.0073465762       1009 quadratic    20 22 0.6728679
```

```
plot_batch(
    x         = fb
  , show      = "cycle"
  , facet     = TRUE
  , highlight = NULL
  )
```



6 Reading the results

`mhp_interpret()` reads a result in plain language, from its numbers, and ends with the sources of the method.

```
mhp_interpret(
    x   = fm
  , top = 3
  )
```

```
Fully modified HP filter

The fully modified HP filter chose lambda = 1,448 (the modified HP filter
alone: 683) and quadratic end-point weights on 22 values at each end, which
cut the end-point loss of the HP filter with the same lambda by 65%.

Findings:
- Cycle (log scale, about the share of the trend): standard deviation 0.0162,
  first-order autocorrelation 0.88, range 0.0871.
- Largest positive cycle values (booms) at 2000Q2, 1978Q4, 1999Q4; largest
  negative ones (recessions) at 2009Q2, 2009Q3, 2009Q1.
- Linear weights: loss 0.855 (alpha = 240); quadratic weights: loss 0.616
  (alpha = 15); the quadratic scheme was kept.

Recommendations:
- Compare the filters on this series with fmhp_compare() (moments of the
  cycles and end-point loss).

Cautions:
- The last values of the cycle are revised as new observations arrive: read
  them with care.

Sources:
- Hanif, M. N., Iqbal, J. and Choudhary, M. A. (2017). Fully modified HP
  filter. SBP Working Paper Series, No. 88, State Bank of Pakistan.
- Hanif, M. N. (2017). Fully modified HP filter function (MATLAB code).
  MATLAB Central File Exchange, 63198.
- Bloechl, A. (2014). Reducing the excess variability of the Hodrick-Prescott
  filter by flexible penalization. Discussion Paper 2014-1, University of
  Munich.
- McDermott, C. J. (1997). An automatic method for choosing the smoothing
  parameter in the HP filter. Unpublished, International Monetary Fund,
  Washington DC.
- Hodrick, R. J. and Prescott, E. C. (1997). Postwar U.S. business cycles: an
  empirical investigation. Journal of Money, Credit and Banking, 29(1), 1-16.
  doi:10.2307/2953682
```

7 Where next

- Vignette 2: the filters, their criteria, gains and loss, and how they are computed.
- Vignette 3: the modified HP filter of Choudhary, Hanif and Iqbal (2014).
- Vignette 4: the fully modified HP filter of Hanif, Iqbal and Choudhary (2017), with the reproduction of its figures and tables.
- Vignette 5: the simulation study.
- Vignette 6: the application, `mhp_app()`.
- Vignette 7: checks of the code against the authors' MATLAB programs, the released version and dense computations.

8 References

Choudhary, M. A., Hanif, M. N. and Iqbal, J. (2014). On smoothing macroeconomic time series using the modified HP filter. *Applied Economics*, 46(19), 2205–2214.

Hanif, M. N., Iqbal, J. and Choudhary, M. A. (2017). Fully modified HP filter. *SBP Working Paper Series*, No. 88, State Bank of Pakistan.

Hodrick, R. J. and Prescott, E. C. (1997). Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit and Banking*, 29(1), 1–16.

McDermott, C. J. (1997). An automatic method for choosing the smoothing parameter in the HP filter. Unpublished, International Monetary Fund.