The modified HP filter

Muhammad Yaseen

2026-11-05

Table of contents

1 Abstract **1**

2 The idea **1**

3 One series **2**

4 Three US series **4**

5 Against the HP filter **5**

6 How stable is the choice? **6**

7 The simulation designs of the paper **7**

8 References **8**

1 Abstract

This vignette works through the modified HP filter of Choudhary, Hanif and Iqbal (2014): the smoothing parameter chosen series by series by the generalised cross-validation of McDermott (1997). It applies the filter to the US quarterly series of the package, sets the results beside those of the paper, and shows how the criterion behaves. The data differ from the paper's (FRED 1959Q1 to 2009Q3 here; IFS 1957Q1 to 2010Q1 there), so the values are compared, not reproduced.

2 The idea

The HP filter fixes λ by the frequency of the data: 100 for annual, 1600 for quarterly and 129600 for monthly series. Choudhary et al. (2014) argue that one value cannot suit every series and every economy, and choose λ for each

series as the minimiser of

$$\mathrm{GCV}(\lambda) = \frac{1}{T}\left(1 + \frac{2T}{\lambda}\right)\sum_{t=1}^{T}\left(x_t - \hat{g}_t(\lambda)\right)^2$$

over $\lambda = 1, 2, ...$ (their eq. 6). The vignette "Theory" derives this criterion and its exact form.

3 One series

```
data(us_macro)
fit <-
  mhp_filter(
      x          = us_macro$gdp
    , max_lambda = 100000
    , as_dt      = FALSE
    , log_scale  = TRUE
    , criterion  = "approximate"
    )
fit
```

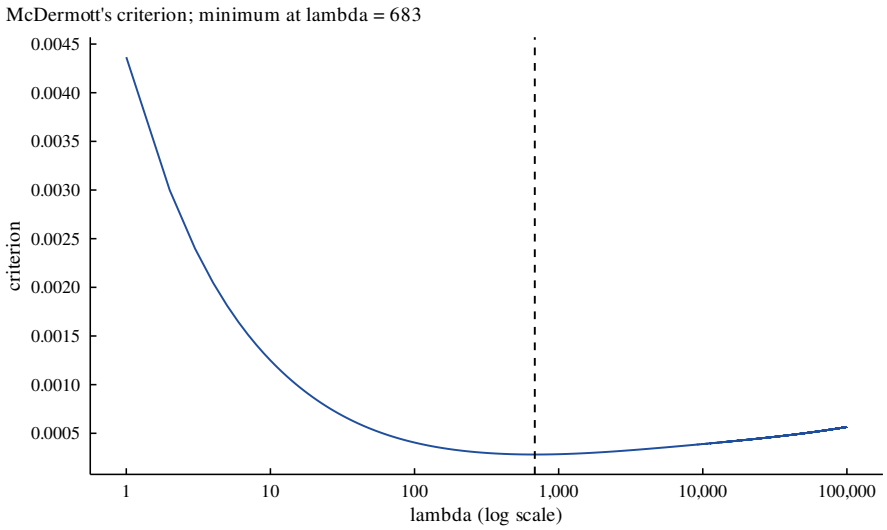
```
Modified HP Filter Result
==============================
Observations: 203
Optimal lambda: 683
GCV: 2.82e-04
Filtered on the log scale; cycle = series - trend in levels

Cycle Statistics:
  Mean: 1.044
  SD:   91.54
  AR1:  0.8554
  Range: 561.9

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).
```

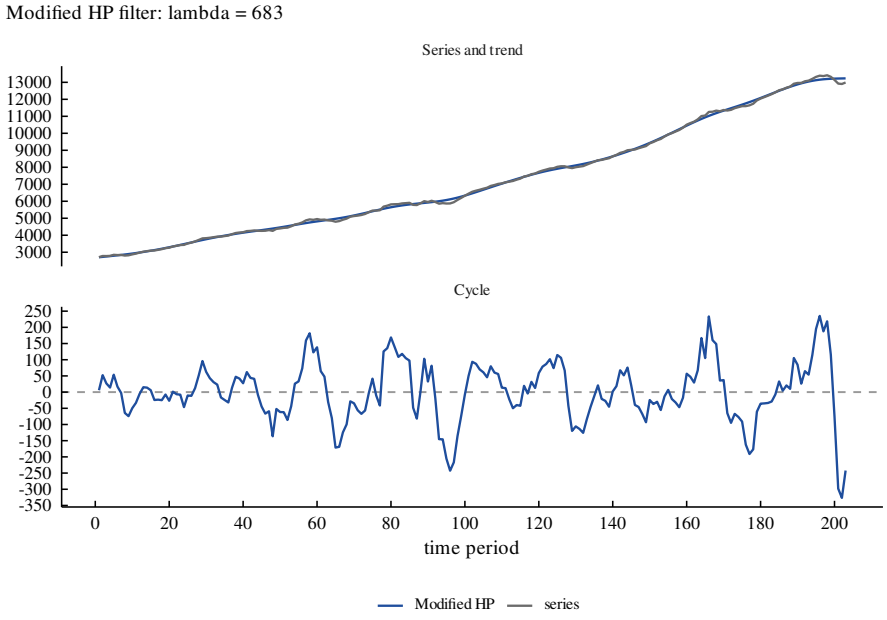
The criterion against λ has a single interior minimum:

```
ggplot2::autoplot(
    object = fit
  , which  = "gcv"
  )
```



and the trend and cycle:

```
ggplot2::autoplot(
    object = fit
  , which  = "decomposition"
  )
```



4 Three US series

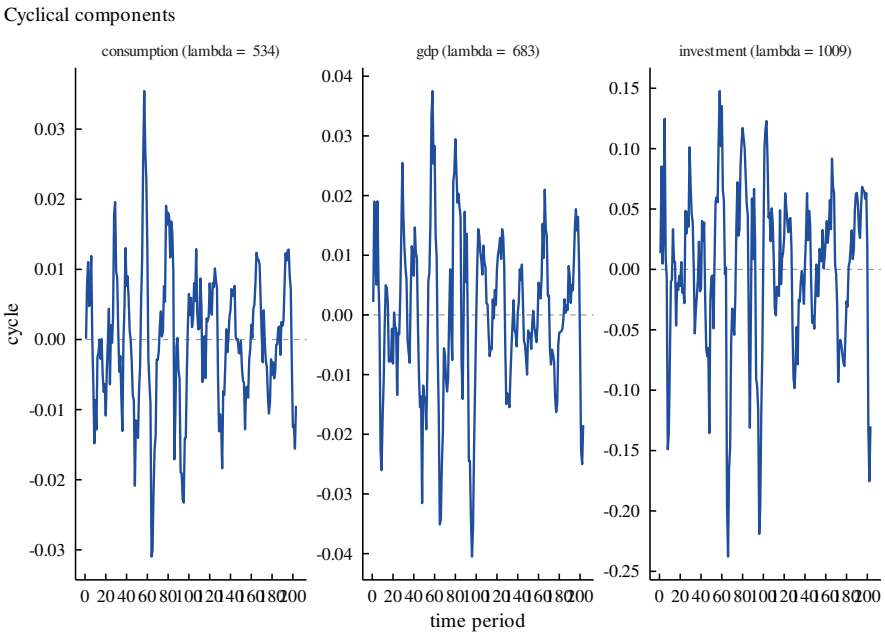
`mhp_batch()` filters the columns of a matrix, each with its own λ.

```
X <- log(as.matrix(us_macro[, c("gdp", "consumption", "investment")]))
b <-
  mhp_batch(
      X          = X
    , max_lambda = 100000
    )
lam <- attr(x = b, which = "lambdas")
lam[, paper := c(941, 841, 488)]
lam
```

```
        series lambda          gcv paper
1:         gdp    683 0.0002819926   941
2: consumption    534 0.0001753795   841
3:  investment   1009 0.0063778004   488
```

The column `paper` holds the values of Table 2 of Choudhary et al. (2014) for the United States. Ours are of the same order, below the 1600 of the HP filter, but differ series by series: the samples and the vintages of the data are not the same, and the criterion is flat near its minimum, so small changes in the data move λ a long way. Investment has the largest λ here and the smallest in the paper.

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



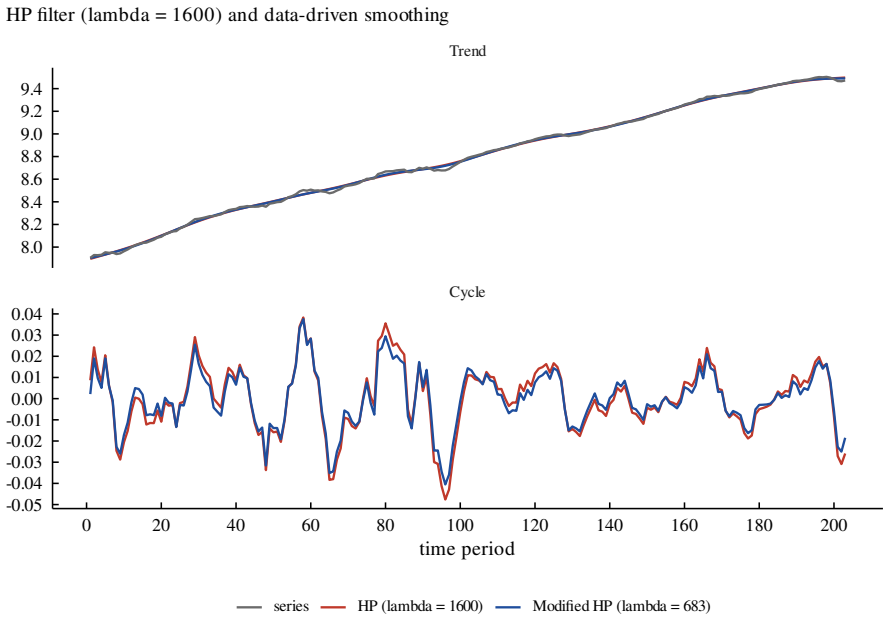
5 Against the HP filter

`mhp_compare()` sets the cycle of the HP filter with the conventional λ beside the cycle of the modified filter.

```
mhp_compare(
    x          = log(us_macro$gdp)
  , frequency  = "quarterly"
  , max_lambda = 100000
  )
```

```
        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
            gcv
1:           NA
2: 0.0002819926
```

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



6 How stable is the choice?

Because λ is estimated, it changes as data arrive. Refitting on samples that end in successive years shows how far:

```
ends <- seq(from = 120, to = nrow(us_macro), by = 8)
data.table::data.table(
    T      = ends
  , lambda = vapply(
        X          = ends
      , FUN        = function(n) {
          get_lambda(
              x = mhp_filter(
                  x          = us_macro$gdp[seq_len(n)]
                , max_lambda = 100000
                , as_dt      = FALSE
                , log_scale  = TRUE
                , criterion  = "approximate"
                )
            )
        }
      , FUN.VALUE = numeric(length = 1)
      )
  )
```

```
      T lambda
 1: 120    375
 2: 128    422
 3: 136    484
 4: 144    476
 5: 152    534
 6: 160    580
 7: 168    629
 8: 176    606
 9: 184    677
10: 192    707
11: 200    743
```

7 The simulation designs of the paper

Choudhary et al. (2014) compare the filters on 30 simulated designs: a random-walk trend with drift and an AR(2) cycle, with six pairs of AR coefficients and five ratios of the standard deviations of the two shocks. `mhp_designs(paper = "2014")` lists them and `mhp_simulate()` draws a series from one of them; the vignette "Simulation" runs the comparison.

```
mhp_designs(
    paper = "2014"
  )
```

```
    model sd_ratio ar1   ar2
 1:     1    10.00 0.9  0.01
 2:     2    10.00 1.2 -0.25
 3:     3    10.00 1.2 -0.40
 4:     4    10.00 1.2 -0.55
 5:     5    10.00 1.2 -0.75
 6:     6     5.00 0.9  0.01
 7:     7     5.00 1.2 -0.25
 8:     8     5.00 1.2 -0.40
 9:     9     5.00 1.2 -0.55
10:    10     5.00 1.2 -0.75
11:    11     1.00 0.9  0.01
12:    12     1.00 1.2 -0.25
13:    13     1.00 1.2 -0.40
14:    14     1.00 1.2 -0.55
15:    15     1.00 1.2 -0.75
16:    16     0.50 0.9  0.01
17:    17     0.50 1.2 -0.25
18:    18     0.50 1.2 -0.40
19:    19     0.50 1.2 -0.55
20:    20     0.50 1.2 -0.75
```

```
21:    21     0.01 0.9  0.01
22:    22     0.01 1.2 -0.25
23:    23     0.01 1.2 -0.40
24:    24     0.01 1.2 -0.55
25:    25     0.01 1.2 -0.75
26:    26    10.00 0.8  0.00
27:    27     5.00 0.8  0.00
28:    28     1.00 0.8  0.00
29:    29     0.50 0.8  0.00
30:    30     0.01 0.8  0.00
    model sd_ratio ar1   ar2
```

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.

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 manuscript, International Monetary Fund.