

Package: epo (via r-universe)

September 27, 2026

Type Package

Title Enhanced Portfolio Optimization (EPO)

Version 0.1.0.9000

Maintainer Bernardo Reckziegel <bernardo_cse@hotmail.com>

Description Implements the Enhanced Portfolio Optimization (EPO) method of Pedersen, Babu and Levine (2021) <doi:10.1080/0015198X.2020.1854543>. EPO stabilizes mean-variance optimization by shrinking correlations toward zero, which addresses estimation error in both the risk model and expected returns. Both the ``Simple EPO" and the ``Anchored EPO" (an anchor/benchmark-aware, Black-Litterman-like extension) are provided.

License MIT + file LICENSE

URL <https://github.com/Reckziegel/epo>,
<https://reckziegel.github.io/epo/>

BugReports <https://github.com/Reckziegel/epo/issues>

Encoding UTF-8

LazyData true

Imports assertthat (>= 0.2.1), dplyr (>= 1.2.1), rlang (>= 1.3.0)

Roxygen list(markdown = TRUE)

Suggests testthat (>= 3.3.2), xts (>= 0.14.2)

Config/testthat/edition 3

Config/roxygen2/version 8.0.0

Repository <https://reckziegel.r-universe.dev>

Date/Publication 2026-07-31 16:10:27 UTC

RemoteUrl <https://github.com/reckziegel/epo>

RemoteRef HEAD

RemoteSha 3ef7d092e398f9273d02b81f108871e2da7d7071

Contents

epo	2
Index	6

epo	<i>Enhanced Portfolio Optimization (EPO)</i>
-----	--

Description

Computes the optimal portfolio allocation using the Enhanced Portfolio Optimization (EPO) method of Pedersen, Babu, and Levine (2021).

Usage

```
epo(
  x,
  signal,
  lambda,
  method = c("simple", "anchored"),
  w,
  anchor = NULL,
  normalize = TRUE,
  endogenous = TRUE
)
```

Default S3 method:

```
epo(
  x,
  signal,
  lambda,
  method = c("simple", "anchored"),
  w,
  anchor = NULL,
  normalize = TRUE,
  endogenous = TRUE
)
```

S3 method for class 'tbl'

```
epo(
  x,
  signal,
  lambda,
  method = c("simple", "anchored"),
  w,
  anchor = NULL,
  normalize = TRUE,
```

```

    endogenous = TRUE
  )

## S3 method for class 'xts'
epo(
  x,
  signal,
  lambda,
  method = c("simple", "anchored"),
  w,
  anchor = NULL,
  normalize = TRUE,
  endogenous = TRUE
)

## S3 method for class 'matrix'
epo(
  x,
  signal,
  lambda,
  method = c("simple", "anchored"),
  w,
  anchor = NULL,
  normalize = TRUE,
  endogenous = TRUE
)

```

Arguments

<code>x</code>	A data-set with asset returns. It should be a tibble, a xts or a matrix.
<code>signal</code>	A double vector with the investor's beliefs about expected returns (signals, forecasts) for each asset in <code>x</code> .
<code>lambda</code>	A double with the investor's (absolute) risk-aversion coefficient, as in the paper's notation. For <code>method = "simple"</code> , the resulting portfolio's Sharpe ratio does not depend on <code>lambda</code> , so any positive value works when <code>normalize = TRUE</code> . For <code>method = "anchored"</code> with <code>endogenous = TRUE</code> , this argument is ignored because the risk-aversion coefficient is calibrated internally.
<code>method</code>	A character. One of: "simple" or "anchored".
<code>w</code>	A double between 0 and 1. The EPO shrinkage parameter: 0 yields standard mean-variance optimization (no shrinkage) and 1 yields maximum shrinkage (the anchor portfolio, for <code>method = "anchored"</code> , or an unoptimized portfolio, for <code>method = "simple"</code>). In practice, <code>w</code> is often chosen empirically, e.g. by picking the value that would have maximized the realized Sharpe ratio using only past (out-of-sample) data.
<code>anchor</code>	A double vector with the anchor (benchmark) portfolio that the allocation should not deviate too much from (e.g. a strategic asset allocation, a market-cap benchmark, or the 1/N portfolio). Only used when <code>method = "anchored"</code> .

normalize	A boolean indicating whether the allocation should be normalized to sum 1 (full-investment constraint). The default is normalize = TRUE.
endogenous	A boolean indicating whether the risk-aversion parameter should be calibrated endogenously from the anchor and signal (paper's footnote 13), rather than taken from lambda. Only used when method = "anchored". The default is endogenous = TRUE.

Details

Standard mean-variance optimization (MVO) is highly sensitive to estimation error in the correlation matrix and in expected returns. This error is concentrated in the least important principal components of the correlation matrix (the "problem portfolios"), whose risk tends to be underestimated and whose expected return tends to be overestimated. EPO fixes this by shrinking the off-diagonal correlations toward zero by a factor w before running MVO, which increases the estimated volatility (and lowers the implied Sharpe ratio) of exactly the problem portfolios.

Two flavors of EPO are implemented, both governed by a single shrinkage parameter, w , between 0 (no shrinkage, i.e. standard MVO) and 1 (maximum shrinkage):

- method = "simple" implements the "Simple EPO" (paper's equation 16). The allocation is given by $x = \frac{1}{\lambda} \Sigma_w^{-1} s$, where Σ_w is the variance-covariance matrix rebuilt from the shrunk correlation matrix $\Omega_w = (1 - w)\Omega + wI$. At $w = 1$ all correlations are set to zero, which is equivalent (up to scaling) to not optimizing at all.
- method = "anchored" implements the "Anchored EPO" (paper's equation 17), which pulls the solution toward a reference/benchmark portfolio, the anchor. At $w = 0$ the solution is standard MVO; at $w = 1$ the solution collapses onto the anchor; values in between produce Black-Litterman-style portfolios in which w controls the confidence placed in the anchor relative to the signal. Unlike Black-Litterman, the anchor need not be the market portfolio.

Value

A numeric vector with the optimal portfolio weights, one per column of x .

References

Pedersen, L. H., Babu, A., and Levine, A. (2021). Enhanced Portfolio Optimization. *Financial Analysts Journal*, 77(2), 124-151. doi:10.1080/0015198X.2020.1854543

Examples

```
x <- diff(log(EuStockMarkets)) # stock returns
s <- colMeans(x) # it could be any signal

#####
### Simple EPO ###
#####

# Traditional Mean-Variance Analysis
epo(x = x, signal = s, lambda = 10, method = "simple", w = 0)

# 100% Shrinkage
```

```
epo(x = x, signal = s, lambda = 10, method = "simple", w = 1)

# 50% Classical MVO and 50% Shrinkage
epo(x = x, signal = s, lambda = 10, method = "simple", w = 0.5)

#####
### Anchored EPO ###
#####

benchmark <- rep(0.25, 4) # 1/N Portfolio

# Traditional Mean-Variance Analysis
epo(x = x, signal = s, lambda = 10, method = "anchored", w = 0.0, anchor = benchmark)

# 100% on the Anchor portfolio
epo(x = x, signal = s, lambda = 10, method = "anchored", w = 1.0, anchor = benchmark)

# Somewhere between the two worlds
epo(x = x, signal = s, lambda = 10, method = "anchored", w = 0.5, anchor = benchmark)
```

Index

epo, [2](#)