

Risk and Performance Estimators Standard Errors (RPESE)

Anthony Christidis, Doug Martin, Xin Chen

October 28, 2019

Abstract

The Risk and Performance Estimators Standard Errors package (**RPESE**) implements a new method for computing accurate standard errors of risk and performance estimators when returns are serially correlated as when they are uncorrelated. The new method makes use of the representation of a risk or performance estimator as a summation of a time series of influence-function (IF) transformed returns, and computes estimator standard errors using a sophisticated method of estimating the spectral density at frequency zero of the time series of IF-transformed returns. The initial release of **RPESE** allows users to compute accurate standard errors for six risk estimators, including the standard deviation, semi-standard deviation, value-at risk and expected shortfall, and eight performance estimators, including the Sharpe ratio, Sortino ratio, and expected shortfall ratio. This vignette provides basic instruction on how to use the **RPESE** package.

1 Introduction

The current finance industry practice in reporting risk and performance estimates for individual assets and portfolios seldom includes reporting estimate standard errors (SEs). For this reason, consumers of such reports have no way of knowing the statistical accuracy of the estimates. As a leading example, one seldom sees SE's reported for Sharpe ratios, and consequently cannot tell whether or not two Sharpe ratios for two different portfolio products are significantly different.

This motivated us to develop a Risk and Performance Estimator Standard Errors (**RPESE**) package for computing risk and performance estimator SE's that are accurate: (1) when

returns are serially correlated as well as when they are uncorrelated, and (2) account for fat-tailed and skewed non-normality of returns distributions. **RPESE** uses a new method for computing risk and performance estimators standard errors due to [Chen and Martin \[2018\]](#), henceforth (CM).

RPESE supports computing SEs for the six risk estimators shown in Table 1, and the eight performance estimators shown in Table 2. For each of the names in the Name column of the two tables, there is a corresponding R function with a slightly different name in **RPESE**, except for *LMP1* and *LPM2* in Table 1 there is a single function with an optional argument to choose between these two risk estimators, and for *SorR. μ* and *SorR.c* in Table 2 there is a single function with an optional argument to choose between these two performance estimators.

Name	Estimator Description
<i>SD</i>	Sample standard deviation
<i>SemiSD</i>	Semi-standard deviation
<i>LPM1</i>	Lower partial moment of order 1
<i>LPM2</i>	Lower partial moment of order 2
<i>ES</i>	Expected shortfall with tail probability α
<i>VaR</i>	Value-at-risk with tail probability α

Table 1: Risk Estimator Names and Descriptions

Name	Estimator Description
<i>Mean</i>	Sample mean
<i>SR</i>	Sharpe ratio
<i>SorR.μ</i>	Sortino ratio with threshold the mean
<i>SorR.c</i>	Sortino ratio with threshold a constant c
<i>ESratio</i>	Mean excess return to ES ratio with tail probability α
<i>VaRratio</i>	Mean excess return to VaR ratio with tail probability α
<i>RachRatio</i>	Rachev ratio with lower upper tail probabilities α and β
<i>OmegaRatio</i>	Omega ratio with threshold c

Table 2: Performance Estimator Names and Descriptions

The `RPESE` package uses the `PerformanceAnalytics` (PA) package to compute the point estimators in Tables 1 and 2, and then computes the point estimate standard errors using the new method, which we now briefly describe.

The New Method

The basic elements of the new method are as follows. For a given risk or performance estimator, the time series of returns used to compute the estimate are transformed using the *influence function* (IF) of the estimator. For an introduction to influence functions, and derivations of the influence functions of the estimators in Tables 1 and 2, see **Zhang, Martin and Christidis (2019)**. It turns out that risk and performance estimators can be represented as the sample mean of the time series of influence-function transformed returns. It is well-known that the standardized (with respect to sample size) sum of a stationary time series has a variance that is well-approximated by the spectral density of the time series at zero frequency. Thus, computing the standard error of a risk or performance estimator reduces to estimating the spectral density at zero frequency of a standardized sum of the influence-function transformed returns. CM developed an effective method of doing so based on first computing the periodogram of the influence-function transformed returns, and then using a regularized general linear model (GLM) method for exponential and gamma distributions to fit a polynomial to the periodogram values. The regularization method used is an elastic net (EN) method that encourages sparsity of coefficients and is well-known in the machine learning literature. The intercept of such GLM fitting provides the an estimate of the spectral

density at zero frequency, and hence a risk or performance estimator standard error. The interested reader can find the details in the CM paper.

2 RPESE Component Packages

The overall structure of RPESE, depicted in Figure 1, shows that RPESE makes use of the following two new packages:

- RPEIF (Influence Functions package)
- RPEGLMEN (Generalized Linear Model fitting with Elastic Net, for exponential and gamma distributions)

The purpose of RPEIF is to provide the analytic formula of the influence function, in order to compute the IF-transformed returns for each of the risk and performance estimators. For each risk and performance estimator in Tables 1 and 2, the RPEGLMEN package fits an EN regularized GLM polynomial fit to the periodogram of the time series of IF-transformed returns, using a GLM for exponential distributions or Gamma distributions. Figure 1 shows the relationship between the above two packages and the overall RPESE package.

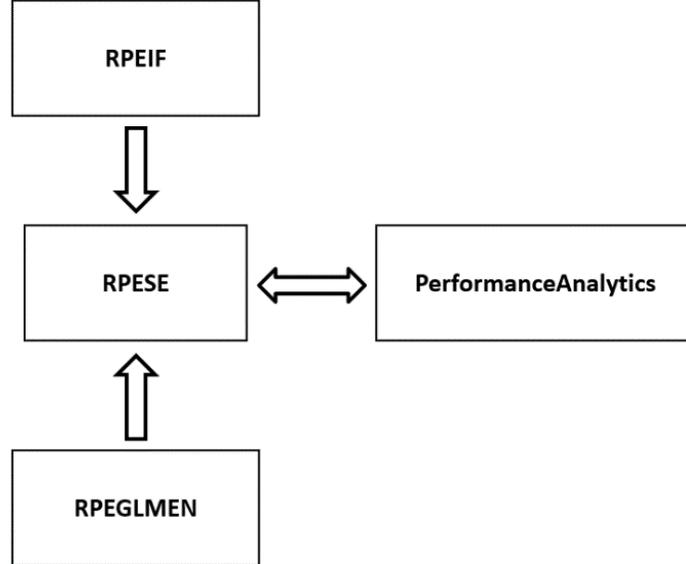


Figure 1: Packages Relations between RPEIF, RPEGLMEN, RPESE and PerformanceAnalytics

3 How to Use RPESE

In the following sections, we show how to use RPESE to compute standard errors of risk and performance measures using the `edhec` hedge fund data.

3.1 Installing and Loading RPESE and Loading an Examples Data Set

In order to use RPESE, you don't need to manually install any other of the dependent packages as they will be installed automatically when RPESE is installed.

```
install.packages("RPESE")
```

To load RPESE, use the code:

```
library(RPESE)
```

We will use from the library `PerformanceAnalytics` the data set `edhec` to demonstrate the functionality of RPESE. The `PerformanceAnalytics` `xts` data set `edhec`, consisting of the returns of a number of hedge funds January, 1997 to August, 2008, is used in the examples in this vignette. Load `edhec`, confirm the object's class and see the names of the hedge funds with the following code.

```
data(edhec, package='PerformanceAnalytics')
class(edhec)

## [1] "xts" "zoo"

names(edhec)

## [1] "Convertible Arbitrage" "CTA Global"
## [3] "Distressed Securities" "Emerging Markets"
## [5] "Equity Market Neutral" "Event Driven"
## [7] "Fixed Income Arbitrage" "Global Macro"
## [9] "Long/Short Equity" "Merger Arbitrage"
## [11] "Relative Value" "Short Selling"
## [13] "Funds of Funds"
```

Since the hedge fund names are too long for convenient display, it will be best to use shorter two or three letter labels, with the following code:

```
names(edhec) <- c("CA", "CTA", "DIS", "EM", "EMN", "ED", "FIA",
                 "GM", "LS", "MA", "RV", "SS", "FOF")
```

3.2 Functions in RPESE

To see what functions are contained in the RPESE package, use the code line:

```
ls("package:RPESE")

## [1] "ES.SE"           "ESratio.SE"      "EstimatorSE"
## [4] "LPM.SE"         "Mean.SE"         "OmegaRatio.SE"
## [7] "printSE"        "RachevRatio.SE" "SemiSD.SE"
## [10] "SharpeRatio.SE" "SortinoRatio.SE" "StdDev.SE"
## [13] "VaR.SE"         "VaRratio.SE"
```

3.3 Basic Functionality

To compute the standard errors of risk and performance measure estimators listed in Table 1, the argument `SE` must be set to `TRUE`. Also there is the option to return any of the standard error computation methods available in RPESE, namely `IFiid`, `IFcor`, `BOOTiid` and `BOOTcor`. See the documentation of the RPESE package for more details.

As a basic example, let's consider the Sharpe Ratio estimate and its computed standard error for the Convertible Arbitrage hedge fund. We print the output using the `printSE` function from RPESE.

```
# Sharpe Ratio SE computation for single hedge fund
Sharpe.out <- SharpeRatio.SE(edhec[,"CA"],
                             se.method=c("IFiid","IFcor","IFcorAdapt",
                                           "BOOTiid","BOOTcor", "none"))

# Print output
printSE(Sharpe.out)
```

```
##      SR IFiid IFcor IFcorAdapt BOOTiid BOOTcor
## CA 0.32 0.123 0.155      0.127   0.125   0.201
```

The risk and performance measures functions in allow the user to return the standard errors for more than one asset or portfolio at the same time.

```
# Sharpe Ratio SE computation for all hedge funds in data set
Sharpe.out <- SharpeRatio.SE(edhec,
                             se.method=c("IFiid","IFcor","IFcorAdapt",
                                           "BOOTiid","BOOTcor", "none"))

# Print output
printSE(Sharpe.out)
```

```
##      SR IFiid IFcor IFcorAdapt BOOTiid BOOTcor
## CA 0.320 0.123 0.155      0.127   0.128   0.191
## CTA 0.258 0.081 0.080      0.080   0.083   0.043
## DIS 0.433 0.118 0.145      0.119   0.121   0.183
## EM 0.214 0.094 0.119      0.116   0.100   0.114
## EMN 0.667 0.180 0.184      0.184   0.201   0.309
## ED 0.415 0.116 0.137      0.129   0.109   0.127
## FIA 0.299 0.130 0.158      0.137   0.151   0.255
## GM 0.451 0.073 0.081      0.081   0.065   0.051
## LS 0.350 0.090 0.114      0.114   0.086   0.123
## MA 0.608 0.133 0.144      0.143   0.130   0.130
## RV 0.508 0.135 0.166      0.144   0.133   0.187
## SS 0.076 0.079 0.093      0.093   0.074   0.059
## FOF 0.325 0.092 0.119      0.106   0.089   0.119
```

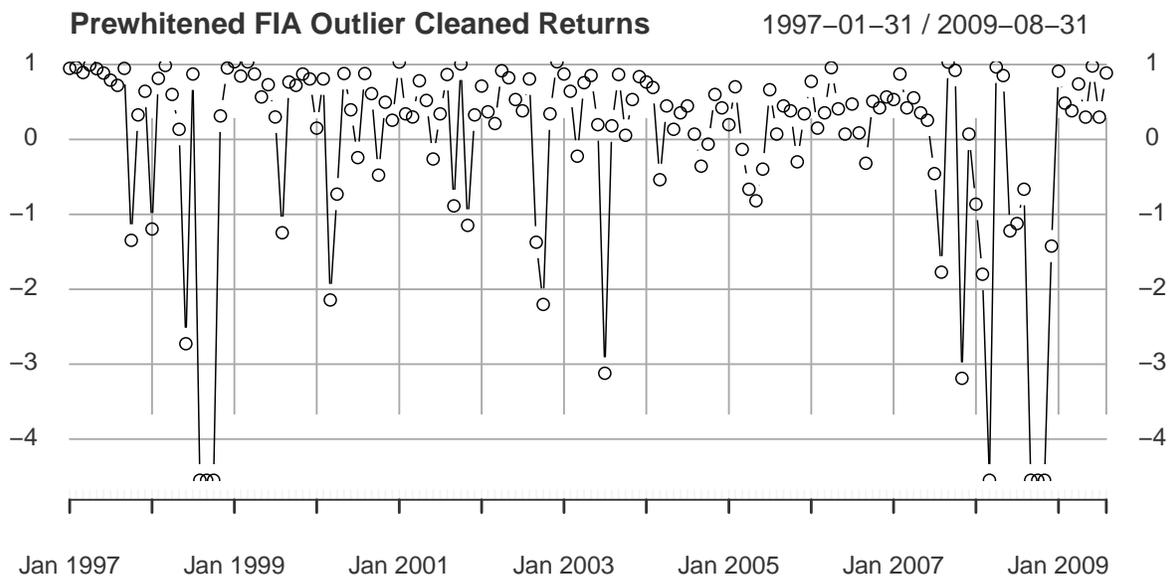
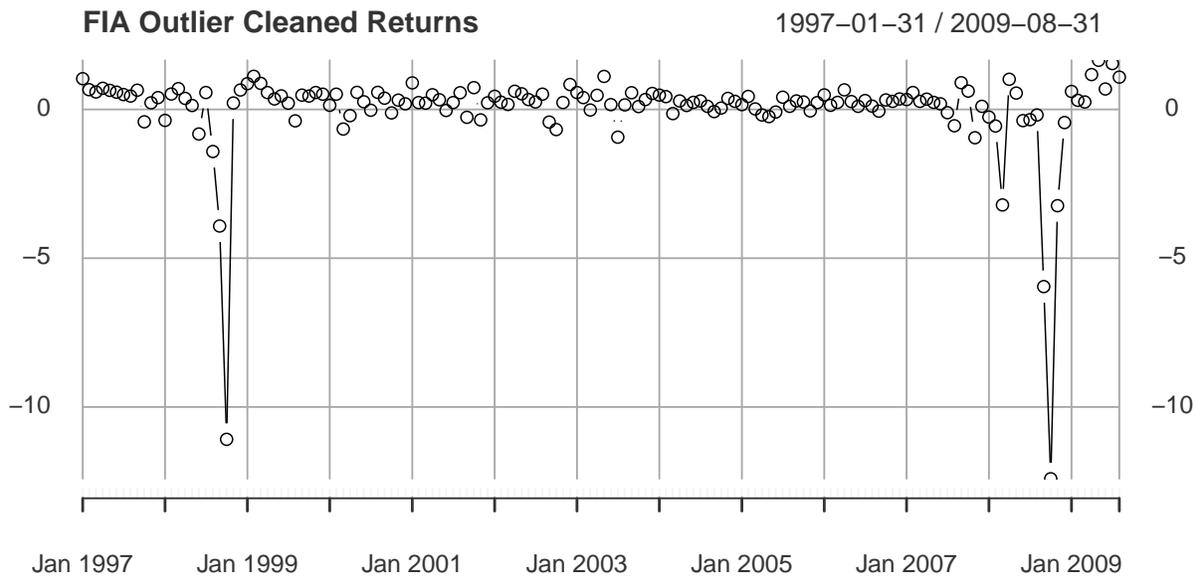
We may look at the documentation of the function for the Sharpe Ratio as well via the following code:

```
# Documentation of the ES.SE function
?SharpeRatio.SE
```

3.4 Outlier Cleaning

There is also the outlier cleaning capability from RPEIF that gets passed on to RPESE. First, notice the difference in the IF TS with and without outlier cleaning.

```
library(RPEIF)
IFout <- IF.SR(returns = edhec[, "FIA"], cleanOutliers = F, IFprint = T)
IFout.clean <- IF.SR(returns = edhec[, "FIA"], cleanOutliers = T, IFprint = T)
par(mfrow = c(2,1))
plot(IFout,type = "b", main = "FIA Outlier Cleaned Returns", lwd = .8)
plot(IFout.clean,type = "b", main = "Prewhitened FIA Outlier Cleaned Returns",
      lwd = .8)
```



```
par(mfrow = c(1,1))
```

We apply outlier cleaning to the computation of the Sharpe Ratio standard errors estimates.

```

# Sharpe Ratio SE computation for all hedge funds in data set
# Use outlier cleaning
clean.out <- SharpeRatio.SE(edhec,
                            se.method=c("IFiid","IFcor","IFcorAdapt",
                                          "BOOTiid","BOOTcor", "none"),
                            cleanOutliers = T)

# Print output
clean.comparison <- cbind(printSE(Sharpe.out)[,4], printSE(clean.out)[,4])
colnames(clean.comparison) <- c("IFcorAdapt", "IFcorAdapt-Clean")
clean.comparison

##           IFcorAdapt IFcorAdapt-Clean
## [1,]           0.127             0.092
## [2,]           0.080             0.080
## [3,]           0.119             0.105
## [4,]           0.116             0.118
## [5,]           0.184             0.137
## [6,]           0.129             0.121
## [7,]           0.137             0.117
## [8,]           0.081             0.083
## [9,]           0.114             0.113
## [10,]          0.143             0.137
## [11,]          0.144             0.117
## [12,]          0.093             0.095
## [13,]          0.106             0.111

```

3.5 Prewhitening

Another option in RPEIF is the availability to prewhiten the influence functions time series. This option is passed onto the RPESE package.

```

# Sharpe Ratio SE computation for all hedge funds in data set
# Use prewhitening
cleanPW.out <- SharpeRatio.SE(edhec,
                              se.method=c("IFiid","IFcor","IFcorAdapt",
                                            "BOOTiid","BOOTcor", "none"),
                              cleanOutliers = T)

```

```

                                cleanOutliers = T, prewhiten = T)
# Print output
cleanPW.comparison <- cbind(printSE(clean.out)[,3],
                             printSE(cleanPW.out)[,3])
colnames(cleanPW.comparison) <- c("IFcor", "IFcor-CleanPW")
rownames(cleanPW.comparison) <- names(edhec)
cleanPW.comparison

##      IFcor IFcor-CleanPW
## CA  0.132      0.079
## CTA 0.080      0.080
## DIS 0.143      0.080
## EM   0.087      0.081
## EMN 0.137      0.098
## ED   0.133      0.088
## FIA 0.140      0.095
## GM   0.083      0.076
## LS   0.113      0.084
## MA   0.138      0.109
## RV   0.137      0.093
## SS   0.095      0.080
## FOF 0.116      0.079

```

If `se.method` is set to `IFcorAdapt`, then the `prewhiten` option has no effect as the standard error estimate is a weighted average between `IFcor` with and without prewhitening. See Section 4.2 in [Chen and Martin \[2018\]](#) for more details.

3.6 Exponential and Gamma Distributions

The `RPEGLMEN` package used in the fitting of the penalized GLM model to the periodogram of the influence functions time series allows for two different distributions: the exponential distribution (default) and the Gamma distribution.

```

# Sharpe Ratio SE computation for all hedge funds in data set
# Use Gamma distribution
Gamma.out <- SharpeRatio.SE(edhec,
                             se.method=c("IFiid", "IFcor", "IFcorAdapt",

```

```

                                "BOOTiid","BOOTcor", "none"),
                                cleanOutliers = T, fitting.method = "Gamma")
# Print output
Gamma.comparison <- cbind(printSE(clean.out)[,4],
                           printSE(Gamma.out)[,4])
colnames(Gamma.comparison) <- c("IFcorAdapt", "IFcorAdapt-Gamma")
rownames(Gamma.comparison) <- names(edhec)
Gamma.comparison

##      IFcorAdapt IFcorAdapt-Gamma
## CA          0.092          0.108
## CTA         0.080          0.050
## DIS         0.105          0.145
## EM          0.118          0.140
## EMN         0.137          0.202
## ED          0.121          0.150
## FIA         0.117          0.133
## GM          0.083          0.083
## LS          0.113          0.107
## MA          0.137          0.093
## RV          0.117          0.163
## SS          0.095          0.094
## FOF         0.111          0.144

```

References

X. Chen and R. D. Martin. Standard errors of risk and performance measure estimators for serially correlated returns. 2018. URL <https://ssrn.com/abstract=3085672>.