

Risk and Performance Estimators Standard Errors (RPESE)

Anthony Christidis, Doug Martin, Xin Chen

October 31, 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 well 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 **RPESE** package 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 similar 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 some of the point estimators in Tables 1 and 2, and computes the others internally. The main function of **RPESE** is to compute the standard errors of the point estimates listed in Tables 1 and 2, using the new method that 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 for risk and performance estimators, and derivations of the influence functions of the estimators in Tables 1 and 2, see [Zhang et al. \[2019\]](#). 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 an appropriately standardized (with respect to sample size) sum of a stationary time series has a variance that is approximated by the spectral density of the time series at zero frequency, with the approximation becoming exact as the sample size tends to infinity. 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 of risk and performance estimators)
- **RPEGLMEN** (Generalized Linear Model fitting with Elastic Net, for exponential and gamma distributions)

The purpose of **RPEIF** is to provide the analytic formulas of influence functions in support of computing the IF transformed returns for 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 the functions in **RPESE** to compute standard errors of risk and performance estimators using time series of monthly hedge fund returns contained in the **PerformanceAnalytics** package.

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 dependent packages as they will be installed automatically when **RPESE** is installed. You can install **RPESE** from CRAN as follows:

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

To load **RPESE**, use the code line:

```
library(RPESE)
```

We will use the `xts` data set `edhec` of hedge fund returns, contained in `PerformanceAnalytics`, in demonstrating the functionality of `RPESE`. The following code loads the `edhec` data, confirms the object's class, lists the names of the hedge funds, and displays the range of dates of the data.

```
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"

library(xts) # Need this for the next line and later use of plot.zoo
range(index(edhec))

## [1] "1997-01-31" "2009-08-31"
```

Since the hedge fund names are too long for convenient display, the following code is used to create shorter two or three letter names:

```
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"
```

The twelve functions whose function name ends in `.SE` compute the standard error of the estimator corresponding to the first part of the function name. Note that there are only five risk estimators instead of the six in Table 1, which is because `LPM.SE` uses an optional argument to choose computing *LPM1* or *LPM2*, and similarly there are only seven performance estimator instead of the eight in Table 2 because `SortinoRatio.SE` uses an optional argument to choose computing *SoRc* or *SoRμ*.

The function `printSE` is a utility function whose use is demonstrated below, and the function `estimatorSE` is another utility function whose use is illustrated in an example in the function's help file.

3.3 Basic Functionality

The arguments of the twelve standard error computation functions are all similar, and are illustrated below for the cases of the `StdDev.SE` and `SharpeRatio.SE` functions using the `args` function:

```
args(StdDev.SE)
```

```
## function (data, se.method = c("IFiid", "IFcor", "IFcorAdapt",
##      "IFcorPW", "BOOTiid", "BOOTcor")[1:2], cleanOutliers = FALSE,
##      fitting.method = c("Exponential", "Gamma")[1], ...)
## NULL
```

```
args(SharpeRatio.SE)
```

```
## function (data, Rf = 0, se.method = c("IFiid", "IFcor", "IFcorAdapt",
##      "IFcorPW", "BOOTiid", "BOOTcor")[c(1, 3)], cleanOutliers = FALSE,
```

```
##      fitting.method = c("Exponential", "Gamma")[1], ...)  
## NULL
```

The only argument that is required for the standard error computing functions is the `data` argument, and if only the data argument is supplied then the function uses the defaults of the other arguments. However the `se.method =` argument is particularly important, and the standard error computation methods for the various choices for this argument are as follows:

- "IFiid": This results in an influence function (IF) method based computation of a standard error assuming i.i.d. returns
- "IFcor": This is the basic IF method computation of a standard error that takes into account serial correlation in the returns
- "IFcorAdapt": This IF based method adaptively interpolates between IFcor and IFcorPW to better account for serial correlation in the returns than with either IFcor or IFcorPW alone
- "IFcorPW": This IF based method uses pre-whitening of the IF transformed returns and is useful when serial correlation is large
- "BOOTiid": This choice results in computing a bootstrap standard error assuming i.i.d. returns
- "BOOTcor": This choice uses a block bootstrap method to compute a standard error that takes into account serial correlation of returns.

Our two default choices of methods are:

- "IFiid" and "IFcor" for risk estimators, and for performance estimators when returns serial correlation are known to be small
- "IFiid" and "IFcorAdapt" for performance estimators when returns correlations are unknown and may be large

Note how these choices are made in the `method =` arguments for the `StdDev.SE` and `SharpeRatio.SE` functions above.

The value of including `IFiid`, along with `IFcor` and `IFcorAdapt` is that it allows the user to see whether or not serial correlation results in a difference in the standard error that assumes

i.i.d. returns and the standard error that takes into account serial correlation. If there is no serial correlation there will not be much difference, but if there is serial correlation the difference can be considerable.

The `BOOTiid` and `BOOTcor` methods are provided for users who want to see how these bootstrap methods of computing standard errors compare with the IF based methods. Our experience to date indicates that `BOOTiid` generally agrees quite well with `IFiid`, but that `BOOTcor` is not as consistent in giving values similar to those of `IFcor`.

As a simple example of using an SE function, the following code computes the standard error (SE) of the standard deviation estimator for the convertible arbitrage (CA) hedge fund, using the `IFiid` and `IFcor` methods:

```
# Standard deviation SE computation for a single hedge fund
StdDev.SE(edhec[, "CA"])

## $SD
##           [,1]
## StdDev 0.02004739
##
## $IFiid
## [1] 0.003443619
##
## $IFcor
## [1] 0.006149056
```

In this case the `IFcor` standard error is considerable 76% larger than that of the `IFiid` standard error.

We note that the result returned by an SE function is a list, and so the result is printed using the default print method for a list. A more compact display of the results, with rounding to three digits by default, can be obtained using the `printSE` function, whose arguments are:

```
args(printSE)

## function (SE.data, round.digit = 3, round.out = TRUE)
## NULL
```

For example:

```
sdSE <- StdDev.SE(edhec[, "CA"])
printSE(sdSE)
```

```
##      SD IFiid IFcor
## 1 0.02 0.003 0.006
```

Of course if you want to compute any subset of the four IF based SE's and the two bootstrap SE's, you can do. For example, you can obtain all six of those standard error estimates as follows:

```
# Sharpe Ratio SE computation for single hedge fund
Sharpe.out <- SharpeRatio.SE(edhec[, "CA"],
                             se.method=c("IFiid", "IFcor",
                                           "IFcorAdapt", "IFcorPW",
                                           "BOOTiid", "BOOTcor"))
```

```
# Print output
```

```
printSE(Sharpe.out)
```

```
##      SR IFiid IFcor IFcorAdapt IFcorPW BOOTiid BOOTcor
## CA 0.32 0.123 0.155      0.262   0.299   0.13   0.201
```

The risk and performance estimator functions allow you to return the standard errors for more than one asset or portfolio, e.g. a portfolio of assets, at the same time. The following code results in standard errors for all thirteen of the `edhec` hedge funds, using the first three method in the set of six methods.

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

```
##      SR IFiid IFcor IFcorAdapt
## CA 0.320 0.123 0.155      0.265
## CTA 0.258 0.081 0.080      0.080
## DIS 0.433 0.118 0.145      0.187
## EM 0.214 0.094 0.119      0.120
## EMN 0.667 0.180 0.184      0.184
```

```
## ED  0.415 0.116 0.137      0.149
## FIA 0.299 0.130 0.158      0.201
## GM  0.451 0.073 0.081      0.081
## LS  0.350 0.090 0.114      0.114
## MA  0.608 0.133 0.144      0.146
## RV  0.508 0.135 0.166      0.194
## SS  0.076 0.079 0.093      0.093
## FOF 0.325 0.092 0.119      0.114
```

Help files for the functions in the **RPESE** are available in the usual ways. For example, you can get the help file for Sharpe Ratio in both of the following ways.

```
?SharpeRatio.SE
help(SharpeRatio.SE)
```

3.4 Outlier Cleaning

There is also the outlier cleaning functionality in the **RPEIF** package that is fully described in Section 7 of CM, and is available in **RPESE**. Here we illustrate the use of the outlier cleaning facility in terms of the influence function transformed returns for the sample mean estimator. It is shown in Section 2 of CM that the influence function for the sample mean estimator is $IF(r; \mu) = r - \mu$. Thus the IF transformed returns time series, computed with the function **IF.mean** is just $r_t - \mu$, where μ would be replaced by the sample mean in applications. The function **IF.mean** is made accessible by loading the package **RPEIF**. The following code produces Figure 2, which illustrate the effect of outlier cleaning relative to no outlier cleaning for the **FIA** hedge fund returns.

```
library(RPEIF)
IFout <- IF.mean(returns = edhec[, "FIA"], cleanOutliers = F, IFprint = T)
IFout.clean <- IF.mean(returns = edhec[, "FIA"], cleanOutliers = T,
                       IFprint = T)
par(mfrow = c(2,1))
ylim = c(-.1, .035)
plot.zoo(IFout, type = "b",
         ylab = expression(paste("Returns - ", mu)),
         main = "FIA Returns", pch = 20, lwd = .8, cex = .9,
```

```

ylim = ylim)
plot.zoo(IFout.clean,type = "b", ylab = expression(paste("Returns - ",mu)),
        main = "FIA Outlier Cleaned Returns", pch = 20, lwd = .8, cex = .9,
        ylim = ylim)
par(mfrow = c(1,1))

```




Figure 2: FIA Returns Versus Outlier Cleaned FIA Returns (with sample mean subtracted)

You can use the following code to compare Sharpe ratio SE's without and with outlier cleaning.

```

# IFcorAdapt SE results with outliers present and with outliers cleaned
SR <- SharpeRatio.SE(edhec,se.method = "IFcorAdapt",cleanOutliers = F)

```

```

SRcl.out <- SharpeRatio.SE(edhec,se.method = "IFcorAdapt",cleanOutliers = T)
clean.compare <- data.frame(SR$IFcorAdapt, SRcl.out$IFcorAdapt)
names(clean.compare) <- c("With Outliers", "Outliers Cleaned")
row.names(clean.compare) <- names(edhec)
round(clean.compare,3)

```

##	With Outliers	Outliers Cleaned
## CA	0.265	0.180
## CTA	0.080	0.080
## DIS	0.187	0.158
## EM	0.120	0.115
## EMN	0.184	0.137
## ED	0.149	0.137
## FIA	0.201	0.168
## GM	0.081	0.083
## LS	0.114	0.113
## MA	0.146	0.137
## RV	0.194	0.156
## SS	0.093	0.095
## FOF	0.120	0.116

It is not surprising that the SE's of the Sharpe ratio are smaller with outlier cleaning than with the outliers in the returns, as outliers generally inflate estimator variability.

3.5 Exponential and Gamma Distributions

The RPEGLMEN package used in the fitting of the elastic net penalized GLM model to the periodogram of the IF transformed return series was developed and tested for exponential distribution GLM models. However some initial work was done for that type of GLM model fitting for the family of Gamma distribution. And initial results for the Gamma distribution, reported in CM, indicates that the Gamma distribution GLM often results in a more parsimonious polynomial model fit the the periodogram. By way of example, the following code, which is very slow, computes standard errors of Sharpe ratio estimators for exponential and Gamma distributions.

```

# Sharpe Ratio SE computation for edhec hedge funds using Gamma distribution
Clean.out <- SharpeRatio.SE(edhec,
                           se.method=c("IFiid","IFcor","IFcorAdapt"),
                           cleanOutliers = T)
GammaClean.out <- SharpeRatio.SE(edhec,
                                se.method=c("IFiid","IFcor","IFcorAdapt"),
                                cleanOutliers = T, fitting.method = "Gamma")
GammaExp.comparison <- cbind(printSE(Clean.out)[,4],
                             printSE(GammaClean.out)[,4])
colnames(GammaExp.comparison) <- c("IFcorAdapt", "IFcorAdapt-Gamma")
rownames(GammaExp.comparison) <- names(edhec)
GammaExp.comparison

##      IFcorAdapt IFcorAdapt-Gamma
## CA          0.184          0.206
## CTA          0.080          0.050
## DIS          0.150          0.199
## EM           0.122          0.144
## EMN          0.137          0.202
## ED           0.137          0.183
## FIA          0.168          0.188
## GM           0.083          0.085
## LS           0.113          0.111
## MA           0.139          0.096
## RV           0.156          0.188
## SS           0.095          0.094
## FOF          0.116          0.150

```

We regard the Gamma family code implemented in `RPEGLMEN`, used for the above computation, to be quite experimental at this stage, and anticipate further development of the code in C++ at a future date.

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>.

S Zhang, R D Martin, and A A Christidis. Influence functions for risk and performance estimators. *Working paper*, 2019.