Lecture 2-d
Bootstrap & Value-at-Risk

 1

Review: Sampling Distribution of $\bar{X}$ and s^2

• Recall that the sampling distribution of the sample mean is:

$$\bar{X} \sim N(\mu, \sigma^2/N)$$

Example: Using Shiller's S&P 500 monthly log returns (N = **1805**, $\bar{X}$ = **0.007378**, s = **0.0404559**)

$$\Rightarrow \bar{X} \sim N(0.007378, .00095^2)$$

• Under normal variates, the sampling distribution of s^2 is:

$$(N-1)\, s^2/\sigma^2 \sim \chi^2_{N-1}.$$

Example: Using Shiller's data.

Estimated Var[s^2] = $2 * \sigma^4/(N-1)$ = 2 * **0.040455**^4/1804 =
= 2.969493e-09

S.E.(s^2) = sqrt(2.969493e-09) = **0.0000545** (or 0.0055%). 2

Review – Returns: Expected Returns & the ERP

• Q: How precise is the computation of ERP = E[($r_{M,t}$ - $r_{f,t}$)]?

Example: We use Shiller's monthly data, with 150 years of data, to produce an estimate of the ERP = E[($r_{M,t}$ - $r_{f,t}$)]:

Annualized Market return = **0.007378** * **12** = **0.088536**

Annualized risk-free rate = **0.04511**

ERP = **0.088536** - **0.04511** = **0.043426** (**4.34%**)

• Many economists would consider this estimated ERP as "*too high*." Why? The standard "neoclassical" model gets an ERP close to **1%.** To get to **4.34%** (or **6%**) the model requires extremely high risk aversion..

• It turns out that ERP is not easy to estimate precisely. Duarte and Rosa (2015) list over 20 models to estimate ERP in the U.S. With 1**960-2013** data, they report estimates from **-0.4%** to **13.1%**.

3

Review – Returns: ERP – Wide Range

• Using **50 years of monthly data**, we report ERP annual estimates for developed markets (in USD), from **0.88%** (Italy) to **11.56** (HK). For EM, there is also big dispersion of ERP estimates (also big SDs!)

• We use the SE as a measure of precision of an estimate. For the sample mean, $\bar{X}$, we have: S.E.($\bar{X}$) = $s/\sqrt{T}$

U.S.:	**15.01**/sqrt(620/12) = 2.0882%	
Germany:	21.48/sqrt(620/12) = 2.9883%	
Hong Kong:	**33.23**/sqrt(620/12) = 4.6230 %	⇐ Effect of *T*
Brazil:	**47.67**/sqrt(404/12) = 8.2157 %	
Russia:	**47.54**/sqrt(320/12) = 9.2061%	
India:	**28.35**/sqrt(344/12) = 5.2950%	

4

Review – Returns: Risk-Return – Sharpe Ratio

- A popular statistics that provides a measure of the risk-return trade-off for an asset is the Sharpe ratio (SR), the ratio of the excess expected return of an asset to its risk, measured by its return volatility (SD). We estimate the SR of asset *i* with

$$\widehat{SR}_i = \frac{\hat{\mu}_i - r_f}{s_i} \qquad (\hat{\mu}_i = \bar{X} \text{ for the return of } i)$$

Interpretation: A 1% change in risk, increases excess returns by SR%.

Using the previous data, we calculate the SR for several markets:

U.S.: **0.0382**/**.1501** = 0.2545

Switzerland: 0.0585/**.1764** = 0.3316

Hong Kong: **0.1156**/**.3323** = 0.3479

Russia: **0.1660**/**.4754** = 0.3492 ⟸ Best trade-off!

India: **0.0760**/**.2835** = 0.2681.

5

Review – Hypothesis Testing

- Testing involves the comparison between two competing hypothesis:
  - H_0: The maintained hypothesis.
  - H_1: The hypothesis considered if H_0.

- Idea: We collect a sample, $X = \{X_1, X_2, \dots, X_n\}$. We construct a statistic $T(X) = f(X)$, called the *test statistic*. Now we have a decision rule:
  - If $T(X)$ is contained in space R, we reject H_0 (& we learn).
  - If $T(X)$ is in the complement of R (R^C), we fail to reject H_0.

Note: $T(X)$, like any other statistic, is a RV. It has a distribution. We use the distribution of $T(X)$ to determine R (& we associate a probability to R).

6

Review – Hypothesis Testing

Example: Suppose $T(X) = \bar{X}$. If data is normal, the distribution of $\bar{X}$ is also normal. Then, under H_0, we build a Rejection Region, *R:*

$$R = [\bar{X} < T_{LB}\ ,\ T_{UB} > \bar{X}]$$

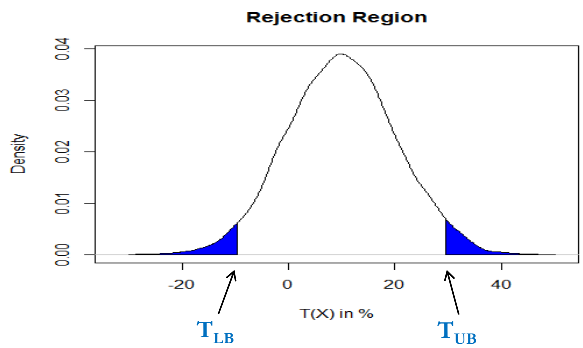

Note: The blue area ("*significance level*") represents the $P[R|H_0]$. For example, if the blue area is 5%, then, T_{LB} = **-1.96** & T_{UB} = **1.96**.

7

Review – Hypothesis Testing: *p-value*

• The *classical approach*, also known as *significance testing*, relies on *p-values:*

p-value is the probability of observing a result at least as extreme as the test statistic, under H_0.

Example: Suppose $T(X) \sim \chi^2_2$. We compute $\widehat{T(X)}$ = **7.378**. Then,

p-value($\widehat{T(X)}$ = **7.378**) = 1 – Prob[$T(X)$ < **7.378**] = 0.025

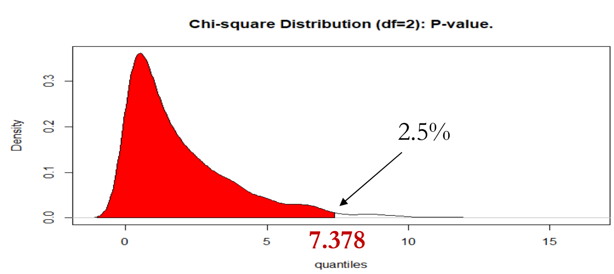

8

Review – Hypothesis Testing: Steps

- Steps for the *classical approach*, also known as *significance testing*:

1. Identify H_0 & set a *significance level* (α%).

2. Determine the appropriate test statistic T(X) and its distribution under the assumption that H_0 is true.

3. Calculate T(X) from the data.

4. Rule: If *p-value* of T(X) < α ⇒ Reject H_0 (& we learn H_0! is not true).
If *p-value* of T(X) > α ⇒ Fail to reject H_0. (No learning.)

Note: In Step **4**, setting α% is equivalent to setting R. Thus, instead of looking at *p-value*, we can look if T(X) falls in R (in the blue area). We do this by constructing a (1 - α)% C.I.

- Mistakes are made. We want to quantify these mistakes.

9

Review – Hypothesis Testing: H_0: $\mu = \mu_0$

Example: We want to test if the mean is equal to μ_0. Then,

1. H_0: $\mu = \mu_0$.
H_1: $\mu \neq \mu_0$.

2. Appropriate T(X): *t-test* (based on σ unknown and estimated by s).
Determine distribution of T(X) under H_0. Sampling distribution of $\bar{X}$, under H_0: $\bar{X} \sim N(\mu_0, \sigma^2/N)$. Then, distribution of T(X) under H_0:

$$t = \frac{\bar{X} - \mu_0}{s/\sqrt{N}} \sim t_{N-1} \qquad \text{–when } N > 30,\ t_N \sim N(0, 1).$$

3. Compute t, $\hat{t}$, using $\bar{X}$, μ_0, s, and N. Get *p-value*($\hat{t}$).

4. Rule: Set an α level. If *p-value*($\hat{t}$) < α ⇒ Reject H_0: $\mu = \mu_0$.
Alternatively, if $|\hat{t}| > t_{N-1,1-\alpha/2}$ ⇒ Reject H_0: $\mu = \mu_0$.

10

Review – H. Testing: Are Excess Returns Zero?

Example 1: We test if the **S&P 500 monthly excess return** is zero.
Data (1871-2021): $\bar{X}$ = **0.003619**, s = **0.04052**, N = 1805. Then,

1. H_0: $\mu = 0$.
   H_1: $\mu \neq 0$.

2. Appropriate test: $t = \dfrac{\bar{X} - \mu_0}{s/\sqrt{N}}$

3. $\hat{t} = \dfrac{0.003619}{0.04052/\sqrt{1805}}$ = **3.7941** & *p-value*($\hat{t}$) = **0.00015**

4. Rule: *p-value*($\hat{t}$) = **0.00015** < α = .05 ⇒ Reject H_0: $\mu = 0$.

Alternatively, $|\hat{t}$ = **3.7941** $| > t_{N-1,1-.05/2}$ = **1.96** ⇒ Reject H_0: $\mu = 0$.

Conclusion: S&P 500 monthly mean excess returns are not equal to zero. ¶

11

Review – H. Testing: Are Excess Returns Zero?

Example 1 (continuation): The observed $\hat{t}$ = **3.7941** is in the rejection region built around H_0: (**-1.96**, **1.96**).

Note: $t_{1789,\,1-.05/2} = z_{1-.05/2}$ = **1.96** (since N > 30.)

And, by symmetry of Normal: $-z_{05/2} = z_{1-.05/2}$ = **1.96**

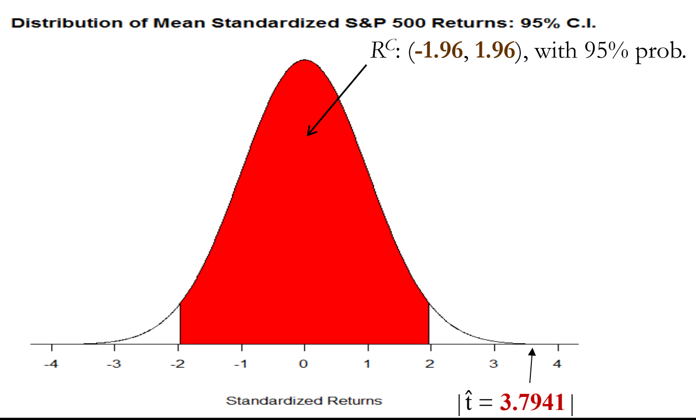

12

Review – Confidence Intervals (C.I.)

• When we estimate parameters with an estimator, $\hat{\theta}$, we get a point estimate for θ. For example, in the previous example, $\bar{X}$ = **0.003619**.

• Broader concept: Estimate a set C_n, a collection of values in R^k. For example, $\mu \in C_n = [L_n; U_n]$, called an *interval estimate* for θ.

• The goal of C_n is to contain the true population value, θ. The wider the interval C_n, the more uncertain we are about our estimate, $\hat{\theta}$.

• Interval estimates C_n are called *confidence intervals* (C.I.), usually noted with the coverage probability $(1 - \alpha)\%$.

13

Review – Confidence Intervals (C.I.)

• When we know the distribution of $\hat{\theta}$, it is easy to construct a C.I. For example, if the distribution of $\hat{\theta}$ is **normal**, then a $(1 - \alpha)\%$ C.I.:

$$C_n = [\hat{\theta} \pm z_{1-\alpha/2} * \text{Estimated SE}(\hat{\theta})] \qquad (-z_{\alpha/2} = z_{1-\alpha/2})$$

Example: We estimate a 95% C.I. for the monthly S&P mean excess return. Assuming, normality, $\bar{X} \sim N(\mu, \sigma^2/N)$. Then, a $(1 - \alpha)\%$ C.I.:

$$C_n = [\bar{X} - z_{1-.05/2} * SD(\bar{X}),\ \bar{X} + z_{1-.05/2} * SD(\bar{X})]$$

$$C_n = [\mathbf{0.003619} - \mathbf{1.96} * [\mathbf{0.04052}/\sqrt{1805}],\ \mathbf{0.003619} + \mathbf{1.96} * [\mathbf{0.04052}/\sqrt{1805}]$$

$= [0.00175, 0.00549] =$ [**0.18%, 0.55%**].

<u>Note</u>: By looking at the 95% C.I., we reject than monthly S&P Composite excess returns are 0, since 0% is outside the 95% C.I.

14

Review – C.I. for Monthly Stock Excess Returns

Example (continuation): Reject H_0: $\mu = 0$, since 0 is outside the observed 95% C.I.

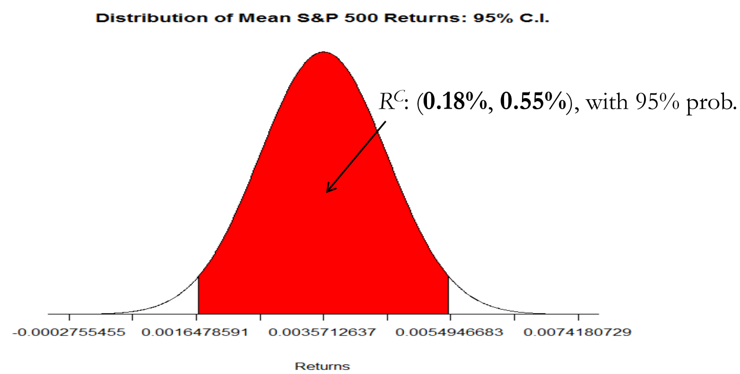

• Note: This 95% C.I. shows why Mehra & Prescott (1985) thought the ERP was *too high*. Their theoretical monthly ERP was, at most, 0.0833% (=1%/12). ¶

15

Review – C.I. for the Variance of Stock Returns

Example: We estimate a 95% C.I. for the variance of **monthly S&P 500 mean total return.** Assuming normality, s , once scaled, follows:

$$(N-1)\, s^2/\sigma^2 \sim \chi^2_{N-1}.$$

First, we rewrite the standard C.I. for a chi-squared variable:

$$P(\chi^2_{\upsilon,\alpha/2} < \chi^2_{\nu} < \chi^2_{\upsilon,1-\alpha/2}) = 1 - \alpha$$

$$P(\chi^2_{N-1,\alpha/2} < (N-1)\, s^2/\sigma^2 < \chi^2_{N-1,1-\alpha/2}) = 1 - \alpha \qquad (\upsilon = N-1)$$

After some easy algebra, we derive:

$$P[(N-1)\, s^2/\chi^2_{\upsilon,1-\alpha/2} < \sigma^2 < (N-1)\, s^2/\chi^2_{\upsilon,\alpha/2}] = 1 - \alpha.$$

Note: This C.I. is not symmetric. But, as the degrees of freedom (N) get large, the χ^2_{N-1} starts to look like a Normal and, thus, more symmetric.

16

Review – C.I. for the Variance of Stock Returns

Example (continuation): From the χ^2_{1804} distribution, we get: $\chi^2_{1804,.025}$ = **1688.2** & $\chi^2_{1804,.975}$ = **1923.6**.

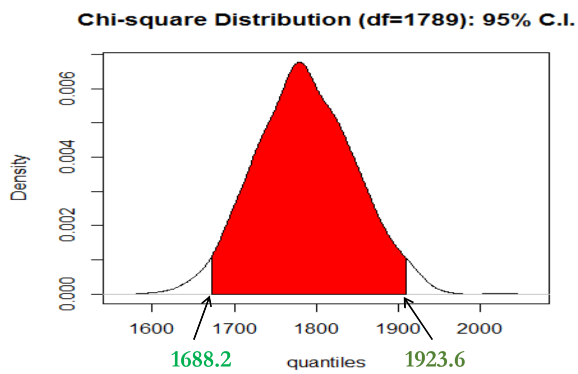

P[1804 * (**0.04046**)²/(**1923.6**) < σ² < 1804 * (**0.04046**)²/(**1688.2**)] = .95
P[0.001535 < σ² < 0.001749] = .95

17

Review – C.I. for the Variance of Stock Returns

Example (continuation):
P[0.001535 < σ² < 0.001749] = .95
Taking square root above delivers a 95% C.I. for σ:
⇒ 95% C.I. for σ is given by (**3.918%**, **4.182%**).

C.I. is compact around the estimated **0.04046** ⇒ σ is measured with accuracy.

<u>Note</u>: Usually *N* is large (N>30). We can use the normal approximation to calculate CIs for the population σ. For the S&P data:

$\text{SE}[s] = s/\sqrt{2 * (N-1)}$ = **0.04046**/$\sqrt{2 * 1804}$ = 0.00067 (or .067%)

A 95% CI for σ is given by:
(**4.046%** ± **1.96** * .067%) = (**3.914%**, **4.178%**). (Very close!)

18

C.I. Application: Using the ED

• In the previous examples, we assumed that we knew the distribution of the data: Stock returns follow a normal distribution.

Q: What happens when the data follows an unknown distribution, *F*?

• Well, we still can use $\bar{X}$ or s^2 as estimates of μ and σ^2, since they have good properties when *N* is large: consistency & asymptotic normality.

• But, if ***F*** is **unknown** and/or ***N*** is **not large enough** (or the normal approximation is not good), we still can build a C.I. for any statistic using a new method: a *bootstrap*.

• The *bootstrap* is a method for estimating the sampling distribution of a statistic, $\theta = \theta(x_1, x_2, \ldots, x_N)$, by resampling from the ED, where

$$x_1, x_2, \ldots, x_N \sim i.i.d.\ F \text{ (unknown)}$$

C.I. Application: Using the ED

Idea: We use the data at hand -the empirical distribution (ED)- to estimate the variation of a statistic θ that is itself computed from the same data. Recall that, for large samples drawn from *F*, the ED approximates the CDF of *F* very well.

• The *bootstrap* is a resampling mechanism used to provide information about the sampling distribution of a statistic θ. Bootstrapping uses the ED –i.e., sample- as if it were the true CDF.

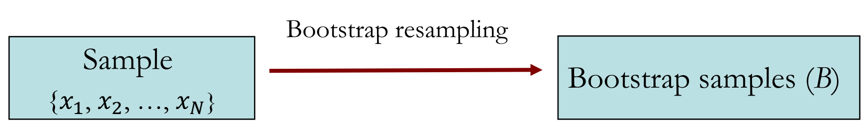

C.I. Application: Using the ED – The Bootstrap

• Suppose we have N *i.i.d.* observations drawn from *F(x)*:

$$\{x_1, x_2, \ldots, x_N\}$$

From the ED, *F**, we sample with replacement N observations:

$$\{x_1^* = x_2, x_2^* = x_4, x_3^* = x_4, x_4^* = x_{55}, \ldots, x_N^* = x_{N-8}\}$$

This is an *empirical bootstrap sample*, which is a resample of the same size N as the original data, drawn from *F**.

• For any statistic θ computed from the original sample data, we can define a statistic θ^* by the same formula, but computed instead using the resampled data.

• θ^* is computed by resampling the original data; we can compute many θ^* by resampling many times from *F**. Say, we resample θ^* B times.

C.I. Application: Using the ED – The Bootstrap

Example: You are interested in the relation between CEO's education (**X**) and firm's long-term performance (***y***). You have 1,500 observations on both variables. You estimate the correlation coefficient, ρ (=θ), with its sample counterpart, r (=$\hat{\theta}$). You find the correlation to be very low.

• Q: How reliable is this result? The distribution of r is complicated. You decide to use a bootstrap to study the distribution of r.

• Randomly construct a sequence of B samples (all with N=1,500). Say,

$B_1 = \{(x_1, y_1), (x_3, y_3), (x_6, y_6), (x_6, y_6), \ldots, (x_{1458}, y_{1458})\} \Rightarrow \hat{\theta}_1^* = r_1$

$B_2 = \{(x_5, y_5), (x_7, y_7), (x_{11}, y_{11}), (x_{12}, y_{12}), \ldots, (x_{1486}, y_{1486})\} \Rightarrow \hat{\theta}_2^* = r_2$

. . . .

$B_B = \{(x_2, y_2), (x_2, y_2), (x_2, y_2), (x_3, y_3), \ldots, (x_{1499}, y_{1499})\} \Rightarrow \hat{\theta}_B^* = r$

C.I. Application: Using the ED – The Bootstrap

• We have a collection of estimated θ^*:

$$\{\hat{\theta}_1^*, \hat{\theta}_2^*, \hat{\theta}_3^*, \ldots, \hat{\theta}_B^*\}.$$

From this collection of $\hat{\theta}^*$'s, we can compute the mean, the variance, skewness, draw a histogram, etc., and confidence intervals.

• Below, we plot the histogram of bootstrapped correlations and defined a C.I. with $[r_{LB}, r_{UB}]$ $\Rightarrow \rho \in [r_{LB}, r_{UB}]$ with some probability

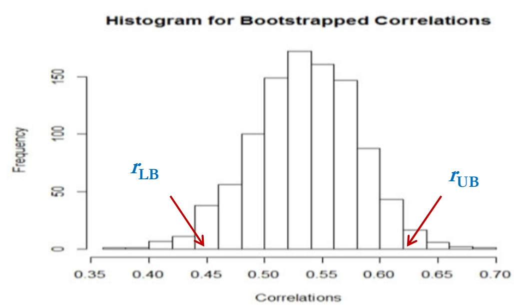

C.I. Application: Using the ED – The Bootstrap

• Using the histogram or the sorted $\{\hat{\theta}_1^*, \hat{\theta}_2^*, \hat{\theta}_3^*, \ldots, \hat{\theta}_B^*\}$, we build a $(1-\alpha)$% C.I. Using the histogram:

- The lower bound leaves $\alpha/2$% of the $\hat{\theta}^*$ to the right
- The upper bound leaves $(1-\alpha/2)$% of the $\hat{\theta}^*$ to the left.

• Bootstrap Steps:

1. From the original sample, draw random sample with size N.
2. Compute statistic θ from the resample in 1: $\hat{\theta}_1^*$.
3. Repeat steps 1 & 2 B times $\Rightarrow$ Get B statistics: $\{\hat{\theta}_1^*, \hat{\theta}_2^*, \hat{\theta}_3^*, \ldots, \hat{\theta}_B^*\}$
4. Compute moments, draw histograms, etc. for these B statistics.

C.I. Application: Using the ED – The Bootstrap

• Results (Bootstrap Principle):

1. With a large enough B, the LLN allows us to use the $\hat{\theta}^*$'s to estimate the distribution of $\hat{\theta}$, $F(\hat{\theta})$.

2. The variation in $\hat{\theta}$ is well approximated by the variation in $\hat{\theta}^*$.

Result 2 is the one that we use to estimate the size of a C.I.

• There are many ways to construct a C.I. using bootstrapping. The easier one is the one described above. Just use the distribution of the $\hat{\theta}^*$'s to compute directly a C.I. This is the *bootstrap percentile method*.

The percentile method uses the distribution of $\hat{\theta}^*$ as an approximation to the distribution of $\hat{\theta}$.

C.I. Application: Bootstrap Percentile Method

• Technical Note: The bootstrap delivers consistent results only.

Example: We construct a 95% C.I. for the variance of S&P 500 monthly returns (continuation of previous examples). Using the boot.ci function, with type=**perc**, from ***boot*** package. First, install boot, use the function **install.packages()** ⇒ install.packages("boot"). Then, call library(boot):

```
Sh_da <- read.csv("https://www.bauer.uh.edu/rsusmel/4397/Shiller_2020data.csv", head=TRUE, sep=",")
SP <- Sh_da$P
T <- length(SP)
lr <- log(SP[-1]/SP[-T])
lr_var <- var(lr)
T_s <- length (lr)

library(boot)
```

C.I. Application: Bootstrap Percentile Method

Example (continuation):

```
# function to obtain the variance from the data
var_p <- function(data, i) {
        d <- data[i]
        return(var(d))
}

library(boot)
boot.ci(boot.samps, type = "perc")

> boot.ci(boot.samps, type = "perc")
BOOTSTRAP CONFIDENCE INTERVAL CALCULATIONS
Based on 1000 bootstrap replicates

CALL :
boot.ci(boot.out = boot.samps, type = "perc")

Intervals :
Level     Percentile
95%   ( 0.0014,  0.0020 )
Calculations and Intervals on Original Scale
```

C.I. Application: Bootstrap Percentile Method

Example (continuation):

• Check results by sorting boot.samps$t.

```
> new <- sort(boot.samps$t)
> new[25]                              # CI's Lower Bound
[1] 0.001398215
> new[975]                             # CI's Upper Bound
[1] 0.001955096
```

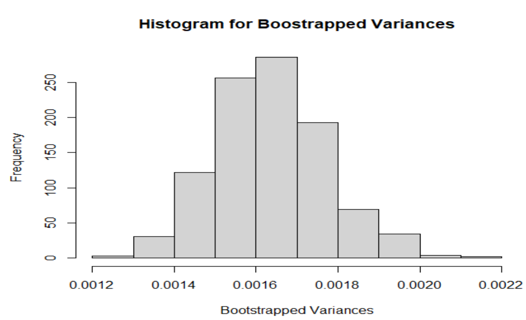

Or for σ, taking square roots: 95% CI ∈ [3.74%, 4.42%].

C.I. Application: Empirical Bootstrap

• The percentile method uses the distribution of $\hat{\theta}^*$ as an approximation to the distribution of $\hat{\theta}$. It is very simple, but there are more appealing methods. In general, a bootstrap based on comparing differences is sounder. This is the key to the *empirical bootstrap.*

• To build a C.I. for θ, we use $\hat{\theta}$, computed from the original sample. As in the previous C.I.'s, we want to know how far is $\hat{\theta}$ from θ. For this, we would like to know the distribution of

$$q = \hat{\theta} - \theta.$$

• If we knew the distribution of $q = \hat{\theta} - \theta$, we build a (1 – α)% C.I., by finding the critical values $q_{\alpha/2}$ & $q_{(1-\alpha/2)}$ to have:

$$\Pr(q_{\alpha/2} \leq \hat{\theta} - \theta \leq q_{(1-\alpha/2)} | \theta) = 1 - \alpha$$

C.I. Application: Empirical Bootstrap

• Or, after some manipulations:

$$\Pr(\hat{\theta} - q_{\alpha/2} \geq \theta \geq \hat{\theta} - q_{(1-\alpha/2)} | \theta) = 1 - \alpha,$$

which gives a (1 – α)% C.I.:

$$C_n = [\hat{\theta} - q_{(1-\alpha/2)}, \hat{\theta} - q_{\alpha/2}]$$

• We do not know the distribution of q, but we can use the bootstrap to estimate it with

$$q^* = \hat{\theta}^* - \hat{\theta}.$$

and, then, to get $q^*_{\alpha/2}$ & $q^*_{(1-\alpha/2)}$:

$$C_n = [\hat{\theta} - q^*_{(1-\alpha/2)}, \hat{\theta} - q^*_{\alpha/2}]$$

• This C.I. is called the *pivotal* C.I.

C.I. Application: Empirical Bootstrap

• Intuition: The distribution of $\hat{\theta}$ is 'centered' at θ, while the distribution of $\hat{\theta}^*$ is centered at $\hat{\theta}$. If there is a significant separation between $\hat{\theta}$ and θ, these two distributions will also differ significantly.

On the other hand, the distribution of $q = \hat{\theta} - \theta$ describes the variation of $\hat{\theta}$ about its center. Similarly, the distribution of $q^* = \hat{\theta}^* - \hat{\theta}$ describes the variation of $\hat{\theta}^*$ about $\hat{\theta}$.

Then, even if the centers are quite different, the two variations about the centers can be approximately equal.

C.I. Application: Empirical Bootstrap

Example: We estimate a 95% C.I. for the variance of monthly returns of the S&P 500.

Note: You need to install R package ***resample***, using the **install.packages()** function ⇒ install.packages("resample").

```
sim_size <- 1000                    # B = size of bootstrap
library(resample)                   # call library resample
data_star <- sample(lr, T_s*sim_size, replace=TRUE)       # create B resamples of size T_s
boot_sample <- matrix(data_star, nrow=T_s, ncol=sim_size)  # organize resamples in matrix
boots_vars  <- colVars(boot_sample)     # compute the variance for each bootsrap sample
q_star <- boots_vars - lr_var           # Compute q* for each bootstrap sample
q <- quantile(q_star, c(0.025, 0.975))  # Find the 0.025 and 0.975 quantile for q_star
ci <- lr_var - c(q[2], q[1])            # Calculate the 95% C.I. for the variance.
cat("Confidence interval: ",ci, "\n")   # Print C.I using cat
```

C.I. Application: Empirical Bootstrap

Example (continuation):

```
> lr_var
[1] 0.001637
> ci
      97.5%        2.5%
0.001376664 0.001909769

> cat("Confidence interval: ", ci, "\n")
Confidence interval:  0.001376664 0.001909769
>
```

Or for σ, the 95% CI is given by (**3.71%**, **4.37%**).

C.I. Application: Empirical Bootstrap

Example: We construct the same 95% C.I. for the variance of monthly S&P 500 returns using the R package ***boot***.

```
library(boot)
boot.samps <- boot(data=lr, statistic=var_p, R=sim_size)    # resampling and θ* estimation
boot.ci(boot.samps, type = "basic")                         # boot computes the CI.
> boot.ci(boot.samps, type = "basic")
CALL :
boot.ci(boot.out = boot.samps, type = "basic")

Intervals :
Level      Basic
95%   ( 0.0014,  0.0019 )
Calculations and Intervals on Original Scale
```

C.I. Application: Empirical Bootstrap

Example (continuation): Check results using step-by-step procedure:

```
q_star <- boot.samps$t - lr_var          # q* = θ* − θ̂
q_ad <- sort(q_star)                      # sort q*
> lr_var - q_ad[975]                      # CI's Lower Bound
[1] 0.001357793
> lr_var - q_ad[25]                       # CI's Upper Bound
[1] 0.001914674
```

We can transform this CI for the variance into a CI for the SD:

```
> sqrt(lr_var - q_ad[975])
[1] 0.03684825
> sqrt(lr_var - q_ad[25])
[1] 0.04375699
```

A 95% CI for σ is given by (**3.68%**, **4.38%**), wider than the CI assuming a Normal distribution for returns.

C.I. Application: Empirical Bootstrap

• It is common to gauge the uncertainty of the estimation of θ by computing the sample standard error, SE($\hat{\theta}^*$):

• Steps

1. Computing the sample variance:

$$\text{Var}(\hat{\theta}^*) = \frac{1}{B-1}\sum_{i=i}^{B}(\hat{\theta}_i^* - \bar{\theta}^*)^2,$$

where $\bar{\theta}^* = \frac{1}{B}\sum_{i=i}^{B}\hat{\theta}_i^*$.

2. Estimate the S.E. of $\hat{\theta}^*$: SE($\hat{\theta}^*$) = sqrt[Var($\hat{\theta}^*$)].

Example: Estimate the SE(s^2):

1. Compute: $\text{Var}(s^2) = \frac{1}{B-1}\sum_{i=i}^{B}(s_i^{2*} - s_B^2)^2$, where $s_B^2 = \frac{1}{B}\sum_{i=i}^{B}s_i^{2*}$.
2. SE(s^2) = sqrt[Var(s^2)].

C.I. Application: Parametric Bootstrap Method

• If we assume the data is from a parametric model (say, from a Normal or a Gamma distribution), we can use the parametric bootstrap to access the uncertainty (variance, C.I.) of the estimated parameter.

A parametric bootstrap generates bootstrap samples from the assumed distribution, based on moments computed from the sample; *not* from the ED.

• Suppose we have a sample with N observations drawn from $F(x;\theta)$:

$$\{x_1, x_2, \ldots, x_N\}$$

We know $F(x;\theta)$, but do not know its parameters. Suppose there is only one unknown parameter, θ. From the sample, we compute $\hat{\theta}$. Then, we bootstrap from $F(x;\hat{\theta})$ and proceed as before to form a C.I.

C.I. Application: Parametric Bootstrap Method

• Steps:

1. Draw B samples of size N from $F(x;\hat{\theta})$.
2. For each bootstrap sample, $\{x_1^*, x_2^*, x_3^*, ..., x_N^*\}$, calculate $\hat{\theta}^*$.
   ⇒ Get B $\hat{\theta}^*$'s.
3. Estimate a C.I. using the previous methods.

Example: We have a sample: $\{x_1, x_2, \ldots, x_N\}$ drawn from a $N(0,\sigma^2)$. We estimate s^2 and use a parametric bootstrap to gauge its uncertainty with a C.I.

Steps:

1. Draw B samples of size N from a $N(0, s^2)$ ⇒ Compute B s^{2*}.
2. Estimate a C.I. using the previous methods.

C.I. Application: Parametric Bootstrap Method

Example: Suppose **S&P 500 monthly returns** follow a N(0, σ^2). We estimate σ^2 with s^2:

```
> lr_var                                           # s^2
[1] 0.001656445
lr_sd <- sqrt(lr_var)

x <- rnorm(T_s*sim_size, mean=0, sd=lr_sd)          # generate normal data
boot_sample <- matrix(x, nrow=T_s, ncol=sim_size)  # organize simulated data
boots_vars  <- colVars(boot_sample)                # compute variances
q_star <- boots_vars - lr_var                      # q* = θ* − θ̂
q <- quantile(q_star, c(0.025, 0.975))
ci <- lr_var - c(q[2], q[1])
> ci
     97.5%        2.5%
0.001547382 0.001760286
```

Or for σ, the 95% CI is given by (**3.94%**, **4.20%**). Very close to the CIs we obtained before assuming a Normal distribution for returns.

C.I. Application: Bootstrapping - Why?

• Q: Why do we need a bootstrap?

- Sample sizes are "small" and asymptotic assumptions do not apply
- DGP assumptions are violated.
- Distributions are complicated.

• Usually, we would not use a bootstrap to compute C.I.'s for the mean, in general, the normal distribution works well, as long as N is large enough.

• The bootstrap is used to generate C.I. & S.E. for estimates of other statistics where the normal distribution is not a good approximation. A typical example is the median, where for non-normal underlying distributions, the SE[median] is complicated to compute.

C.I. Application: Value-at-Risk

• Q: What is the most an investor can lose with a particular investment over a given time framework? Or, what is the worst case scenario?

• *Value-at-Risk* (VaR) provides one answer, a (lower) bound with a probability attached to it.

• So far, we have measured risk of an asset/investment with its volatility.

• Volatility is calculated including positive (right tail) and negative (left tail) returns. Investors, however, love the right tail of the returns distribution, but dislike the other tail. VaR focuses on the left tail.

• VaR gives a formal definition of "worst case scenario" for an asset over a time period.

C.I. Application: Value-at-Risk

• VaR gives a formal definition of "worst case scenario" for an asset.

VaR: *Maximum expected amount (loss)* in a given time interval within a (*one-sided*) (1 - α)% C.I.:

VaR(1 - α) = Amount exposed * (1 + worst % change scenario in C.I.)

It is common to express the "*expected loss*" relative to today's expected value of asset/investment:

VaR-mean(1 - α) = VaR – E[Amount exposed]

• There are different ways to compute the worst case scenario within a time interval. We go over two approaches:
- Assuming a probability distribution (normal, in our case).
- Using the empirical distribution (a bootstrap, using the past).

C.I. Application: Value-at-Risk

Example: $\alpha = .05$
VaR = Amount exposed * (1 + worst change scenario in 97.5% C.I.).
VaR-mean(97.5%) = VaR – E[Amount exposed]

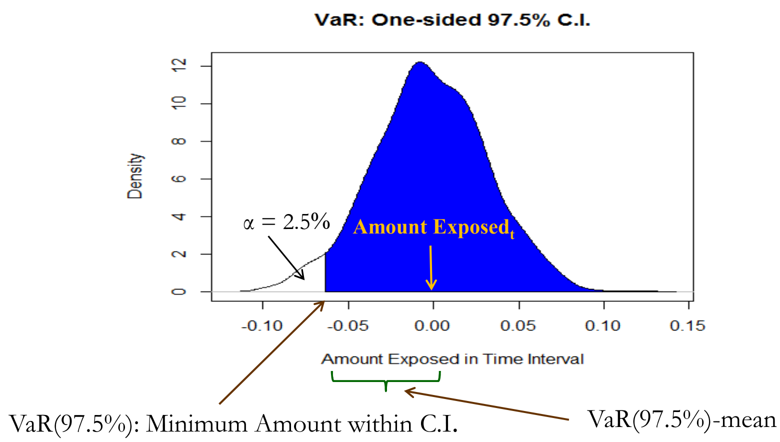

C.I. Application: VaR in FX Markets

- When a company is involved with transactions denominated in foreign currency (FC), it is exposed to *currency risk*. Transaction exposure (TE) provides a simple measure of this exposure:

$$TE_t = \text{Value of a fixed future transaction in FC} * S_t$$

where S_t is the exchange rate expressed as units of domestic currency (USD for us) per unit of FC (say, EUR).

Example: A Swiss company, Swiss Cruises, sells packages in USD.

Amount = **USD 1 million**.

Payment: 30 days.

S_t = **0.92 CHF/USD**

$$\Rightarrow \mathbf{TE_t = USD\ 1M * 0.92\ CHF/USD = CHF\ 0.92M}.$$

If S_t is described by a Random Walk, then TE_t is a forecast of the value of the transaction in 30 days (TE_{t+30}). ¶

C.I. Application: VaR in FX Markets (Normal)

- Swiss Cruises wants a measure of the uncertainty related to the amount to receive in CHF in 30 days, since S_{t+30} is unknown.

We can use a range to quantify this uncertainty, we want to say

$TE_{t+30} \in [TE_{LB}, TE_{UB}]$ with high probability.

To determine this range for TE, we **assume** that (log) changes in S_t, $e_{f,t}$, are **normally distributed**: $e_{f,t} \sim N(\mu, \sigma^2)$.

Then, we build a (1 - α)% interval around the mean: $[\mu \pm z_{1-\alpha/2} * \sigma]$.

Usual α's in interval calculations: $\alpha = .05 \Rightarrow |z_{.025}| = \mathbf{1.96}$ (≈2)

$\alpha = .02 \Rightarrow |z_{.01}| = \mathbf{2.33}$

As usual, we estimate (μ, σ) using $(\bar{X}, s)$.

C.I. Application: VaR in FX Markets (Normal)

Example: Range estimate based on a **Normal distribution**.

Assume Swiss Cruises believes that CHF/USD monthly changes follow a normal distribution. Swiss Cruises estimates the mean and variance using the last 15 years of data.

$\bar{X}$ = Monthly mean = **-0.00152** ≈ -0.15%

s^2 = Monthly variance = 0.001014 (⇒ s = **0.03184**, or 3.18%)

$e_{f,t} \sim N(-0.00152, 0.03184^2)$ $e_{f,t}$ = CHF/USD log changes.

Swiss Cruises constructs a 95% CI for CHF/USD monthly changes.

Recall that a 95% C.I. for $e_{f,t+30}$ (which applies to any t) is given by:

$e_{f,t} \in [-0.00152 \pm 1.96 * 0.03184] = [-0.06393; 0.06089]$.

Based on this range for $e_{f,t}$, we can build a 95% C.I. for S_{t+30} and, then, for TE_{t+30} (= **USD 1M** * S_{t+30}).

C.I. Application: VaR in FX Markets (Normal)

Example (continuation):

$e_{f,t+30} \in$ [-0.00152 ± 1.96 * 0.03184] = [-0.06393; 0.06089].

Now, based on the range $e_{f,t}$, we derive a range for S_{t+30}:

(A) Upper bound
$S_{t+30,UB} = S_t$ * (1+ $e_{f,UB}$) = **0.92 CHF/USD** * (1 + **0.06089**) =
= **0.97602**

(B) Lower bound
$S_{t+30,LB} = S_t$ * (1+ $e_{f,LB}$) = **0.92 CHF/USD** * (1 - **0.06393**)] =
= **0.86118**

⇒ $S_{t+30} \in$ [**0.97602 CHF/USD**; **0.86118 CHF/USD**].

C.I. Application: VaR in FX Markets (Normal)

Example (continuation): Finally, we derive the bounds for the TE:

(A) Upper bound ($S_{t+30,UB} = S_t$ * (1+ $e_{f,UB}$) = **0.97602 CHF/USD**)
TE_{UB}: **USD 1M** * **0.97602 CHF/USD** = **CHF 976,019**.

(B) Lower bound ($S_{t+30,LB} = S_t$ * (1+ $e_{f,LB}$) = **0.86118 CHF/USD**)
TE_{LB}: **USD 1M** * **0.86118 CHF/USD** = **CHF 861,184**.

⇒ $TE_{t+30} \in$ [**CHF 0.861 M**; **CHF 0.976 M**]. ¶

• The lower bound, for a receivable, represents the worst case scenario within the interval. This is the *Value-at-Risk* (VaR) interpretation:

VaR: *Maximum expected amount (loss)* in a given time interval within a (*one-sided*) confidence interval.

VaR = Amount exposed * (1+ worst % change scenario in C.I.)

C.I. Application: VaR in FX Markets (Normal)

VaR = Amount exposed * (1+ worst % change scenario in C.I.).
Then, in our application: VaR = $\mathbf{TE_t}$ * (1 + $e_{f,LB}$)

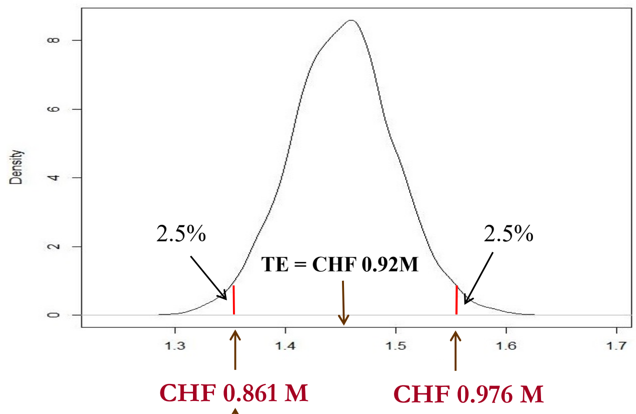

VaR(97.5%): Minimum revenue within a 97.5% C.I.

C.I. Application: VaR in FX Markets (Normal)

It is common to express the "*expected loss*" relative to today's expected value of transaction (or asset):

VaR-mean = VaR – $\mathbf{TE_t}$ = $\mathbf{TE_t}$ * (1 + $e_{f,LB}$) – $\mathbf{TE_t}$
= $\mathbf{TE_t}$ * $e_{f,LB}$

Or just

VaR-mean = Amount exposed * worst case scenario

Example (continuation): The minimum revenue to be received by SC in the next 30 days, within a 97.5% CI.

VaR(97.5%) = **CHF 0.92M** * [1+ (-**0.00152** – **1.96** * **0.03184**)]
= **CHF 0.8612M**.

Interpretation of VaR: If SC expects to cover expenses with this USD inflow, the maximum amount in CHF to cover, within a 97.5% one-sided CI, should be **CHF 0.8612M**.

C.I. Application: VaR in FX Markets (Normal)

• Relative to today's valuation (or *expected valuation*, according to RWM), the maximum expected loss in 30 days within a 97.5% one-sided C.I. is:

VaR-mean(.975) = **CHF 0.8612M** – **CHF 0.92M** = **CHF -0.0588M**.

Note that we can also compute the VaR-mean as:

VaR-mean(.975) = **CHF 0.92M** * (-0.00152 – **1.96** * **0.03184**)
= **CHF -0.0588M**. ¶

• Technically speaking, the VaR is a *quantile*, where a quantile is the fraction of observations that lie below a given value (in this case the VaR).

In the previous example, the 0.025 quantile (or 2.5% quantile) for expected loses is **CHF -0.0588M**.

C.I. Application: VaR in FX Markets (Normal)

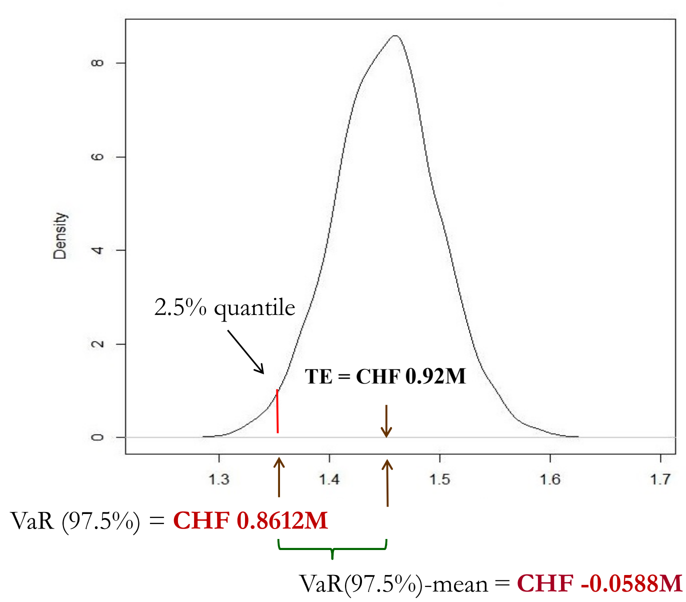

C.I. Application: VaR in FX Markets (Normal)

Example (continuation):

Note: We could have used a different quantile –i.e. a different significant level- to calculate the VaR, for example 1% (⇒ $z_{.99}$ = **2.33**). Then,

VaR(99%) = **CHF 0.92M** * [1+ (**-0.00152** – **2.33** * **0.03184**)]
= **CHF 0.8503M** (A more *conservative* bound.)

⇒ VaR-mean (.99) = **CHF 0.8503M** – **CHF 0.92M** = **CHF -0.0697M**

Interpretation of VaR-mean: Relative to today's valuation (or *expected valuation*, according to RWM), the maximum *expected loss* with a 99% "chance" is **CHF -0.0697M**.

Note: As the C.I. gets wider, Swiss Cruises can spend less CHF on account of the USD 1M receivable. ¶

C.I. Application: VaR in FX Markets (Normal)

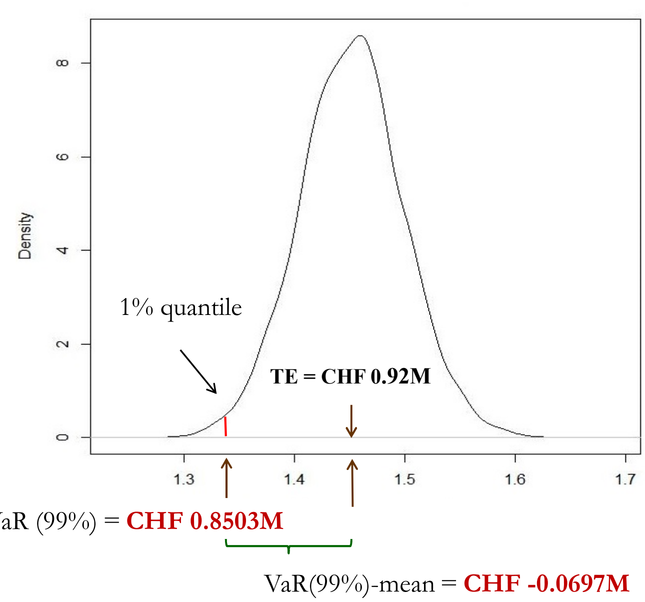

C.I. Application: VaR in FX Markets (Bootstrap)

- VaR is a statistic –a function of the data. We can do an empirical bootstrap to calculate the mean, SE (=SD), C.I., etc.

Example: We want to calculate the average VaR(97.5%) and its S.E., using all CHF/USD data from 1990:Jan - 2023:July. Then,

```
chfusd <- read.csv("http://www.bauer.uh.edu/rsusmel/4386/chfusd.csv",sep=",")  # Data
S <- chfusd$CHF_USD                          # Extract CHF_USD column of the data
T <- length(S)                               # Check total T (1971:1 to 2023:7)
Tstart <- 229                                # Start of sample period: 1990:1
SP <- S[Tstart: T]
T <- length(SP)
Val <- 1000000                               # Value of transaction in FC (in M)
S_0 <- S[T]                                  # Today's S_t
e_f <- log(SP[-1]/SP[-T])
T_s <- length(e_f)
alpha = .05                                  # Specify alpha level for VaR
T_s_low <- round(T_s*alpha/2)                # Obs corresponding to alpha/2*T_s
TE_o <- Val*S_0*(1+e_f)                      # calculate Original TE values
STE_o <- sort(TE_o)                          # sort Original TE
```

C.I. Application: VaR in FX Markets (Bootstrap)

Example (continuation):

```
VaR_o <- STE_o[T_s_low]                      # Original VaR
> VaR_o
[1] 860293

# function to obtain VaR from the data
varisk <- function(data, i) {
d <-data[i]
TE <- Val*S_0*(1+d)                                   # calculate R TE values
STE <- sort(TE)                                       # sort TE
VaR <- STE[T_s_low]
return(VaR)
}
library(boot)
sim_size <- 1000
boot.samps <- boot(data=e_f, statistic=varisk, R=sim_size)
```

C.I. Application: VaR in FX Markets (Bootstrap)

Example (continuation):

> boot.samps
Bootstrap Statistics :
original bias std. error
t1* 860293 1929.305 4870.733

> boot.ci(boot.samps, type = "basic")
Intervals :
Level Basic
95% (849352, 867587)

> mean(boot.samps$t)
[1] 862222.3
> sd(boot.samps$t)
[1] 4870.733

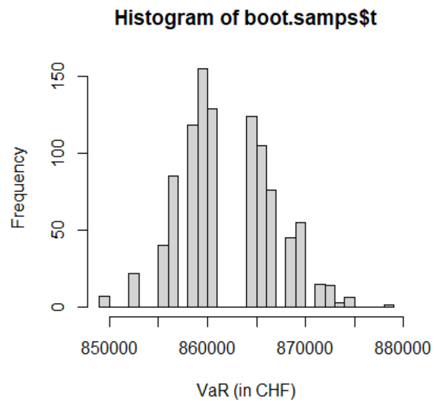

hist(boot.samps$t, xlab="VaR (in CHF)", breaks=30)

Bootstrap estimated VaR(97.5%) = **CHF 0.8622M**