File Edit View Run Kernel Tabs Settings Help

Launcher | Exam2_fake.ipynb

Code | Julia 1.7.0

Load Packages

•••

[1]: TaskLocalRNG()

•••

Introduction

This exam explores how autocorrelation ought to change how we test statistical hypotheses.

Task 1

Code a function for simulating T observations from an AR(1) series

$y_t = (1-\rho)\mu + \rho y_{t-1} + \varepsilon_t \sigma$ where ε_t is N(0,1).

That is, generate $y_1, y_2, \ldots, y_T$ from this formula.

To make also the starting value (y_0) random, simulate $T + 100$ data points, but then discard the first 100 values of y_t.

Generate a single "sample" using `(T,ρ,σ,μ) = (500,0,3,2)`. Calculate and report the average (mean) and the first 5 autocorrelations (hint: `autocor()`) of this sample. Redo a 2nd time, but with `ρ=0.75`.

•••

[3]: SimAR1 (generic function with 1 method)

•••

```
fake results
average from one sample with ρ=0     2.369

autocorrelations with ρ=0
1    -0.177
2     0.056
3    -0.010
4    -0.000
5    -0.038
```

•••

```
fake results
```

```
average from one sample with ρ=0.75     2.335

autocorrelations with ρ=0.75
1     0.266
2     0.099
3     0.034
4     0.004
5    -0.041
```

Task 2

Do a Monte Carlo simulation. Use the parameters `(T,ρ,σ,μ) = (500,0,3,2)`.

1. Generate a sample with T observations and calculate the average. Repeat $M = 10,000$ times and store the estimated averages in a vector of length M. (The rest of the question uses the symbol μ_i to denote the average from sample i.)

2. What is average μ_i across the M estimates? (That is, what is $\frac{1}{M}\sum_{i=1}^{M}\mu_i$?) *Report* the result.

3. What is the standard deviation of μ_i across the M estimates? Compare with the theoretical standard deviation (see below). *Report* the result.

4. Does the distribution of μ_i look normal? *Plot* a histogram and compare with the theoretical pdf (see below).

...basic stats (the theoretical results)

says that the sample average of an iid ("independently and identically distributed") data series is normally distributed with a mean equal to the true (population) mean μ and a standard deviation equal to $s = \sigma_y/\sqrt{T}$ where σ_y is the standard deviation of y.

To compare with our simulation results, you could estimate σ_y from a single simulation with very many observations (say 10'000).

•••

```
fake results
Average across the simulations:      1.998

Std across the samples (with ρ=0) and in theory:
simulations    theory
     0.198     0.222
```

•••

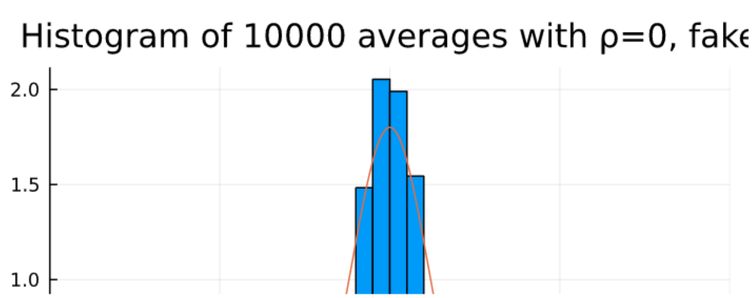

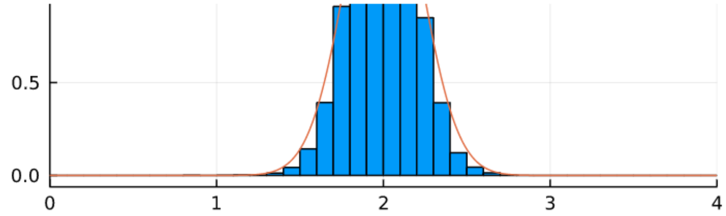

Task 3

Redo task 2, but now use `ρ=0.75` (the other parameters are unchanged).

•••

```
fake results
Average across the simulations:      2.005

Std across the samples (with ρ=0.75) and in theory:
simulations    theory
     0.308     0.351
```

•••

Histogram of 10000 averages with ρ=0.75, fa

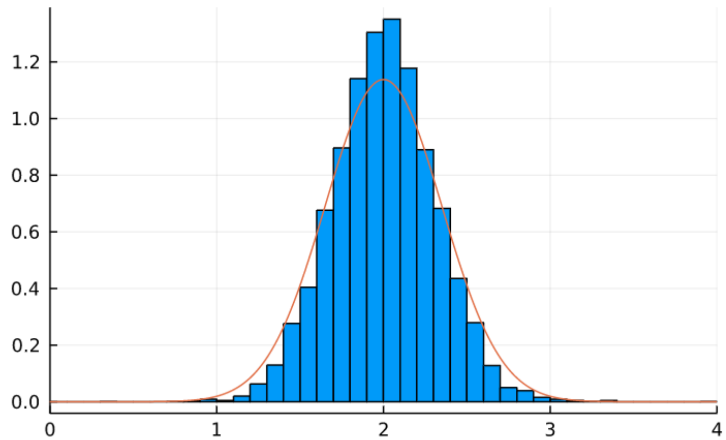

Task 4

You decide to test the hypothesis that $\mu = 2$. Your decision rule is

- reject the hypothesis if $|(\mu_i - 2)/s| > 1.645$ with $s = \sigma_y/\sqrt{T}$

With this decision rule, you are clearly assuming that the theoretical result (definition of s) is correct.

Estimate both μ_i and σ_y from each sample.

In what fraction of the M simulation do you reject your hypothesis when $\rho = 0$ and when $\rho = 0.75$? For the other parameters, use `(T,σ,μ) = (500,3,2)` (same as before).

•••

•••

```
fake results
Frequency of rejections:
     with ρ=0  with ρ=0.75
        0.043        0.211
```

•••