

Load Packages

In []:

```
using Printf, Statistics, StatsBase, Random, Distributions
include("jlFiles/printmat.jl")
Random.seed!(678)           #set the random number generator to this starting point
```

```
TaskLocalRNG()
```

In []:

```
using Plots

gr(size=(480,320))
default(fmt = :svg)
```

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, \dots, 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 $(T, \rho, \sigma, \mu) = (500, 0.3, 2)$. Calculate and report the average (mean) and the first 5 autocorrelations (hint: `autocor()`) of this sample. Redo a 2nd time, but with $\rho=0.75$.

In []:

```
function SimAR1(T, ρ, σ, μ)
    y = fill(NaN, T + 100)
    e = rand(Normal(0, 1), T + 100)
    y[1] = 1
    for i = 2:T + 100
        y[i] = (1-ρ)μ + ρ*y[i-1] + e[i]*σ
    end
    return y[101:end]
end
```

```
SimAR1 (generic function with 1 method)
```

In []:

```
y1 = SimAR1(500,0,3,2)
printmat("average from one sample with  $\rho=0$ ", mean(y1))
printmat("autocorrelations with  $\rho=0$ ")
printmat(1:5, autocor(y1)[2:6])
```

average from one sample with $\rho=0$ 1.986

autocorrelations with $\rho=0$

1	-0.016
2	0.027
3	0.004
4	0.004
5	0.008

In []:

```
y2 = SimAR1(500,0.75,3,2)
printmat("average from one sample with  $\rho=0.75$ ", mean(y2))
printmat("autocorrelations with  $\rho=0.75$ ")
printmat(1:5, autocor(y2)[2:6])
```

average from one sample with $\rho=0.75$ 3.187

autocorrelations with $\rho=0.75$

1	0.766
2	0.580
3	0.495
4	0.441
5	0.372

Task 2

Do a Monte Carlo simulation. Use the parameters $(T, \rho, \sigma, \mu) = (500, 0, 3, 2)$.

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).

In []:

```
M = 10_000

(T, rho, sigma, mu) = (500, 0, 3, 2)

mu1l = fill(NaN, M)
sigma1l = fill(NaN, M)

# Monte Carlo simulation
for i = 1:M
    y = SimAR1(T, rho, sigma, mu)
    mu1l[i] = mean(y)
    sigma1l[i] = std(y)
end

# Theoretical results
y1 = SimAR1(10_000, rho, sigma, mu)
sigma_y1 = std(y1)
s1 = sigma_y1/sqrt(T)

printmat("Average across the simulations:", mean(mu1l))
println("Std across the samples (with rho=0) and in theory:")
printmat(["simulations", std(mu1l)], ["theory", s1])
```

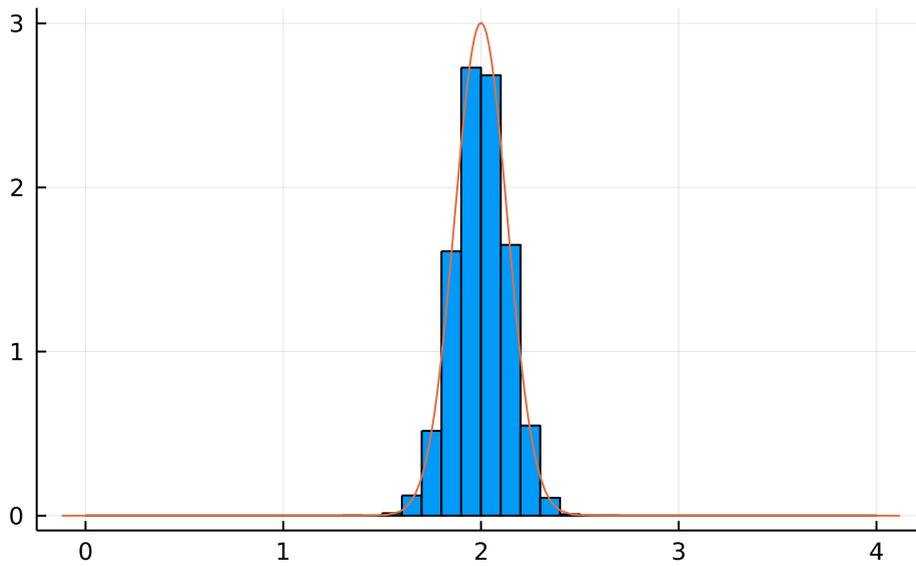
```
Average across the simulations:      2.001
```

```
Std across the samples (with rho=0) and in theory:
simulations      theory
    0.134        0.133
```

In []:

```
histogram( $\mu_1$ ,bins=0:0.1:4,normalize=true,legend=false,title="Histogram of 10000  
averages with  $\rho=0$ ")  
plot!( $\mu_1 \rightarrow$ pdf(Normal( $\mu$ ,  $s_1$ ),  $\mu_1$ ))
```

Histogram of 10000 averages with $\rho=0$



Task 3

Redo task 2, but now use $\rho=0.75$ (the other parameters are unchanged).

In []:

```
M = 10_000

(T,q2,σ,μ) = (500,0.75,3,2)

μi2 = fill(NaN, M)
σi2 = fill(NaN, M)

# Monte Carlo simulation
for i = 1:M
    y = SimAR1(T, q2, σ, μ)
    μi2[i] = mean(y)
    σi2[i] = std(y)
end

# Theoretical results
y2 = SimAR1(10_000,q2,σ,μ)
σy2 = std(y2)
s2 = σy2/sqrt(T)

printmat("Average across the simulations:", mean(μi2))
println("Std across the samples (with q=0.75) and in theory:")
printmat(["simulations", std(μi2)], ["theory", s2])
```

Average across the simulations: 1.989

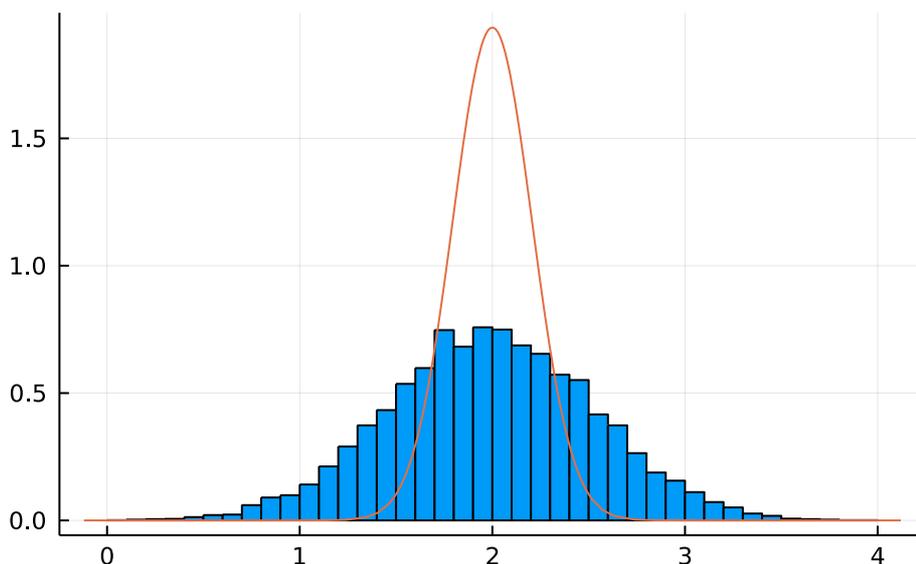
Std across the samples (with q=0.75) and in theory:

simulations	theory
0.534	0.206

In []:

```
histogram(μi2, bins=0:0.1:4, normalize=true, legend=false, title="Histogram of 10000 averages with q=0.75")
plot!(μi2->pdf(Normal(μ, s2), μi2))
```

Histogram of 10000 averages with $\rho=0.75$



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 $(T, \sigma, \mu) = (500, 3, 2)$ (same as before).

In []:

```
(T,σ,μ) = (500,3,2)

rejection1 = 0
rejection2 = 0

# Count how many times we reject the hypothesis
for i = 1:M
    # rejections for ρ = 0
    s1 = σi1[i]/sqrt(T)
    if abs((μi1[i] - 2)/s1) > 1.645
        rejection1 += 1
    end

    # rejections for ρ = 0.75
    s2 = σi2[i]/sqrt(T)
    if abs((μi2[i] - 2)/s2) > 1.645
        rejection2 += 1
    end
end

println("Frequency of rejections:")
printmat(["with ρ=0 ", rejection1 / (M)], ["with ρ=0.75 ", rejection2 / (M)])
```

Frequency of rejections:

with ρ=0	with ρ=0.75
0.098	0.536