Homework 2

Due: Thursday 2/7/19 by 12:00pm (noon)

Grading Scheme:

- Maximum of 2 points for 1., determined as follows:
  - 0 points for no solutions whatsoever or `R` output only
  - 1 point for an honest effort but very few correct answers
  - 2 points for the correct plot and written answers that are on the right track
- Maximum of 2 points for 2., determined as follows:
  - 0 points for no solutions whatsoever or `R` output only
  - 1 point for an honest effort but very few correct answers
  - 2 points for the correct plot and written answers that are on the right track
- 1 point for 3. if autocovariances and plot are correct and written answers are on the right track

Solutions are given below in blue.

An Overview of Level-α Tests

This homework is going to ask you to conduct a level-α tests of a null hypothesis, which requires that you combine a few bits of information from class.

Let's call the null hypothesis H and the alternative hypothesis K. Suppose we have a test statistic $\hat{t}$, that is a function of the data, and that we know either exactly or approximately what the distribution of $\hat{t}$ is under the null H. Then we can construct a level-α test of the null hypothesis H by comparing $\hat{t}$ to the $1-\alpha/2$ and $\alpha/2$ quantiles of the distribution of $\hat{t}$ under the null H. If $\hat{t}$ is within those quantiles, we **accept** the null hypothesis H, otherwise we reject it.

This general idea is something that I expect have seen in your previous statistics classes, and we referenced it in our review of linear regression when we talked about testing the null hypothesis H that $\hat{\beta}_j$ is exactly equal to a specific value for some j. Consider testing the null hypothesis $H: \beta_1 = 0$ against the alternative $K: \beta_1 \neq 0$. For such a problem, our test statistic is

$$\hat{t} = \frac{\hat{\beta}_1}{s_w\sqrt{\left(\boldsymbol{Z}'\boldsymbol{Z}\right)^{-1}_{11}}},$$

and we know that $\hat{t} \sim \mathcal{T}_{n-q}$ under H, where q is the number of covariates we have in our regression model (number of columns of $\boldsymbol{Z}$). We perform a level-α test of H by comparing $\hat{t}$ to the $1-\alpha/2$ and $\alpha/2$ quantiles of a $\mathcal{T}_{n-q}$ distribution. If $\hat{t}$ is within these quantiles, we accept $H: \beta_1 = 0$, otherwise we reject it. Here's a bit of `R` code to demonstrate what I mean, in this example:

```
# Let's work through this example with our chicken data
library(astsa)
data(chicken)
n <- length(chicken)
reg <- lm(chicken~time(chicken))
s.w <- summary(reg)$sig
q <- length(coef(reg))
# Let's test if the time coefficient is equal to zero
# First, construct the test statistic
ZtZ.inv <- solve(crossprod(model.matrix(reg)))
```

```r
t.hat <- coef(reg)["time(chicken)"]/(s.w*sqrt(ZtZ.inv[2, 2]))
# Note: Another way to get this would be to set
# t.hat = summary(reg)$coef["time(chicken)", "t value"]
#
# In this case, we know that the test statistic should be t-distributed under the
# null with n-q degrees of freedom. The quantiles for a level alpha = 0.05 test will be
alpha <- 0.05
q <- qt(c(alpha/2, 1 - alpha/2), df = n - q)
# Compare the test statistic to these quantiles, do we accept?
t.hat >= q[1] & t.hat <= q[2] # Accept null if TRUE, reject otherwise
# Note: There was a typo that I made here before,
# t.hat %in% q does *not* do the right thing. I apologize for any confusion!
# Points will not be taken off if t.hat %in% q was used in
# place of t.hat >= q[1] & t.hat <= q[2]
```

To apply this idea to these problems, ask yourself: What is a sample quantity that we can calcuate for a time series $\boldsymbol{x}$ and use as a test statistic $\hat{t}$ that:

- We talked about in class;
- Is relevant to testing a hypothesis about $\rho_x(1)$;
- We know the approximate or exact distribution of under the null that $\boldsymbol{x}$ is a white noise time series, with $\rho_x(1) = 0$?

The AR(1) Model

1. This problem will ask you to work with the autoregressive (AR) model.

(a) Describe what R returns when you run `x <- arima.sim(n = 100, list(ar=1), sd = 1)`, and why this occurs.

R returns the error `Error in arima.sim(n = 100, list(ar = 1), sd = 1) :   'ar' part of model is not stationary`, because an AR(1) model is not stationary when $\phi_1 = 1$. Specifically, when $\phi_1 = 1$, x_t will have infinite variance.

(b) Simulate 1,000 **AR**(1) time series of length $n = 100$ with $\sigma_w^2 = 1$ for values of $\phi_1 = \{-0.5, -0.25, -0.125, 0, 0.125, 0.25, 0.5\}$. For each value of ϕ_1, compute the percent of simulations in which a level-0.05 test of the null hypothesis that the time series is white noise, with $\rho_x(1) = 0$, rejects the null, using $\hat{\rho}_x(1)$ from class and 3.(h) in Homework 1. Plot the percent of simulations in which a test of the null hypothesis rejects the null against ϕ_1.

```r
set.seed(1)
phi1s <- c(-0.5, -0.25, -0.125, 0, 0.125, 0.25, 0.5)
n <- 100
xs <- array(dim = c(5000, n, length(phi1s)))
for (i in 1:dim(xs)[1]) {
  for (k in 1:length(phi1s)) {
    phi1 <- phi1s[k]
    if (phi1 != 0) {
      xs[i, , k] <- arima.sim(n = n, list(ar=phi1), sd = 1)
    } else {
      xs[i, , k] <- rnorm(n)
    }
  }
}
```

```r
acf1s <- apply(xs, c(1, 3), function(x) acf(x, plot = FALSE)$acf[2])
ci <- qnorm(c(0.025, 0.975), mean = 0, sd = 1/sqrt(n))
ar.rej.null <- apply(acf1s, 2, function(x) {
  mean(!(x >= ci[1] & x <= ci[2]))
})
plot(phi1s, 100*ar.rej.null, type = "b",
     ylab = "% Simulations Reject", xlab = expression(phi[1]))
```

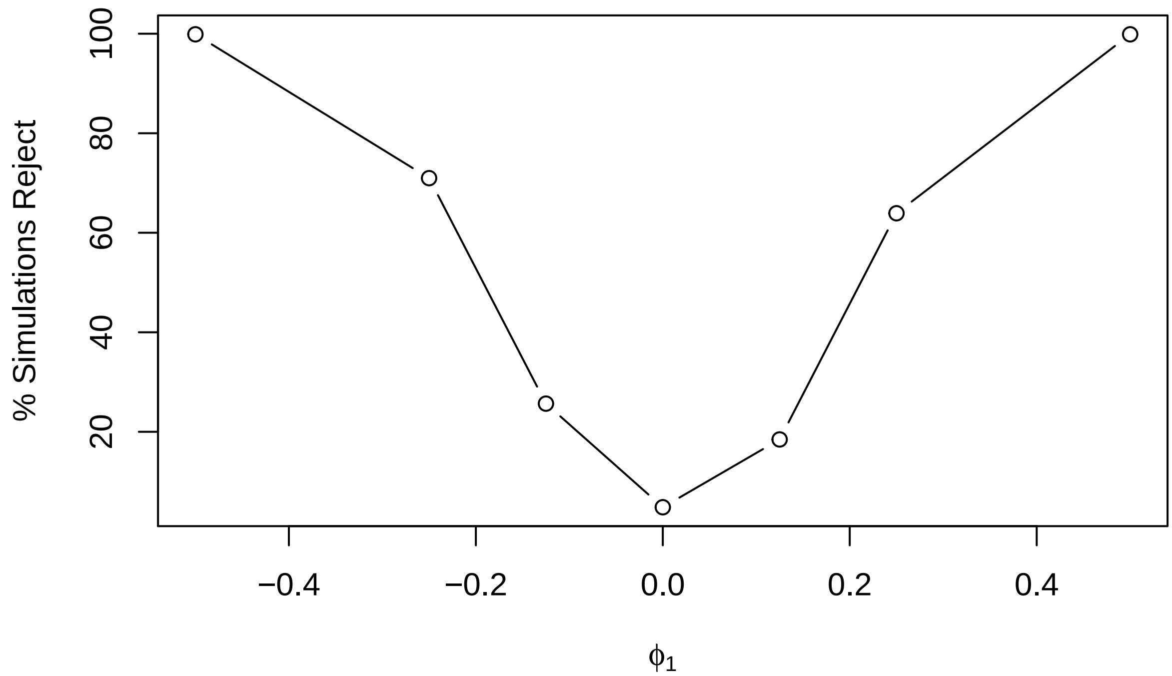

(c) When $\phi_1 \neq 0$, the percent of simulations in which a test of the null hypothesis that $\rho_x(1) = 0$ rejects the null estimates the **power** of the test. When $\phi_1 = 0$, the percent of simulations in which a test of the null hypothesis that $\rho_x(1) = 0$ rejects the null estimates the **level** of the test. Is the estimated level 0.05, as we would expect from a level-0.05 test? If not, why not?

When $\phi_1 = 0$, a test of the null hypothesis that $\rho_x(1) = 0$ rejects the null hypothesis in `4.84`% of simulations. This is a bit lower than what we would expect. The discrepancy could be due to the fact that the test is based on an **approximate** distribution for $\hat{\rho}_x(1)$, as opposed to the exact distribution.

(d) Describe in at most two sentences how the power of the test relates to the true value ϕ_1. Intuitively, does this make sense?

The power of the test is increasing with the magnitude of ϕ_1. Intuitively, this makes sense - the test is better able to detect departures when the null when the data are more different than the null, as is the case when ϕ_1 is large in magnitude.

The MA(1) Model

2. This problem will ask you to work with the moving average (MA) model.
   (a) Without using the `arima.sim` function or any other third party function for simulating an MA time series, simulate a length-100 time series $\boldsymbol{x}$ according to the MA model:

$$x_t = 0.5w_{t-1} + w_t,\ w_t \overset{i.i.d.}{\sim} \mathcal{N}(0, 1)$$

```r
n <- 100
ma <- 0.5
w <- rnorm(n + 1, mean = 0, sd = 1)
x <- numeric(length(n))
for (i in 1:n) {
  x[i] <- ma*w[i] + w[i + 1]
}
```

(b) Using the code you wrote in (a) or `arima.sim`, simulate $1,000$ **MA** (1) time series of length $n = 100$ with $\sigma_w^2 = 1$ for values of $\theta_1 = \{-1, -0.268, -0.127, 0, 0.127, 0.268, 1\}$. For each value of θ_1, compute the percent of simulations in which a test of the null hypothesis that the time series is white noise, with $\rho_x(1) = 0$, rejects the null, using $\hat{\rho}_x(1)$ from class and 3.(h) in Homework 1. Plot the percent of simulations in which a test of the null hypothesis rejects the null against θ_1.

```r
theta1s <- c(-1, -0.268, -0.127, 0, 0.127, 0.268, 1)
n <- 100
xs <- array(dim = c(5000, n, length(theta1s)))
for (i in 1:dim(xs)[1]) {
  for (k in 1:length(theta1s)) {
    theta1 <- theta1s[k]
    if (theta1 != 0) {
      xs[i, , k] <- arima.sim(n = n, list(ma=theta1), sd = 1)
    } else {
      xs[i, , k] <- rnorm(n)
    }
  }
}
acf1s <- apply(xs, c(1, 3), function(x) acf(x, plot = FALSE)$acf[2])
ci <- qnorm(c(0.025, 0.975), mean = 0, sd = 1/sqrt(n))
ma.rej.null <- apply(acf1s, 2, function(x) {
  mean(!(x >= ci[1] & x <= ci[2]))
})
plot(theta1s, 100*ma.rej.null, type = "b",
     ylab = "% Simulations Reject", xlab = expression(theta[1]))
```

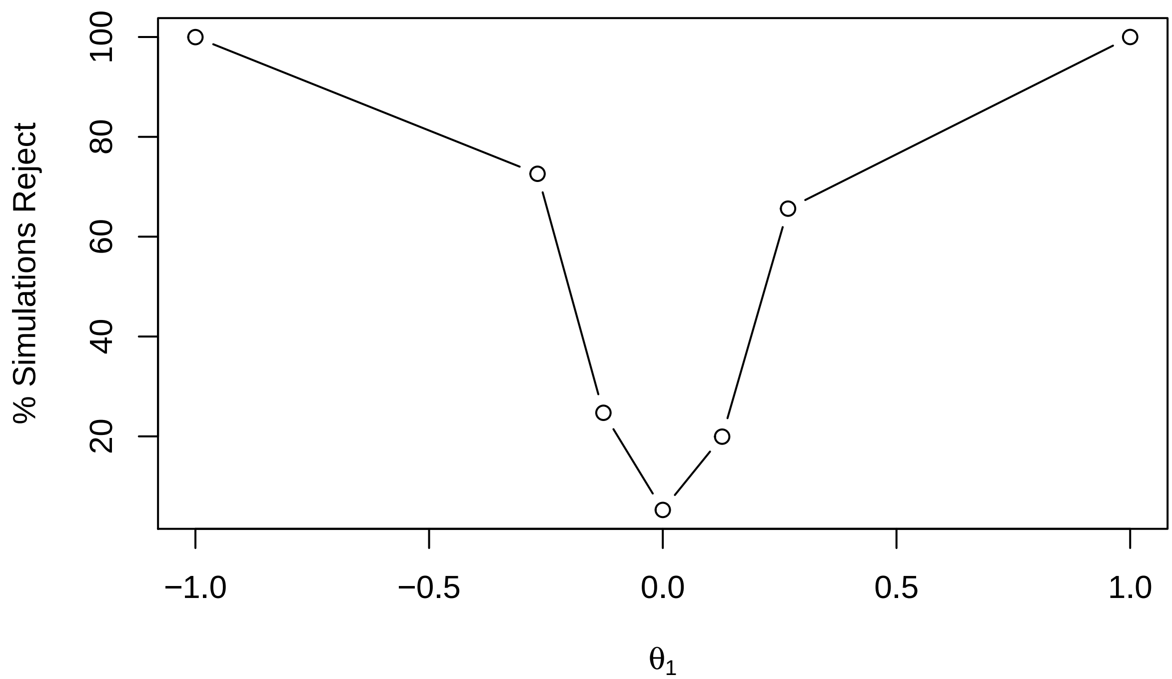

(c) When $\theta_1 \neq 0$, the percent of simulations in which a test of the null hypothesis that $\rho_x(1) = 0$ rejects the null estimates the **power** of the test. When $\theta_1 = 0$, the percent of simulations in which a test of the null hypothesis that $\rho_x(1) = 0$ rejects the null estimates the **level** of the test. Is the estimated level 0.05? If not, why not?

When $\theta_1 = 0$, a test of the null hypothesis that $\rho_x(1) = 0$ rejects the null hypothesis in `5.26`% of simulations. This is a bit lower than what we would expect. Again, the discrepancy could be due to the fact that the test is based on an **approximate** distribution for $\hat{\rho}_x(1)$, as opposed to the exact distribution.

(d) Describe in at most two sentences how the power of the test relates to the true value θ_1. Intuitively, does this make sense?

The power of the test is increasing with the magnitude of θ_1. Intuitively, this makes sense - the test is better able to detect departures when the null when the data are more different than the null, as is the case when θ_1 is large in magnitude.

Comparing AR(1) and MA(1) Models

3. This problem asks you to compare what you observed in 1. (b)-(d) to what you observed in 2. (b)-(d).

(a) Compute the true lag-one autocorrelation $\rho_x(1)$ under for an **AR**(1) model with $\phi_1 = \{-0.5, -0.25, -0.125, 0.125, 0.25, 0.5\}$.

The true lag-one autocorrelations are given in the following table.

ϕ_1	$\rho_x(1)$
-0.5	-0.5
-0.25	-0.25
-0.125	-0.125
0	0
0.125	0.125
0.25	0.25
0.5	0.5

(b) Compute the true lag-one autocorrelation $\rho_x(1)$ under for an $\mathbf{MA}(1)$ model with $\theta_1 = \{-1, -0.268, -0.127, 0, 0.127, 0.268, 1\}$.

The true lag-one autocorrelations are given in the following table, rounded to the nearest thousandth.

θ_1	$\rho_x(1)$
-1	-0.5
-0.268	-0.25
-0.127	-0.125
0	0
0.127	0.125
0.268	0.25
1	0.5

(c) Plot the the percent of simulations in which a test of the null hypothesis rejects the null against the true autocorrelation $\rho_x(1)$ for both the $\mathbf{AR}(1)$ and $\mathbf{MA}(1)$ simulations on a single plot. You should have two lines or sets of points, one for the $\mathbf{AR}(1)$ simulations and one for the $\mathbf{MA}(1)$ simulations.

```
plot(phi1s, ar.rej.null*100, type = "b", xlab = expression(rho[x](1)),
     ylab = "% Simulations Reject")
points(theta1s/(1 + theta1s^2), ma.rej.null*100, col = "blue", type = "b")
legend("bottomright", col = c("black", "blue"), pch = c(1, 1),
       legend = c("AR(1)", "MA(1)"))
```

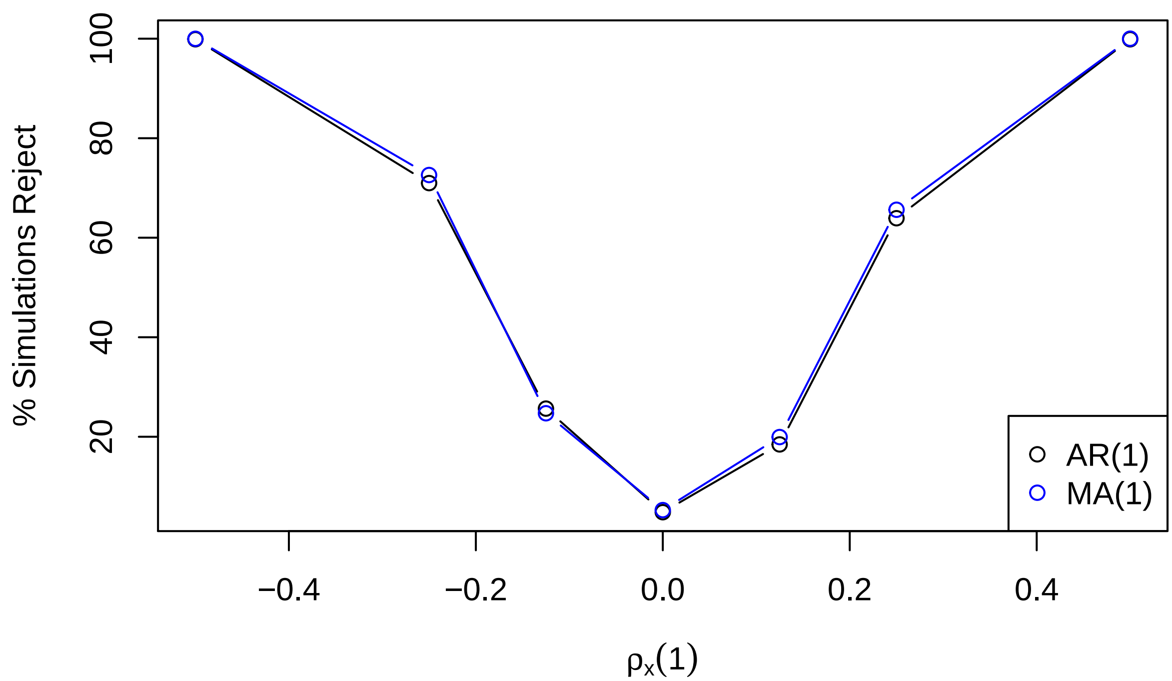

(d) In one sentence, interpret what you observe in (c), taking what you find in (a) and (b) into account.

We observe that the results are almost exactly the same for fixed $\rho_x(1)$ regardless of whether or not an $\mathbf{AR}(1)$ or $\mathbf{MA}(1)$ model is used. Intuitively, this is somewhat surprising because the $\mathbf{AR}(1)$ and $\mathbf{MA}(1)$ are different models, specifically an $\mathbf{AR}(1)$ model induces dependence of x_t on the current value w_t infinitely many past values $w_{t-1}, \dots, w_{t-j}, \dots$, whereas an $\mathbf{AR}(1)$ model incudes dependence of x_t on just the current and most recent past values w_t and w_{t-1}. However these results suggest that, at least for $n = 100$, the sample lag-one autocorrelation behaves similarly under both models.