Preliminary (take home) exam, given Friday 3/31, due Friday 4/7

Exercise P1. (20%) Consider the Gaussian location-scale model (Model III), for sample size n, i. e. observations are i.i.d $X_1, \ldots, X_n$ with distribution $N(\mu, \sigma^2)$ where $\mu \in \mathbb{R}$ and $\sigma^2 > 0$ are unknown. For a certain $\sigma_0^2 > 0$, consider hypotheses $H : \sigma^2 \leq \sigma_0^2$ vs. $K : \sigma^2 > \sigma_0^2$.

Find an α-test with rejection region of form (c, ∞) (i.e. a one-sided test) where c is a quantile of a χ^2-distribution. (Note: it is not asked to find the LR test; but the test should observe level α, on *all* parameters in the hypothesis $H : \sigma^2 \leq \sigma_0^2$.)

Hint: A good estimator of σ^2 might be a starting point.

Solution. A good estimator of σ^2 is the bias corrected sample variance:

$$\hat{S}_n^2 = \frac{1}{n-1} \sum_{i=1}^{n} (X_i - \bar{X}_n)^2.$$

According to Theorem 6.2, p. 55 it can be represented as

$$\hat{S}_n^2 = \frac{\sigma^2}{n-1} \sum_{i=1}^{n-1} \xi_i^2 \tag{1}$$

where $\xi_1, \ldots, \xi_{n-1}$ are independent standard normals. Hence $(n-1)\hat{S}_n^2/\sigma^2$ has a χ^2-distribution with $n-1$ degrees of freedom. If the hypothesis were $H : \sigma^2 = \sigma_0^2$ then

$$V_0 = \frac{n-1}{\sigma_0^2} \hat{S}_n^2$$

is a reasonable test statistic: the distribution under H is χ^2_{n-1}, which does not depend on the unspecified nuisance parameter μ. Thus the test is

$$\varphi(X) = \begin{cases} 1 \text{ if } V_0 > \chi_{n-1;1-\alpha} \\ 0 \text{ otherwise} \end{cases}$$

(where $\chi_{n-1;1-\alpha}$ is the lower $1-\alpha$ quantile). This is an α-test under $H : \sigma^2 = \sigma_0^2$; we have to ensure now that under the larger hypothesis $H : \sigma^2 \leq \sigma_0^2$.it has still level α. Now the representation (1) is true for any $\sigma^2 > 0$ when σ^2 is the true parameter. Let us write probabilities as $P_{\mu,\sigma^2}(\cdot)$ where μ, σ^2 are the pertaining "true"parameters; then

$$\begin{aligned} P_{\mu,\sigma^2}(\varphi(X) &= 1) = P_{\mu,\sigma^2}\left(\frac{n-1}{\sigma_0^2}\hat{S}_n^2 > \chi_{n-1;1-\alpha})\right) \\ &= P\left(\frac{\sigma^2}{\sigma_0^2}\sum_{i=1}^{n-1}\xi_i^2 > \chi_{n-1;1-\alpha})\right) \\ &= P\left(\sum_{i=1}^{n-1}\xi_i^2 > \frac{\sigma_0^2}{\sigma^2}\chi_{n-1;1-\alpha})\right) \end{aligned}$$

where the probability without index $P(\cdot)$ refers to $\xi_1, \ldots, \xi_{n-1}$ which have a given law (not depending on any parameters). Since under H we have $\sigma_0^2/\sigma^2 \geq 1$, it follows

$$P\left(\sum_{i=1}^{n-1} \xi_i^2 > \frac{\sigma_0^2}{\sigma^2}\chi_{n-1;1-\alpha})\right) \leq P\left(\sum_{i=1}^{n-1} \xi_i^2 > \chi_{n-1;1-\alpha})\right) = \alpha$$

thus φ has indeed level α for $H : \sigma^2 \leq \sigma_0^2$.

Exercise P2 (Two sample problem, F-test for variances). Let $X_1, \ldots, X_n$ be independent $N(\mu_1, \sigma_1^2)$ and $Y_1, \ldots, Y_n$ be independent $N(\mu_2, \sigma_2^2)$, also independent of $X_1, \ldots, X_n$ ($n > 1$) where μ_1, σ_1^2 and μ_2, σ_2^2 are all unknown. Define the statistics

$$F = F(X,Y) = \frac{S_X^2}{S_Y^2}, \tag{2}$$

$$S_X^2 = n^{-1}\sum_{i=1}^{n}\left(X_i - \bar{X}_n\right)^2, \; S_Y^2 = n^{-1}\sum_{i=1}^{n}\left(Y_i - \bar{Y}_n\right)^2$$

(here (X, Y) symbolizes the total sample).

Define the **F-distribution with k_1, k_2 degrees of freedom** (denoted F_{k_1,k_2}) as the distribution of Z_1/Z_2 where Z_i are independent r.v.'s having χ^2-distributions of k_1 and k_2 degrees of freedom, respectively.

i) (15%) Show that $F(X,Y)$ has an F-distribution if $\sigma_1^2 = \sigma_2^2$, and find the degrees of freedom.

ii) (20%) For hypotheses $H : \sigma_1^2 \leq \sigma_2^2$ vs. $K : \sigma_1^2 > \sigma_2^2$, find an α-test with rejection region of form (c, ∞) (i.e. a one-sided test) where c is a quantile of an F-distribution. (Note: it is not asked to find the LR test; but the test should observe level α, on *all* parameters in the hypothesis $H : \sigma_1^2 \leq \sigma_2^2$).

Remark. *The above definition of the F-distribution was misquoted; it is not the one found in the literature. The F-distribution F_{k_1,k_2} is commonly defined as the distribution of $k_1^{-1}Z_1/k_2^{-2}Z_2$ where Z_i are independent r.v.'s having χ^2-distributions of k_1 and k_2 degrees of freedom, respectively. Thus numerator and denominator have to be divided by the respective degree of freedom; in the handout this is given correctly. .This error in the problem formulation does not affect the solution however.*

Solution. (i) By 1 we have

$$S_X^2 = \frac{\sigma_1^2}{n}\sum_{i=1}^{n-1}\xi_i^2, \; S_Y^2 = \frac{\sigma_2^2}{n}\sum_{i=1}^{n-1}\eta_i^2 \tag{3}$$

where $\eta_1, \ldots, \eta_{n-1}$ are also independent standard normals $\eta_1, \ldots, \eta_{n-1}$. Since X_i and Y_i are independent, the ξ_i and η_i are also independent. Both $\sum_{i=1}^{n-1}\xi_i^2$ and $\sum_{i=1}^{n-1}\eta_i^2$ have a distribution χ_{n-1}^2. Introduce the (unobservable) random variable

$$F_0 = \frac{(n-1)^{-1}\sum_{i=1}^{n-1}\xi_i^2}{(n-1)^{-1}\sum_{i=1}^{n-1}\eta_i^2};$$

it has a distribution $F_{n-1,n-1}$. Then under $H:\sigma_1^2=\sigma_2^2$

$$F=F(X,Y)=\frac{S_X^2}{S_Y^2}=\frac{(n-1)^{-1}\sum_{i=1}^{n-1}\xi_i^2}{(n-1)^{-1}\sum_{i=1}^{n-1}\eta_i^2}=F_0$$

has exactly this F-distribution (since both degrees of freedom are $n-1$, the above error in the definition of the F-distribution is irrelevant).

(ii) The representation 3 is valid for general σ_1^2, σ_2^2; then

$$F(X,Y)=\frac{\sigma_1^2}{\sigma_2^2}\frac{(n-1)^{-1}\sum_{i=1}^{n-1}\xi_i^2}{(n-1)^{-1}\sum_{i=1}^{n-1}\eta_i^2}=\frac{\sigma_1^2}{\sigma_2^2}F_0.$$

This suggests $F(X,Y)$ as a reasonable test statistic for $H:\sigma_1^2\leq\sigma_2^2$, and the hypothesis is rejected when $F(X,Y)$ is too large. Thus we consider the test

$$\varphi(X)=\{\begin{array}{l}1 \text{ if } F(X,Y)>F_{n-1,n-1;1-\alpha}\\ 0 \text{ otherwise}\end{array}$$

where $F_{n-1,n-1;1-\alpha}$ is the lower $1-\alpha$-quantile of $F_{n-1,n-1}$. According to part (i), this certainly an α-test for $H:\sigma_1^2=\sigma_2^2$. It remains to be shown that also under the larger hypothesis $H:\sigma_1^2\leq\sigma_2^2$ it has level α. The proof is analogous to problem P1. Let us write probabilities as $P_{\boldsymbol{\mu},\sigma_1^2,\sigma_2^2}(\cdot)$ where $\boldsymbol{\mu}=(\mu_1,\mu_2)$, σ_1^2, σ_2^2 are the pertaining "true"parameters; then

$$\begin{aligned}P_{\boldsymbol{\mu},\sigma_1^2,\sigma_2^2}(F(X,Y) &> F_{n-1,n-1;1-\alpha})=P_{\boldsymbol{\mu},\sigma_1^2,\sigma_2^2}\left(\frac{\sigma_1^2}{\sigma_2^2}F_0>F_{n-1,n-1;1-\alpha}\right)\\ &= P_{\boldsymbol{\mu},\sigma_1^2,\sigma_2^2}\left(F_0>\frac{\sigma_2^2}{\sigma_1^2}F_{n-1,n-1;1-\alpha}\right).\end{aligned}$$

Under $H:\sigma_1^2\leq\sigma_2^2$, we have $\sigma_2^2/\sigma_1^2\geq 1$ and hence

$$P_{\boldsymbol{\mu},\sigma_1^2,\sigma_2^2}\left(F_0>\frac{\sigma_2^2}{\sigma_1^2}F_{n-1,n-1;1-\alpha}\right)\leq P_{\boldsymbol{\mu},\sigma_1^2,\sigma_2^2}\left(F_0>F_{n-1,n-1;1-\alpha}\right)=\alpha$$

thus indeed the test has level α for $H:\sigma_1^2\leq\sigma_2^2$.

Exercise P3 (25 %) **(Two sample t-test).** Let $X_1,\ldots,X_n$ be independent $N(\mu_1,\sigma^2)$ and $Y_1,\ldots,Y_n$ be independent $N(\mu_2,\sigma^2)$, also independent of $X_1,\ldots,X_n$ ($n>1$), where μ_1,μ_2, and σ^2 are all unknown. Consider hypotheses $H:\mu_1=\mu_2$ vs. $K:\mu_1\neq\mu_2$.

Show that the likelihood ratio test is equivalent to a certain t-test, i.e. a test where the critical value is chosen as the upper $\alpha/2$-quantile of a t-distribution. (*Equivalence* of two tests based on the same sample here means that they result in the same decisions, for all values of the sample).

Hint: Cp. also homework exercise H6.1. The likelihood ratio computation is similar to proposition 7.5 (ii), p. 81-82 handout.

Solution. Theorem 9.8. p. 119 handout.

Exercise P4. (10%) *Complete the proof of proposition 7.5 handout by establishing claim (i).* In detail: consider the Gaussian location-scale model (Model III), for sample size n, i. e. observations are i. i. d. $X_1, \ldots, X_n$ with distribution $N(\mu, \sigma^2)$ where $\mu \in \mathbb{R}$ and $\sigma^2 > 0$ are unknown. Consider hypotheses $H : \mu \leq \mu_0$ vs. $K : \mu > \mu_0$, and the one sided t-test which rejects when the t-statistic

$$T_{\mu_0}(X) = \frac{\left(\bar{X}_n - \mu_0\right) n^{1/2}}{\hat{S}_n}.$$

is too large (with a proper choice of critical value, such that an α-test results). Show that the one sided t-test is the likelihood ratio test for this problem.

Hints: a) When maximizing $p_{\mu,\sigma^2}(x)$ over the alternative, the supremum is not attained ($\mu > \mu_0$ is an open interval). However the supremum is the same as the maximum over $\mu \geq \mu_0$ which is attained by certain maximum likelihod estimators $\hat{\mu}_1, \hat{\sigma}_1^2$ (find these, and also MLE's $\hat{\mu}_0, \hat{\sigma}_0^2$ under H)

b) Note that this time, in difference to part (ii), we have to show that the likelihood ratio is a monotone increasing function of $T_{\mu_0}(X)$ itself, not of its absolute value. Establish this by considering separately the cases of positive and nonnegative values of the t-statistic $T_{\mu_0}(X)$.

Solution. The density p_{μ,σ^2} of the data $x = (x_1, \ldots, x_n)$ is again

$$p_{\mu,\sigma^2}(x) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{S_n^2 + (\bar{x}_n - \mu)^2}{2\sigma^2 n^{-1}}\right), \tag{4}$$

$$S_n^2 = n^{-1} \sum_{i=1}^{n} (x_i - \bar{x}_n)^2. \tag{5}$$

To find MLE's of μ and σ^2 under $\mu > \mu_0$, we first maximize for fixed σ^2 over all possible μ. When $\bar{x}_n > \mu_0$. the solution is $\hat{\mu} = \bar{x}_n$. When $\bar{x}_n \leq \mu_0$, the problem is to minimize $(\bar{x}_n - \mu)^2$ under $\mu > \mu_0$. This minimum is not attained (μ can be selected arbitrarily close to μ_0, such that still $\mu > \mu_0$, which makes $(\bar{x}_n - \mu)^2$ arbitrarily close to $(\bar{x}_n - \mu_0)^2$, never attaining this value). However

$$\inf_{\mu > \mu_0} (\bar{x}_n - \mu)^2 = \min_{\mu \geq \mu_0} (\bar{x}_n - \mu)^2 = (\bar{x}_n - \mu_0)^2 .$$

Thus the MLE of μ under $\mu \geq \mu_0$ is $\hat{\mu}_1 = \max(\bar{x}_n, \mu_0)$. This is not the MLE under K, but gives the supremal value of the likelihood under K for given σ^2. To continue, we have to maximize in σ^2. Now

$$(\bar{x}_n - \hat{\mu}_1)^2 = (\min(0, \bar{x}_n - \mu_0))^2$$

and defining

$$S_{n,1}^2 := S_n^2 + (\bar{x}_n - \hat{\mu}_1)^2$$

we obtain

$$\sup_{\mu > \mu_0, \sigma^2 > 0} p_{\mu,\sigma^2}(x) = \sup_{\sigma^2 > 0} \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{S_{n,1}^2}{2\sigma^2 n^{-1}}\right).$$

The maximization in σ^2 is now analogous to the argument for part (ii) of Proposition 7.9, (p. 81 handout). The maximizing value is $\hat{\sigma}_1^2 = S_{n,1}^2$ and the maximized likelihood (which is also the supremal likelihood under K) is

$$\sup_{\mu>\mu_0,\sigma^2>0} p_{\mu,\sigma^2}(x) = \frac{1}{(\hat{\sigma}_1^2)^{n/2}} \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2n^{-1}}\right).$$

Now under the hypothesis, since $\mu \le \mu_0$ is a closed interval, the MLE's can straightforwardly be found. An analogous argument to the one above gives

$$\begin{aligned}
\hat{\mu}_0 &= \min(\bar{x}_n, \mu_0), \\
\min_{\mu\le\mu_0} (\bar{x}_n - \mu)^2 &= (\bar{x}_n - \hat{\mu}_0)^2 = (\max(0, \bar{x}_n - \mu_0))^2, \\
\hat{\sigma}_0^2 &= S_{n,0}^2 \text{ where } S_{n,0}^2 := S_n^2 + (\bar{x}_n - \hat{\mu}_0)^2, \\
\sup_{\mu\le\mu_0,\sigma^2>0} p_{\mu,\sigma^2}(x) &= \frac{1}{(\hat{\sigma}_0^2)^{n/2}} \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2n^{-1}}\right).
\end{aligned}$$

Thus the likelihood ratio is

$$L(x) = \left(\frac{\hat{\sigma}_0^2}{\hat{\sigma}_1^2}\right)^{n/2} = \left(\frac{S_n^2 + (\bar{x}_n - \hat{\mu}_0)^2}{S_n^2 + (\bar{x}_n - \hat{\mu}_1)^2}\right)^{n/2}$$

Suppose first that the t-statistic $T_{\mu_0}(X)$ has values ≤ 0; this is equivalent to $\bar{x}_n \le \mu_0$. In this case $\hat{\mu}_0 = \bar{x}_n$, $\hat{\mu}_1 = \mu_0$, hence

$$\begin{aligned}
L(x) &= \left(\frac{S_n^2}{S_n^2 + (\bar{x}_n - \mu_0)^2}\right)^{n/2} = \left(\frac{1}{1 + (\bar{x}_n - \mu_0)^2 / S_n^2}\right)^{n/2} \\
&= \left(\frac{1}{1 + (T_{\mu_0}(X))^2/(n-1)}\right)^{n/2}.
\end{aligned}$$

Thus for nonpositive values of $T_{\mu_0}(X)$, the likelihood ratio $L(x)$ is a monotone decreasing function of the absolute value of $T_{\mu_0}(X)$, which means it is *monotone increasing in* $T_{\mu_0}(X)$, on values $T_{\mu_0}(X) \le 0$.

Consider now nonnegative values of $T_{\mu_0}(X)$: $T_{\mu_0}(X) \ge 0$. Then $\bar{x}_n \ge \mu_0$, hence $\hat{\mu}_0 = \mu_0$, $\hat{\mu}_1 = \bar{x}_n$ and

$$\begin{aligned}
L(x) &= \left(\frac{\hat{\sigma}_0^2}{\hat{\sigma}_1^2}\right)^{n/2} = \left(\frac{S_n^2 + (\bar{x}_n - \mu_0)^2}{S_n^2}\right)^{n/2} \\
&= \left(1 + (T_{\mu_0}(X))^2/(n-1)\right)^{n/2}.
\end{aligned}$$

Thus for values $T_{\mu_0}(X) \ge 0$, the likelihood ratio $L(x)$ is a monotone increasing function of the absolute value of $T_{\mu_0}(X)$, which means it is *monotone increasing in* $T_{\mu_0}(X)$, on these values $T_{\mu_0}(X) \ge 0$.

The two areas of values of $T_{\mu_0}(X)$ we considered do overlap (in $T_{\mu_0}(X) = 0$); and we showed that $L(x)$ is a monotone increasing function of $T_{\mu_0}(X)$ on both of these. Hence $L(x)$ is a monotone increasing function of $T_{\mu_0}(X)$.

Exercise P5 (F-test for equality of variances). Consider the two sample problem of exercise P2, but hypotheses $H : \sigma_1^2 = \sigma_2^2$ vs. $K : \sigma_1^2 \ne \sigma_2^2$.

i) (5%) Find the likelihood ratio test and show that it is equivalent to a test which rejects if the F-statistic (2) is outside a certain interval of form $[c^{-1}, c]$.

ii) (5%) Show that the c of i) can be chosen as the upper $\alpha/2$ quantile of the distribution $F_{r,r}$ for a a certain $r > 0$.

Solution. The likelihood is the product of the one for $X_1, \ldots, X_n$ and for $Y_1, \ldots, Y_n$. The means μ_1, μ_2 are nuisance parameters, they are estimated by their respective MLE under both hypothesis amd alternative. After this first stage, the likelihood for $X_1, \ldots, X_n$ becomes

$$\frac{1}{(2\pi\sigma_1^2)^{n/2}} \exp\left(-\frac{S_X^2}{2\sigma_1^2 n^{-1}}\right),$$

and the likelihood for $Y_1, \ldots, Y_n$ analogously

$$\frac{1}{(2\pi\sigma_2^2)^{n/2}} \exp\left(-\frac{S_Y^2}{2\sigma_2^2 n^{-1}}\right),.$$

To obtain the likelihood under K:$\sigma_1^2 \neq \sigma_2^2$, we first maximize in σ_1^2 and σ_2^2 without any restriction. This gives estimates

$$\hat{\sigma}_1^2 = S_X^2,\ \hat{\sigma}_2^2 = S_X^2.$$

These estimates are independent and have a continuous distribution (they each have a density, since the χ^2 has a density), hence they have a joint continuous distribution, and the event $\hat{\sigma}_1^2 = \hat{\sigma}_2^2$ has zero probability (for all values of the parameters). Hence with probability one, $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$ are the MLE under K, an the maximized likelihood is

$$\frac{1}{(\hat{\sigma}_1^2\hat{\sigma}_2^2)^{n/2}} \frac{1}{(2\pi)^n} \exp\left(-\frac{1}{n^{-1}}\right).$$

Now under $H : \sigma_1^2 = \sigma_2^2$ (call the common value σ^2) the joint likelihood of $X_1, \ldots, X_n$, $Y_1, \ldots, Y_n$ is

$$\frac{1}{(2\pi\sigma^2)^n} \exp\left(-\frac{(S_X^2 + S_Y^2)/2}{2\sigma^2(2n)^{-1}}\right).$$

This gives an MLE

$$\hat{\sigma}^2 = \frac{1}{2}(S_X^2 + S_Y^2)$$

and a maximized likelihood under H

$$\frac{1}{(\hat{\sigma}^2)^n} \frac{1}{(2\pi)^n} \exp\left(-\frac{1}{n^{-1}}\right).$$

Thus the likelihood ratio is

$$\begin{aligned} L(X,Y) &= \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_1\hat{\sigma}_2}\right)^n = \left(\frac{(S_X^2 + S_Y^2)/2}{(S_X^2)^{1/2}(S_Y^2)^{1/2}}\right)^n \\ &= \frac{1}{2^n}\left(\left(\frac{S_X^2}{S_Y^2}\right)^{1/2} + \left(\frac{S_Y^2}{S_X^2}\right)^{1/2}\right)^n \\ & \quad \frac{1}{2^n}\left((F(X,Y))^{1/2} + (F(X,Y))^{-1/2}\right)^n. \end{aligned}$$

Consider the function $g(t) = t^{1/2} + t^{-1/2}$, for $t > 0$. . The LR test rejects when either $F(X,Y)$ is too large ($(F(X,Y))^{1/2}$ large) or too small ($(F(X,Y))^{-1/2}$ large).The inequality $L(X,Y) > c_0$ is equivalent to

$$g(F(X,Y)) > c^*$$

The function g has the property $g(t) = g(t^{-1})$, thus the set $\{t : g(t) > c^*\}$ has the property that if u is in this set then also u^{-1} is in it. Hence (noting also that g is monotone increasing for $t > 1$) there must be a c such that

$$\{t : g(t) \leq c^*\} = [c^{-1}, c].$$

Then rejection is equivalent to

$$F(X,Y) > c \text{ or } F(X,Y) < c^{-1}$$

which proves **i)** .

(ii) Under H, $F(X,Y)$ has a distribution $F_{n-1,n-1}$ (P2 (i)), and this distribution has the property that $(F(X,Y))^{-1}$ has the same distribution (this follows immediately from the definition of the F-distribution). Hence selecting $c = F_{n-1,n-1;1-\alpha/2}$ (lower $1-\alpha/2$-quantile) ensures that under H

$$\begin{aligned} P(F(X,Y) &> c) = \alpha/2, \\ P(F(X,Y) &< c^{-1}) = P((F(X,Y))^{-1} > c) = \alpha/2 \end{aligned}$$

which proves the claim.