Statistics 582, Problem Set 9 Solutions

Wellner; 3/9/2011

1. Consider the Locally Most Powerful test ϕ for testing $H : \theta \le 0 \equiv \theta_0$ versus $K : \theta > 0 = \theta_0$ in Example 6.1.5.
(a) Suggest two different approximations to the power of this test, one for local alternatives (of the form $\theta_n = t/\sqrt{n}$ with $t > 0$), and the other for fixed alternatives, $\theta > 0$.
(b) What is the behavior of each of these two approximations for large values of θ? Which of them shows that the power function decreases to 0 as $\theta \to \infty$? Why?
(c) Find a test ϕ of H versus K which has monotone increasing power function $\beta_\phi(\theta)$.

Solution: (a) The test is "reject H if $\sqrt{n}\,\overline{Y}_n > 2^{-1/2} z_\alpha$" where $Y_i \equiv 2X_i/(1+X_i^2)$ are i.i.d. and $X_i \sim$ Cauchy $(\theta, 1)$. Thus under P_θ, by using contour integration and Cauchy's formula, or by using Mathematica, Maple, or your favorite symbolic manipulation program,

$$
\begin{aligned}
m(\theta) \equiv E_\theta Y_i &= \int_{-\infty}^{\infty} \frac{2x}{1+x^2} p_\theta(x) dx = \int_{-\infty}^{\infty} \frac{2x}{1+x^2} \frac{1}{\pi} \frac{1}{1+(x-\theta)^2} dx \\
&= \frac{2\theta}{4+\theta^2},
\end{aligned}
$$

and

$$
\begin{aligned}
\sigma^2(\theta) &\equiv Var_\theta(Y_i) = E_\theta Y_i^2 - m^2(\theta) = \frac{2(4+3\theta^2)}{(4+\theta^2)^2} - \left(\frac{2\theta}{4+\theta^2}\right)^2 \\
&= \frac{2}{4+\theta^2}\,.
\end{aligned}
$$

For local alternatives $\theta = \theta_n = t/\sqrt{n}$, we have

$$
\begin{aligned}
\text{Power}(\theta_n) &= P_{\theta_n}(\sqrt{n}\,\overline{Y}_n > 2^{-1/2} z_\alpha) \\
&= P_{\theta_n}(\sqrt{n}(\overline{Y}_n - m(\theta_n)) \ge 2^{-1/2} z_\alpha - \sqrt{n}(m(\theta_n) - m(0))) \\
&\to P(2^{-1/2} Z \ge 2^{-1/2} z_\alpha - m'(0)t)
\end{aligned}
$$

where

$$
\begin{aligned}
m'(0) &= \int_{-\infty}^{\infty} \frac{2x}{1+x^2} \frac{d}{d\theta} p_\theta(x)|_{\theta=0} dx \\
&= \int_{-\infty}^{\infty} \dot{l}_\theta(x;0) \dot{l}_\theta(x;0) p_\theta(x;0) dx = I(\theta) = 1/2\,.
\end{aligned}
$$

Hence we have

$$\text{Power}(\theta_n) \to P(Z > z_\alpha - 2^{-1/2}t) = 1 - \Phi(z_\alpha - 2^{-1/2}t)\,.$$

This approximation to the power function increases monotonically from α at $t = 0$ to 1 at $t = \infty$ (effectively when $t > 2^{1/2} \cdot 4$). Note that this result is very much in qualitative agreement with corollary 4.2.4 from Statistics 581.

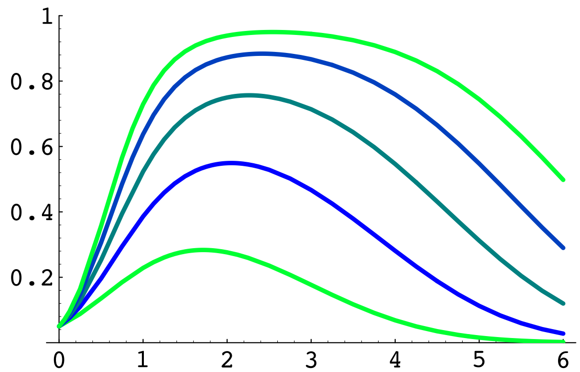

Figure 1: Plots of fixed θ power approximations for $n = 3, 6, 9, 12, 15$.

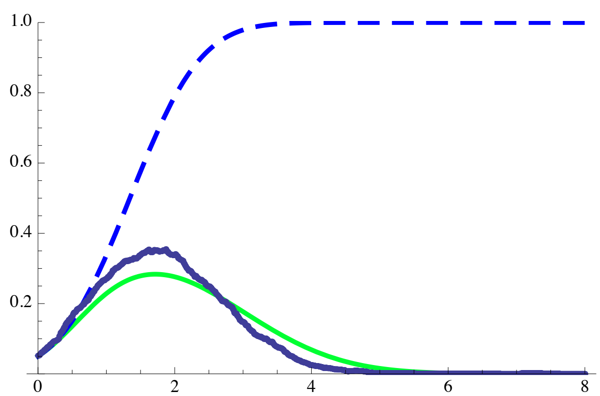

Figure 2: Local (dashed) and fixed (green) θ power approximations, $n = 3$; Monte Carlo estimate of true power (blue)

(b) For fixed alternatives $\theta > 0$ we have

$$\begin{aligned}
\text{Power}(\theta) &= P_\theta(\sqrt{n}\,\overline{Y} > 2^{-1/2}z_\alpha) \\
&= P_\theta(\sqrt{n}(\overline{Y}_n - m(\theta)) > 2^{-1/2}z_\alpha - \sqrt{n}m(\theta)) \\
&\doteq P(Z > (2^{-1/2}z_\alpha - \sqrt{n}m(\theta))/\sigma(\theta)) \\
&= 1 - \Phi((2^{-1/2}z_\alpha - \sqrt{n}m(\theta))/\sigma(\theta))\,.
\end{aligned}$$

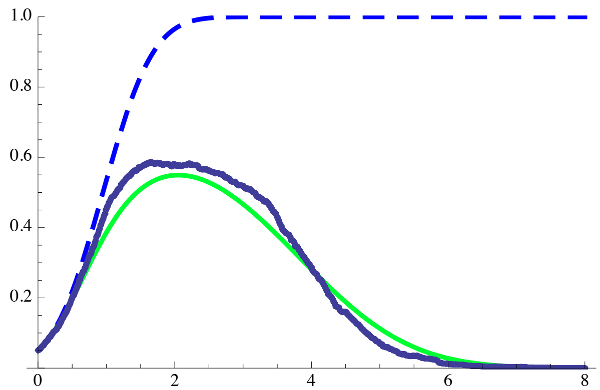

Figure 3: Local (dashed) and fixed (green) θ power approximations, $n = 6$; Monte Carlo estimate of true power (blue)

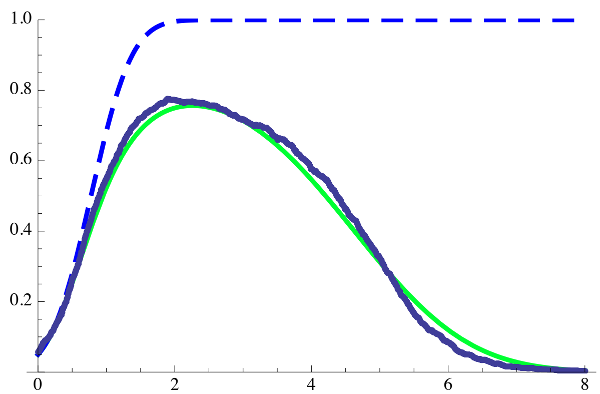

Figure 4: Local (dashed) and fixed (green) θ power approximations, $n = 9$; Monte Carlo estimate of true power (blue)

where

$$\frac{m(\theta)}{\sigma(\theta)} = \frac{\sqrt{2\theta^2}}{\sqrt{4+\theta^2}}$$

which increases from 0 at $\theta = 0$ to a maximum of $1/\sqrt{2}$ at $\theta = 2$, and then tends to $\sqrt{2}$ as θ increases to ∞. Thus our fixed alternative approximation to the power function is completely determined by the function $(2^{-1/2}z_\alpha - \sqrt{n}m(\theta))/\sigma(\theta)$, and this first decreases as θ increases (as it should if the power is to increase), but then it reaches a minimum and increases thereafter (approximately as $2^{-1}z_\alpha\theta - \sqrt{2n}$ as $\theta \to \infty$). Hence this approximation to the power decreases to zero just as we argued that it must in class. See Figures **??**-6 which give plots of the two approximations for $\alpha = .05$ and $n = 3, 6, 9, 12, 15$, together with Monte-Carlo estimates of the true power function baed on $m = 1000$ Monte-Carlo replications

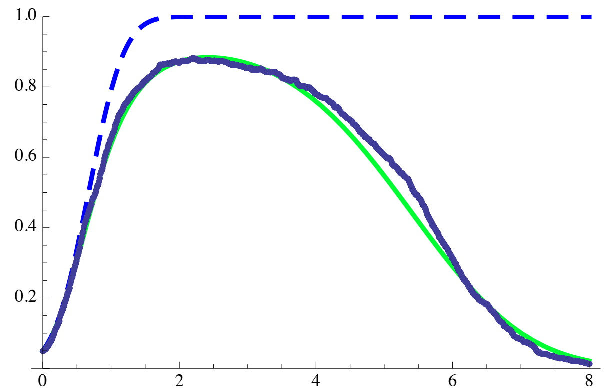

Figure 5: Local (dashed) and fixed (green) θ power approximations, $n = 12$; Monte Carlo estimate of true power (blue)

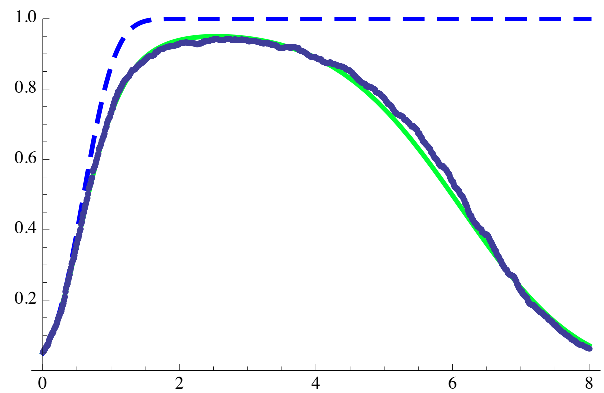

Figure 6: Local (dashed) and fixed (green) θ power approximations, $n = 15$; Monte Carlo estimate of true power (blue)

at each sample size (dark blue). Note that the two approximations agree for θ's close to 0, but the local approximation is always monotone increasing, while the approximations with θ fixed show the approximate power decreasing to 0 as $\theta \to \infty$ as we know it must. Also note that fixed alternative approximation to the power function based on the CLT is reasonably accurate for $n \geq 9$.

(c) Consider $\overline{X}_n \equiv n^{-1} \sum_{i=1}^n X_i$. Then, since $X_i \stackrel{d}{=} V_i + \theta$ where V_i are i.i.d. standard Cauchy random variables with $Ee^{itV_1} = \exp(-|t|)$ (see e.g. Shorack, *Probability for Statisticians*, page 343), it follows that

$$\begin{aligned}
\phi_{\overline{X}_n}(t) &= E\exp(itn^{-1}\sum_1^n X_i) = \phi_{X_1}(t/n)^n = \big(E\exp(itn^{-1}X_1)\big)^n \\
&= \big(E\exp(itn^{-1}(V_1+\theta))\big)^n = \phi_{V_1}(t/n)^n \cdot e^{it\theta} = \\
&= \exp(-|t|)e^{it\theta} = \phi_{X_1}(t);
\end{aligned}$$

i.e. $\overline{X}_n \stackrel{d}{=} X_1$. Thus the test $\phi(\underline{X}) = 1\{\overline{X}_n > k\}$ has

$$\alpha = E_0\phi(\underline{X}) = P_0(\overline{X}_n > k = P_0(X_1 > k)$$

if $k = \tan((1/2-\alpha)\pi) \equiv k_\alpha$, and then

$$\begin{aligned}
\beta_\phi(\theta) &= P_\theta(\overline{X}_n > k_\alpha) = P_\theta(X_1 - \theta > k_\alpha - \theta) \\
&= P_0(V_1 > k_\alpha - \theta) = \int_{k_\alpha-\theta}^\infty \frac{1}{\pi}\frac{1}{1+y^2}dy
\end{aligned}$$

which is monotone increasing as a function of θ; see Figure 7. On the other hand, this test does not yield increasing power as the sample size n increases: it has the same power function for all n!

Here is another possibility: Let $Z_i \equiv 1\{X_i > 0\}$ for $i = 1, \ldots, n$. Thus the Z_i's are independent Bernoulli $p = p_\theta$ random variables with

$$\begin{aligned}
p_\theta &= P_\theta(X_i > 0) = P_\theta(X_i - \theta > -\theta) = P_0(Y_i > -\theta) = \int_{-\theta}^\infty \frac{1}{\pi}\frac{1}{1+y^2}dy \\
&\nearrow 1 \ \ \text{as } \theta \to \infty.
\end{aligned}$$

Note that testing $H : \theta \le \theta_0 = 0$ versus $K : \theta > 0$ is equivalent to testing $H' : p \le p_0 \equiv 1/2$ versus $K' : p > p_0 = 1/2$. Now

$$T_n \equiv \sum_1^n Z_i = n(1 - \mathbb{F}_n(0)) \sim \text{Binomial}(n, p)$$

is a one-parameter exponential family with density function $p(\cdot; p)$ with respect to counting measure on $\mathbb{N}$ given by

$$\begin{aligned}
p(y;p) &= P_p(T_n = y) = \binom{n}{y} p^y (1-p)^{n-y} = \binom{n}{y}\left(\frac{p}{1-p}\right)^y (1-p)^n \\
&= c(p)\exp(Q(p)y)h(y)
\end{aligned}$$

with $Q(p) \equiv \log(p/(1-p))$, $c(p) \equiv (1-p)^n$, and $h(y) = \binom{n}{y}$, so it has monotone likelihood ratio in y. Thus by the Karlin-Rubin theorem, the test

$$\varphi(\underline{X}) = \phi(\underline{Z}) = \begin{cases} 1 & \text{if } T_n = \sum_1^n Z_i > k \\ \gamma & \text{if } T_n = k \\ 0 & \text{if } T_n < k, \end{cases}$$

where k and γ are chosen so that $E_{p_0=1/2}\phi(\underline{Z}) = \alpha$, has power function

$$\beta_\phi(p_\theta) \equiv \tilde{\beta}_\varphi(\theta)$$

which is monotone increasing in p_θ and θ. This test improves with increasing sample size n; see Figure 7

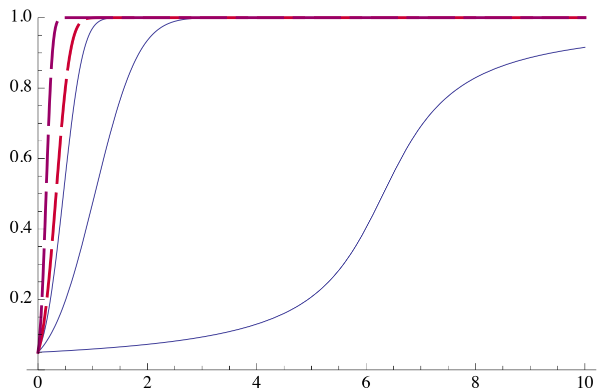

Figure 7: Plot comparing local power approximations for the Binomial test (solid), locally most powerful test (dashed), and the mean test, for $n = 10,\ 50$.

2. Let X and Y be random variables with joint density

$$p_{X,Y}(x,y) = \lambda\mu\exp(-\lambda x - \mu y)1_{(0,\infty)}(x)1_{(0,\infty)}(y).$$

(a) Find a UMP unbiased test of size $\alpha = .2$ for testing $H_0 : \lambda \le \mu + 1$ versus $H_1 : \lambda > \mu + 1$.
(b) Find a UMP unbiased test of size $\alpha = .2$ for testing $H_0 : \lambda = \mu$ versus $H_1 : \lambda \neq \mu$.
(c) Find a UMP unbiased test of size $\alpha = .2$ for testing $H_0 : \lambda \ge 2\mu$ versus $H_1 : \lambda < 2\mu$.
(d) What happens when $X_1, \ldots, X_m$ are i.i.d. Exponential(λ) and $Y_1, \ldots, Y_n$ are i.i.d. Exponential(μ)?

Solution: When $X \sim \exp(\lambda)$ and $Y \sim \exp(\mu)$ we have

$$p_{\lambda,\mu}(x,y) = \lambda\mu\exp(-\lambda x - \mu y)1_{(0,\infty)}(x)1_{(0,\infty)}(y)\,.$$

(a) For testing $H : \lambda \le \mu + 1$ versus $K : \lambda > \mu + 1$ we rewrite the density as follows:

$$\begin{aligned} p_{\lambda,\mu}(x,y) &= \lambda\mu\exp(-\lambda x - \mu y) \\ &= \lambda\mu\exp((\lambda-\mu)y - \lambda(x+y)) \\ &= \lambda\mu\exp(\theta U(x,y) + \xi T(x,y)) \end{aligned}$$

where $\theta \equiv \lambda - \mu$, $U(x,y) \equiv y$, $\xi \equiv -\lambda$, and $T(x,y) \equiv x+y$. Since $\lambda \le \mu + 1$ is equivalent to $\lambda - \mu = \theta \le 1 \equiv \theta_0$, our theory for exponential families applies, and the UMP unbiased test of H versus K is given by

$$\phi(X,Y) = \begin{cases} 1 & \text{if } Y > c_\alpha(T) \\ \gamma(T) & \text{if } Y = c_\alpha(T) \\ 0 & \text{if } Y < c_\alpha(T) \end{cases}$$

where c_α and $\gamma(\alpha)$ satisfy $E\{\phi(X,Y)|T\} = \alpha$. In this case, the conditional distribution of Y given $T = X+Y$ on the boundary $\Theta_B = \{(\lambda-1,\lambda) : \lambda \ge 1\}$ is given by

$$f_{Y|T}(y|t) = \frac{e^y}{e^t - 1} 1_{[0,t]}(y)\,.$$

Therefore $1 - F_{Y|T}(y|t) = 1 - (e^y - 1)/(e^t - 1)$ and for $\alpha = .2$ the critical point for the conditional test is given by

$$c_\alpha(T) = \log(\exp(T) - (\exp(T)-1)/5) = \log((4/5)\exp(T) + 1/5)\,, \qquad \gamma(T) = 0\,.$$

(b) For testing $H : \lambda = \mu$ versus $K : \lambda \ne \mu$, the same rewrite of the density as in (a) works. Now we have $\lambda = \mu$ is equivalent to $\mu - \lambda = 0 \equiv \theta_0$, and $\lambda \ne \mu$ is equivalent to $\mu - \lambda \ne 0 \equiv \theta_0$, so our theory for exponential families applies, and the UMP unbiased test of H versus K is given by

$$\phi(X,Y) = \begin{cases} 1 & \text{if } Y > c_2(T) \text{ or } Y < c_1(T) \\ \gamma_i(T) & \text{if } Y = c_1(T) \text{ or } Y = c_2(T) \\ 0 & \text{if } c_1(T) < Y < c_2(T) \end{cases}$$

where the c_1, c_2, γ_1 and γ_2 are determined so that $E\{\phi(X,Y)|T\} = \alpha$. In this case the conditional distribution of Y given T on $\Theta_B = \{(\lambda,\lambda) : \lambda \ge 0\}$ is Uniform$(0,T)$, and hence the conditional distribution of Y/T given T is Uniform$(0,1)$, and this is independent of T. Hence the UMPU test of H versus K of size .2 is given by "reject H if $Y/T < .1$ or $Y/T > .9$".
(c) For testing $H : \lambda \ge 2\mu$ versus $K : \lambda < 2\mu$, we need a somewhat different rewrite of the joint density. Now

$$\begin{aligned} p_{\lambda,\mu}(x,y) &= \lambda\mu\exp(-\lambda x - \mu y) \\ &= \lambda\mu\exp(-(\lambda - 2\mu)x - \mu(2x+y)) \\ &= \lambda\mu\exp(\theta U(x,y) + \xi T(x,y)) \end{aligned}$$

where $\theta \equiv 2\mu - \lambda$, $U(x,y) \equiv x$, $\xi \equiv -\mu$, and $T(x,y) \equiv 2x + y$. Since $\lambda \ge 2\mu$ is equivalent to $2\mu - \lambda \equiv \theta \le 0 \equiv \theta_0$, (and $\lambda < 2\mu$ is equivalent to $2\mu - \lambda = \theta >$

$0 \equiv \theta_0$), our theory for exponential families applies, and the UMP unbiased test of H versus K is given by

$$\phi(X,Y) = \begin{cases} 1 & \text{if} \quad X > c_\alpha(T) \\ \gamma(T) & \text{if} \quad X = c_\alpha(T) \\ 0 & \text{if} \quad X < c_\alpha(T) \end{cases}$$

where $c_(T)$ and $\gamma(T)$ satisfy $E\{\phi(X,Y)|T\} = \alpha$. In this case the conditional distribution of X given T is Uniform$(0, T/2)$, so $2X/T$ is Uniform$(0,1)$ and independent of T. Hence the UMPU test of H versus K of size $\alpha = .2$ is given by "reject H if $2X/T > .8$".

When we observe $X_1, \ldots, X_m$ are i.i.d. Exponential(λ) and $Y_1, \ldots, Y_n$ are i.i.d. Exponential(μ), then the distribution of the observations is given by

$$\begin{aligned} p_{\lambda,\mu}(\underline{x}, \underline{y}) &= \lambda^m \mu^n \exp(-\lambda \sum_{i=1}^m x_i - \mu \sum_{j=1}^n y_j) \\ &= \lambda^m \mu^n \exp((\lambda - \mu) \sum_{j=1}^n y_j - \lambda(\sum_{i=1}^m x_i + \sum_{j=1}^n y_j)) \\ &= \lambda^m \mu^n \exp(\theta U(\underline{x}, \underline{y}) + \xi T(\underline{x}, \underline{y})) \end{aligned}$$

where $\theta \equiv \lambda - \mu$, $U(\underline{x}, \underline{y}) \equiv \sum y_j$, $\xi \equiv -\lambda$, and $T(\underline{x}, \underline{y}) \equiv \sum x_i + \sum y_j)$. This rewrite works for (a) and (b), and a similar rewrite works for (c) with the new $U = \sum x_i$, $T = 2\sum x_i + \sum y_j$. The form of the tests in (a) - (c) remains the same with the new U and T, and all that remains is to calculate the conditional distributions of U given T. In (a) this density is given by

$$f_{U|T}(u|t) = \frac{u^{n-1}(t-u)^{m-1}e^u}{\int_0^t v^{n-1}(t-v)^{m-1}e^v dv} \, .$$

In (b) it is easily found that $V \equiv U/T \sim \text{Beta}(n,m)$, and the test can be carried out unconditionally using tables of the Beta distributions. In (c) $2U \equiv \sum 2X_i \sim$ Gamma(m, μ) and $V \equiv \sum Y_j \sim$ Gamma(n, μ) on the boundary $\lambda = 2\mu$, so $2\mu U \sim \chi^2_{2m}$ and $\mu V \sim \chi^2_{2n}$. Therefore

$$\frac{2U/(2m)}{V/(2n)} = \frac{n}{m} 2 \frac{U}{V} \sim F_{2m,2n}$$

and since $2U/T = 2U/(2U+V) = (2U/V)(1+(2U/V))$ is a monotone increasing function of $2U/V$, the UMPU test can be carried out unconditionally using tables of the $F_{2m,2n}$ distributions.

3. Lehmann and Romano, TSH, problem 4.3, page 139: Let X have the binomial distribution $Bin(n, p_)$ and consider the hypothesis $H : p = p_0$ at level of significance α. Determine the boundary values of the UMP unbiased test for $n = 10$ with $\alpha = .1$, $p_0 = .2$ and with $\alpha = .05$, $p_0 = .4$, and in each case graph the power functions of both the unbiased and the equal-tails test.

   Solution: For the first scenario: $n = 10$, $\alpha = .1$, $p_0 = .2$, we find, using the calculations in Lehmann and Romano, TSH, page 113, the UMP unbiased test is $\phi(x) = \gamma_1 1\{x = 0\} + \gamma_2 1\{x = 4\} + 1\{x \geq 5\}$ with $c_1 = 0$, $c_2 = 4$, $\gamma_1 = .5590$, and $\gamma_2 = .0815$. The corresponding equal tails test has $c_1 = 0$, $c_2 = 4$, $\gamma_1 = .5/.107374 = .4657$, and $\gamma_2 = 0.19535$. The power functions of these two tests are shown in Figure 8.

   For the second scenario: $n = 10$, $\alpha = .05$, $p_0 = .4$, we find, using the calculations in Lehmann and Romano, TSH, page 113, the UMP unbiased test is $\phi(x) = \gamma_1 1\{x = 1\} + \gamma_2 1\{x = 7\} + 1\{x \geq 8\}$ with $c_1 = 1$, $c_2 = 7$, $\gamma_1 = 0.5034$, and $\gamma_2 = 0.2677$. The corresponding equal tails test has $c_1 = 1$, $c_2 = 7$, $\gamma_1 = 0.4702$, and $\gamma_2 = 0.2992$. The power functions of these two tests are shown in Figure 9, and do not differ substantially.

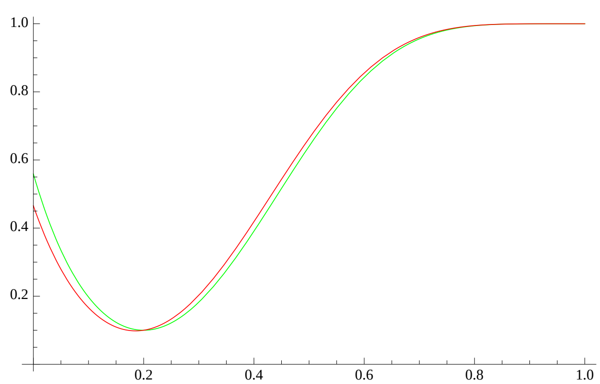

Figure 8: Power functions, UMP unbiased test (green) and equal tails test (red): $n = 10$, $p_0 = .2$, $\alpha = .10$

4. (From Wasserman, *All of Statistics*, page 171.) In 1961, 10 essays appeared in the *New Orleans Daily Crescent*. They were signed "Quintus Curtius Snodgrass" and some people suspected they were actually written by Mark Twain. To investigate this, we will consider the proportion of three letter words founds in an author's work. From eight Twain essays we have

   .225, .262, .217, .240, .230, .229, .235, .217

   From 10 Snodgrass essays we have:

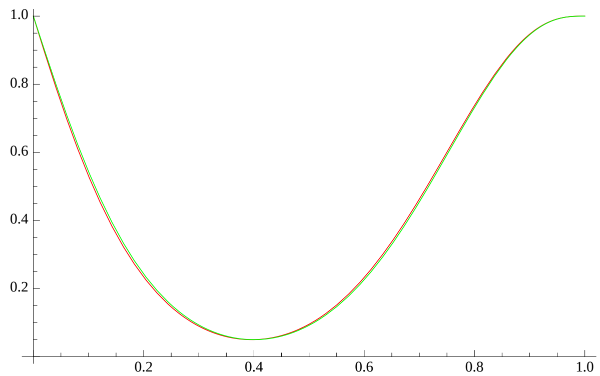

Figure 9: Power functions, UMP unbiased test (green) and equal tails test (red): $n = 10$, $p_0 = .4$, $\alpha = .05$

$$.209, \quad .205, \quad .196, \quad .210, \quad .202, \quad .207, \quad .224, \quad .223, \quad .220, \quad .201$$

(a) Perform a Wald test for equality of the means. Give a $p-$value and a 95% confidence interval for the difference of means. What conclusion do you reach?
(b) Now use a permutation test to avoid the use of large - sample methods. What is your conclusion?

Solution: (a) Labelling the Twain proportions as X's and the Snodgrass proportions as Y's, we find that $\overline{X}_m = .231875$, $\overline{Y}_n = .2097$, $S_X = .01456$, and $S_Y = .00966$. Assuming that $X_i \sim N(\mu, \sigma^2)$ and $Y_j \sim N(\nu, \tau^2)$ with $\sigma \neq \tau$, the Wald statistic becomes

$$\begin{aligned}
W_{m,n} &= \left\{ \frac{\sqrt{\frac{mn}{N}}(\overline{X}_m - \overline{Y}_n)}{\sqrt{(n/N)S_X^2 + (m/N)S_Y^2}} \right\}^2 \\
&= \left\{ \frac{\sqrt{\frac{8 \cdot 10}{18}}(.231875 - .2097)}{\sqrt{(10/18)(.000212125) + (8/18)(.0000933444)}} \right\}^2 \\
&= 3.70355^2 = 13.7163
\end{aligned}$$

and the (approximate) p-value is $P(\chi_1^2 > 13.7163) = .000213$. If we use Welch's approximate $t-$test (see e.g. Lehmann and Casella, TSH, page 447), then the degrees of freedom f becomes, with $R \equiv mS_X^2/(nS_Y^2)$

$$\frac{1}{f} = \left(\frac{R}{1+R}\right)^2 \frac{1}{m-1} + \frac{1}{(1+R)^2} \frac{1}{n-1} = 1/13.6148.$$

Thus the approximate p-value using Welch's approximation is $P(|t_{13.61}| \geq 13.7163) =$

.00265. A 95% confidence interval for $\mu - \nu$ based on normal theory is given by

$$\begin{aligned}
&\overline{X}_m - \overline{Y}_n \pm z_{.025}\sqrt{S_X^2/m + S_Y^2/n} \\
&= .022175 \pm 1.95996\sqrt{.000212125/8 + .0000933444/10} \\
&= (0.0104397, 0.0339103)
\end{aligned}$$

The conclusion based on either of these tests is to reject the null hypothesis: from this evidence we would conclude that the Snodgrass and Twain essays were written by different authors.

(b) If we do an exact permutation t-test using the statistic introduced in class (involving the assumption of equal variances in the alternative), there are $\binom{18}{8} =$ 43758 combinations to consider, and the observed value of the statistic (in the form $(\overline{X}_m - \overline{z})/\sigma_N$) is 2.78917. By my calculations the exact one-sided p-value is 0.000525618, and the exact two-sided p-value is 0.000777. In contrast, by drawing $10^5 = 100,000$ random permutations, the estimated p-values were 0.00058 and 0.00088 respectively.

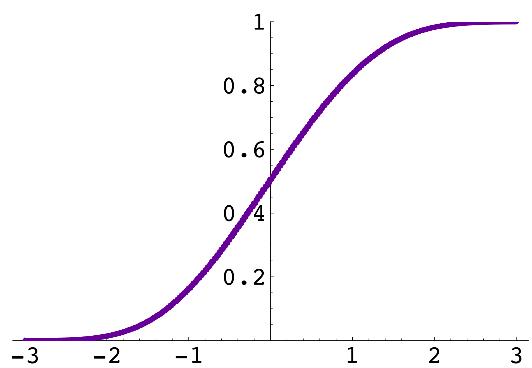

Figure 10: Exact permutation distribution, two-sample $t-$statistic $(\overline{X} - \overline{z})/\sigma_N$

I have not yet programmed the exact permutation test based on the Wald - type statistic used in part (a), but without squaring, which allows for the possibility of different variances. There are still $\binom{18}{8} = 43758$ combinations to consider, and the observed value of the test statistic is 3.70355. I have however, programmed the approximate permutation test based on sampling from the permutation distribution. By drawing $10^5 = 100,000$ random permutations, the estimated p-values I calculated are 0.00461 and 0.01064 respectively. It seems that the permutation distribution of the unequal variances version of the unsquared form of the Wald statistic is more nearly normal than that of the classical $t-$statistic. Note that while the two-sided permutation test still rejects at level $\alpha = .05$, this two-sided p-valued (0.01064)is not nearly as small as the estimated p-value of the

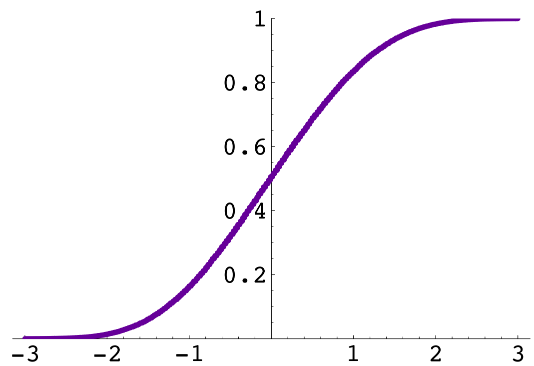

Figure 11: Approximate permutation distribution, two-sample $t-$statistic $(\overline{X}-\overline{z})/\sigma_N$, based on 10^5 random permutations

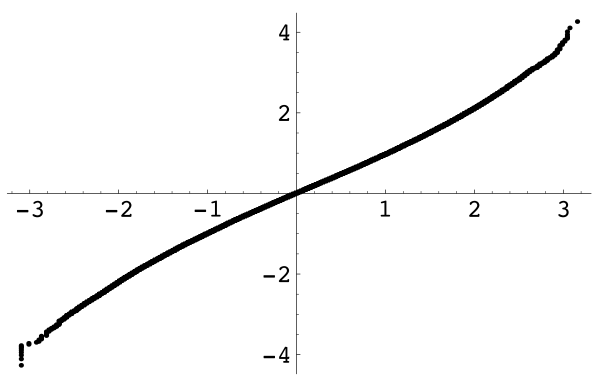

Figure 12: QQ-plot, Approximate permutation distribution, two-sample $t-$statistic $(\overline{X}-\overline{z})/\sigma_N$, based on 10^5 random permutations

permutation $t-$ test noted above (0.00088). We would continue to reject the null hypothesis at the level 0.05, but not at 0.01.

5. For observations $\underline{X} = (X_1, \ldots, X_n)$, let $X_{(1)} \leq \ldots \leq X_{(n)}$ denote the *order statistics* of the X_i's ($X_{(i)} \equiv \mathbb{F}_n^{-1}(i/n)$, $i = 1, \ldots, n$) and let $\underline{R} = (R_1, \ldots, R_n)$ denote the *ranks*; defined by $X_i = X_{(R_i)}$, $i = 1, \ldots, n$ (if $X_i = X_j$ for some $i < j$, define the ranks by $R_i < R_j$ and $X_i = X_{(R_i)}$).
(a) Suppose that $X_1, \ldots, X_n$ are i.i.d. $F \in \mathcal{F}_{ac}$ (the absolutely continuous df's F on R) with density f. Show that the order statistics $\underline{X}_{(\cdot)} \equiv (X_{(1)}, \ldots, X_{(n)})$ are independent of the ranks $\underline{R}$ and that the order statistics have joint density $\bar{p}$ given by

$$\bar{p}(\underline{x}_{(\cdot)}) = n! \prod_{i=1}^{n} f(x_{(i)}), \qquad -\infty < x_{(1)} < \ldots < x_{(n)} < \infty$$

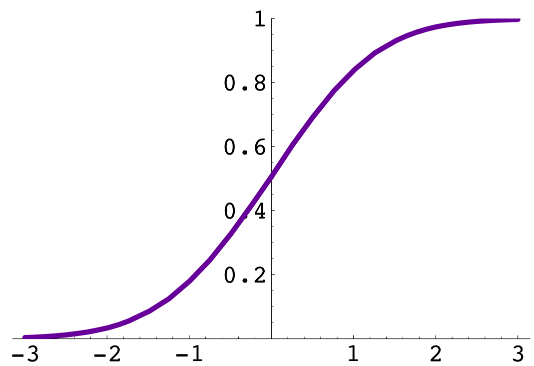

Figure 13: Approximate permutation distribution, one-sided Wald statistic $\overline{X} - \overline{Y})/\sqrt{S_X^2/m + S_Y^2/n}$ based on 10^5 random permutations

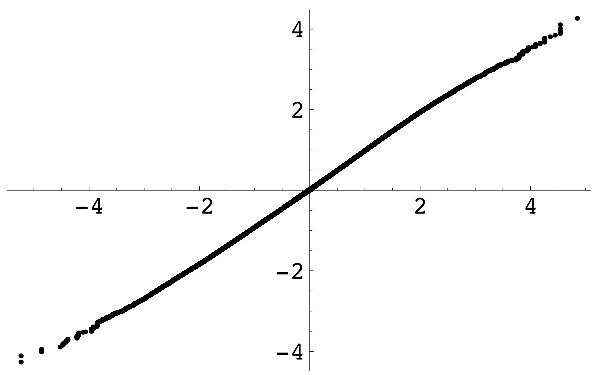

Figure 14: QQ-plot, Approximate permutation distribution, one-sided Wald statistic $\overline{X} - \overline{Y})/\sqrt{S_X^2/m + S_Y^2/n}$ based on 10^5 random permutations

while

$$P(\underline{R} = \underline{r}) = \frac{1}{n!}, \qquad \underline{r} \in \Pi \equiv \{ \text{ all permutations of } \{1, \ldots, n\}\} .$$

(b) Show that if the density f of the X_i's is log-concave, then the joint density $\bar{p}$ of the order statistics $\underline{X}_{(\cdot)}$ is log-concave; i.e. show that if $f((x+y)/2)^2 \ge f(x)f(y)$ for all $x, y \in \mathbb{R}$, then $\bar{p}((\underline{x}+\underline{y})/2)^2 \ge \bar{p}(\underline{x})\bar{p}(\underline{y})$ for all $\underline{x}, \underline{y} \in \mathcal{O}_n \equiv \{\underline{x} \in \mathbb{R}^n : \ x_1 \le x_2 \le \cdots \le x_n\}$.
(c) Show that (a) continues to hold for any joint distribution p of the $\underline{X}$ which is symmetric with respect to permutation of its coordinates: $p(\pi\underline{x}) = p(\underline{x})$ for all $\underline{x}$ and $\pi \in \Pi$ where $\pi x \equiv (x_{\pi(1)}, \ldots, x_{\pi(n)})$.
(d) If the joint distribution p of $\underline{X}$ is general (not permutation symmetric), show that the joint density $\bar{p}$ of the order statistics is given by

$$\bar{p}(\underline{x}_{(\cdot)}) = \sum_{\pi \in \Pi} p(\pi \underline{x}_{(\cdot)}) ,$$

and

$$P(\underline{R} = \underline{r} | \underline{X}_{(\cdot)} = \underline{x}_{(\cdot)}) = \frac{p(r\underline{x}_{(\cdot)})}{\bar{p}(\underline{x}_{(\cdot)})} \ .$$

Solution: I will prove (c) first; then (a) and (b) follow as corollaries:
(c) Suppose that $\underline{X}$ has joint density p. Then for any set Borel set $A \subset \{\underline{x} \in \mathbb{R}^n : \ x_1 < x_2 < \cdots < x_n\}$

$$\begin{aligned}
P(\underline{X}_{(\cdot)} \in A) &= \int_{[\underline{x}_{(\cdot)} \in A]} p(x_1, \ldots, x_n) \ dx_1 \ldots dx_n \\
&= \sum_{r \in \Pi} \int_{[R(\underline{x}) = r, \ \underline{x}_{(\cdot)} \in A]} p(x_1, \ldots, x_n) \ dx_1 \ldots dx_n \\
&= \sum_{r \in \Pi} \int_A p(x_{(r_1)}, \ldots, x_{(r_n)}) \ dx_{(1)} \ldots dx_{(n)} \\
&= \int_A \overline{p}(x_{(1)}, \ldots, x_{(n)}) \ dx_{(1)} \ldots dx_{(n)}
\end{aligned}$$

where we have used the fact that the correspondence between $(x_1, ..., x_n)$ and $(x_{(1)}, \ldots, x_{(n)})$ is one-to-one and linear with Jacobian $= 1$ on each subset $[R = r]$, $r \in \Pi$. This proves that

$$\overline{p}(\underline{x}_{(\cdot)}) = \sum_{\pi \in \Pi} p(\pi \underline{x}_{(\cdot)}) \ .$$

Similarly,

$$\begin{aligned}
P(R = r, \ \underline{X}_{(\cdot)} \in A) &= \int_{[R = r, \ \underline{x}_{(\cdot)} \in A]} p(x_1, \ldots, x_n) \ dx_1 \cdots dx_n \\
&= \int_A p(x_{(r_1)}, \ldots, x_{(r_n)}) \ dx_{(1)} \ldots dx_{(n)} \\
&= \int_A \frac{p(x_{(r_1)}, \ldots, x_{(r_n)})}{\overline{p}(x_{(1)}, \ldots, x_{(n)})} \overline{p}(x_{(1)}, \ldots, x_{(n)}) \ dx_{(1)} \ldots dx_{(n)}
\end{aligned}$$

since $\overline{p}(x_{(1)}, \ldots, x_{(n)}) = 0$ implies $p(x_{r(1)}, \ldots, x_{r(n)}) = 0$ for each $r \in \Pi$. This implies that

$$P(\underline{R} = \underline{r} | \underline{X}_{(\cdot)} = \underline{x}_{(\cdot)}) = \frac{p(r\underline{x}_{(\cdot)})}{\overline{p}(\underline{x}_{(\cdot)})} \ .$$

(b) When $p(\underline{x}) = p(\pi \underline{x})$ for all $\pi \in \Pi$, then

$$\overline{p}(\underline{x}_{(\cdot)}) = n! p(\underline{x}_{(\cdot)}),$$

and

$$P(\underline{R} = \underline{r} | \underline{X}_{(\cdot)} = \underline{x}_{(\cdot)}) = \frac{p(r\underline{x}_{(\cdot)})}{\overline{p}(\underline{x}_{(\cdot)})} = \frac{p(r\underline{x}_{(\cdot)})}{n!p(\underline{x}_{(\cdot)})} = \frac{1}{n!}.$$

Hence R is independent of $\underline{X}_{(\cdot)}$, and $P(R = r) = 1/n!$ for each $r \in \Pi$.
(a) This follows easily from (b) since, in this case, for any permutation π

$$p(\pi\underline{x}) = \prod_{i=1}^{n} f(x_{\pi(i)}) = \prod_{i=1}^{n} f(x_i) = p(\underline{x}).$$