


Exam Retrospective
 Psychology 311
 Midterm Exam
 February, 2013

Overall performance on this difficult exam was very good. The average was 22 out of 27.

Machine readable files were received from 15 out of 16 students, and were used to analyze performance on the various items. Here are the number correct for each item.



As you can see, performance overall was substantially better on the last 15 items than on the first 12.

Here are the numbers for items that anyone got wrong.

X11	X3	X27	X12	X10	X6	X7	X5	X8	X20	X25	X26	X1	X2	X4	X16	X17	X18	X19	X24
4	6	6	7	9	10	11	12	13	13	13	13	14	14	14	14	14	14	14	14

We'll take a look at the items that produced the poorest performance.

11. Suppose the A has a chi-square distribution with 14 degrees of freedom, and B has a chi-square distribution with 22 degrees of freedom. If A and B are independent, what is the standard deviation of the random variable $A + 2B$?
- (a) 7.616
 - (b) 10.77
 - (c) 14.283
 - (d) 204
 - (e) 408

Several facts that are required to solve this item:

A chi-square variables with ν degrees of freedom has a variance of 2ν .

If A and B are independent, by the heuristic rule, the random variable $A + 2B$ has a variance of

$$\text{Var}(A + 2B) = \sigma_A^2 + 4\sigma_B^2 \quad (1.1)$$

A must have a variance of $2 \times 14 = 28$, and B must have a variance of $2 \times 22 = 44$. So

$$\text{Var}(A + 2B) = 28 + 4(44) = 28 + 176 = 204 \quad (1.2)$$

Consequently, the random variable $A + 2B$ has a standard deviation of $\sqrt{204} = 14.283$.

The majority of people who got this item wrong chose alternative (b). This item was created as a deliberate distractor and had the incorrect formula

$$\sqrt{\sigma_A^2 + 2\sigma_B^2} \quad (1.3)$$

People who chose this alternative forgot, in the speed of the moment, to square the linear weight of 2.

You *could* have checked out your answer with R in less than a minute. Remember our random number generation routines? Lets create a million examples of $A + 2B$ in R.

```
> set.seed(12345)
> data<- rchisq(1000000,14)+2*rchisq(1000000,22)
> sd(data)
[1] 14.27338
```

This settles the issue completely.

3. Suppose you sample sets of 79 numbers independently and at random from a normal distribution with a mean of $\mu = 106$ and a standard deviation of $\sigma = 11$. For each set, you compute the sample mean \bar{Y} . As you do this over and over, the distribution of the sample means will be revealed to be normal, and 95% of the means will be lower than _____.

- (a) 108.426
- (b) 108.036
- (c) 107.586
- (d) 127.667
- (e) 106

Countless times in the course, we emphasized that the distribution of the sample mean based on n observations from a normal distribution with $\mu = 106$ is normal, with mean μ and standard deviation σ / \sqrt{n} . So in this case, we can say that \bar{Y} has a mean of 106 and a standard deviation of $11 / \sqrt{79}$. We want the 95th percentile of this distribution, which is computable directly in R as

```
> qnorm(0.95,106,11/sqrt(79))  
[1] 108.0357
```

So (b) is the correct choice. Alternative (d) was the most popular item wrongly chosen. I deliberately made it stand out from all the other numbers because I thought that would make it a good distractor.

27. In a 2-Way 3×3 ANOVA with $n = 12$ per cell, you observe an F statistic of 3.89 for the A main effect. Construct a 90% confidence interval for the RMSSE.

- (a) 0.0879, 0.5086
- (b) 0.1055, 0.6103
- (c) 0.6333, 1.054
- (d) 0.0765, 0.4425
- (e) 0.0976, 0.5645

This was a tough one. Start by getting a confidence interval for the noncentrality parameter λ . In order to do that, you need to know the degrees of freedom for the A main effect F statistic. They are $A - 1 = 3 - 1 = 2$, and $AB(n - 1) = 3 \times 3 \times (12 - 1) = 99$. So the confidence limits for λ are

```
> out <- conf.limits.ncf(F.value = 3.89, conf.level = .90,
+   df.1 = 2, df.2 = 99)
> lambda.limits <- c(out$Lower.Limit,out$Upper.Limit)
> lambda.limits
[1] 0.5561023 18.6208907
```

To convert these to limits on the RMSSE, we need to (carefully!) get the relationship between λ and the RMSSE. This relationship is given in the lecture slides as shown below.

<ul style="list-style-type: none"> Introduction An Introductory Example Basic ANOVA Analysis — Main Effects and Interaction Simple Main Effects Calculation The General Linear Model Calculating Population Model Coefficients Source Table and Expected Mean Squares The E(MS) Test Construction Principle Calculations: The Wiley-Voss Example Effect Size Estimation Power Calculations 	<ul style="list-style-type: none"> The RMSSE Partial ω^2 Cohen's f Key Relationships between Effect Size Measures Effect Size Confidence Intervals Effect Size Confidence Intervals
--	--

The RMSSE

- Steiger (2004) discusses several measures of effect size and their relationship to each other and the noncentrality parameter λ of the F statistic used for testing a particular effect θ .
- The Root Mean Square Standardized Effect (RMSSE) is

$$\Psi_{\theta} = \sqrt{\frac{\sum(\theta/\sigma)^2}{df_{\theta}}} \quad (12)$$

$$= \sqrt{\frac{\lambda_{\theta}}{n_{\theta}df_{\theta}}} \quad (13)$$

where is n_{θ} equal to n (the number of observations in each cell of the design) multiplied by the product of the numbers of levels in all the factors not represented in the effect currently under consideration; df_{θ} is the numerator degrees of freedom parameter for the effect under consideration.

James H. Steiger	2-Way Completely Randomized Design
------------------	------------------------------------

In the above formula, $n_{\theta} = 12 \times 3 = 36$ and $df_{\theta} = 2$, and so $\Psi_{\theta} = \sqrt{\lambda_{\theta} / 72}$.

Consequently, we can convert confidence intervals on λ into confidence intervals on the RMSSE by dividing them by 72 and taking the square root.

This is easy to do in R.

```
> sqrt(lambda.limits/72)
[1] 0.08788426 0.50855037
```

12. Consider the 2-sample, independent sample t test, and compare it to the 2-sample correlated sample t test. Which of the following is an assumption of the correlated sample t test that is actually a stronger (more restrictive) assumption than the assumptions of the independent sample test?
- (a) The assumption of equal variances
 - (b) The assumption of bivariate normality
 - (c) The assumption of independence of observations within samples
 - (d) The assumption of homoskedasticity
 - (e) none of the above

The correct answer is (b). Why is it correct? Because, for the *difference* between scores to have the normal distribution that is required for a 1-sample t test, the two measurements must have a bivariate normal distribution.

Let me comment on the distracters.

Alternative (a) cannot be correct because there is no assumption of equal variances in the correlated sample t . Indeed, there is no need for it. The two columns of numbers are reduced to one, and so the test is actually performed as a 1-sample t .

Alternative (c). Independence of observations is required of both tests.

Alternative (d). Homoskedasticity means “equal variances,” and so is equivalent to (a).

Alternative (e) is not correct because alternative (b) is correct.

10. If two sample means \bar{Y}_1 and \bar{Y}_2 are based on independent samples of size $n_1 = 20$ and $n_2 = 40$ from two normal distributions with means of 0 and variances of $\sigma_1^2 = 50$ and $\sigma_2^2 = 100$, what is the sampling variance of the quantity $\bar{Y}_1 - \bar{Y}_2$?

- (a) 83.621
- (b) 6
- (c) 4
- (d) 5
- (e) 12

By the heuristic rule, the variance of $\bar{Y}_1 - \bar{Y}_2$, assuming the samples are independent, is equal to

$$\sigma_{\bar{Y}_1}^2 + \sigma_{\bar{Y}_2}^2 = \frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2} = \frac{50}{20} + \frac{100}{40} = 5$$

By the way, you could easily simulate this in R. After a million samples, the choice of (d) is quite clear.

```
> set.seed(12345)
> data <- replicate(1000000, mean(rnorm(20, 0, sqrt(50))) -
mean(rnorm(40, 0, sqrt(100))))
> var(data)
[1] 4.991689
```

7. In a 1-Way ANOVA with 3 groups, and a sample size of $n = 12$ per cell, you observe an F -statistic of 4.47, with 2 and 33 degrees of freedom. On the basis of this statistic, the endpoints of a 90% confidence interval for Cohen's f are _____ and _____.
- (a) 0.148, 0.765
 - (b) 0.195, 0.935
 - (c) 0.229, 0.872
 - (d) 0.039, 0.693
 - (e) 0.155, 0.929

We begin by getting a confidence interval for λ , the noncentrality parameter. This may be done in R as follows:

```
> out <- conf.limits.ncf(4.47,0.90,2,33)
> lambda.limits <- c(out$Lower.Limit,out$Upper.Limit)
```

Since, for a 1-Way ANOVA,

$$f = \sqrt{\frac{\lambda}{na}} \quad (1.4)$$

we can transform the confidence interval

```
> n <- 12; a <- 3
> f.limits <- sqrt(lambda.limits/(n*a))
> f.limits
[1] 0.1479499 0.7647664
```

5. Now suppose that, in the previous problem, your t statistic of 3.20 actually represented a 2-sample Student's t test on the hypothesis that $\mu_1 = \mu_2$, based on two independent samples of 11 independent observations from two normally distributed populations. Construct a 90% confidence interval on the standardized effect size

$$E_s = \frac{\mu_1 - \mu_2}{\sigma}$$

- (a) 0.582, 2.246
- (b) 0.416, 2.286
- (c) 0.424, 2.407
- (d) 0.125, 2.583
- (e) 0.565, 2.135

The confidence limits for the noncentrality parameter δ are

```
> out <- conf.limits.nct(3.2,20,.90)
> delta.limits <- c(out$Lower.Limit,out$Upper.Limit)
> delta.limits
[1] 1.323981 5.005975
```

From the equation

$$E_s = \delta \sqrt{\frac{n_1 + n_2}{n_1 n_2}} \quad (1.5)$$

we get

```
> n.1 <- 11
> n.2 <- 11
> Es.limits <- delta.limits * sqrt((n.1 + n.2)/(n.1*n.2))
> Es.limits
[1] 0.5645473 2.1345551
```