Sociology 592 - Research Statistics I
Exam 1 Answer Key
September 26, 2003

Where appropriate, show your work - partial credit may be given. (On the other hand, don't waste a lot of time on excess verbiage.) Do not spend too much time on any one problem. It is legitimate (and probably essential) to refer to results that have previously been proven in class or homework, without re-proving them - for example, you wouldn't need to prove that P(-1.96 ≤ Z ≤ 1.96) = .95, since we have already shown that in class. Likewise, you are free to refer to anything that was demonstrated in the homework or handouts.

---

1. (4 points each, 20 points total). Indicate whether the following statements are true or false. If you think the statement is false, indicate how the statement could be corrected. For false statements, do not just say that you could substitute not equals for equals. For example, the statement P(Z ≤ 0) = .7 is false. To make it correct, don't just say P(Z ≤ 0) <> .7, instead say P(Z ≤ 0) = .5 or P(Z ≤ .525) = .7.

A. If X and Y are random variables, $V(XY) = X^2V(Y)$.

False. If X were a constant, this would be true, but X is a variable, so the above statement makes no sense. Whatever V(XY) equals, it equals a constant, while $X^2V(Y)$ is a random variable and not a constant.

B. If N is large and X has a uniform distribution, then $\overline{X}$ will also have a uniform distribution.

False. If N is large, $\overline{X}$ will have a normal distribution, regardless of what the distribution of X is.

C. $P(Z \le -1.7) = .95543457$

False. True statements would be $P(Z \le 1.7) = .95543457$ or $P(Z \le -1.7) = 1 - F(1.7) = 1 - .95543457 = .0446$.

D. A fair die has been rolled two times, each time producing a 6. If the die is rolled again, the probability is 1/216 that you will once again get a 6.

False. Before you roll any die, the probability is 1/216 that you will roll 3 6s in a row. But, after you roll 2 6s, the probability is 1/6 that you will roll another one. Die tosses are independent and the first two rolls have no effect on how the third roll comes out.

E. Four hopelessly confused graduate students can't decide whether the answer to problem 1D is True or False. If they all flip a fair coin to decide, there is a 1/8 chance that they will all get the same answer.

True. Note that they will all get the same answer if they all roll heads or all roll tails. The probability of 4 heads is $.5^4 = 1/16$, the probability of 4 tails is also 1/16, so the probability of all heads or all tails is 1/8.

2. (10 points each, 30 points total) Answer <u>three</u> of the following. The answers to most of these are fairly straightforward, so do not spend a great deal of time on any one problem. NOTE: I will give up to 5 points extra credit for each additional problem you do correctly.

A. $\overline{X}$ = 40, N = 9. Determine the 95% confidence interval when

a. $\hat{\sigma}$ = 12

For s = 12, look at Appx E, Table 3, v = 8, 2Q = .05:

$$\bar{x} \pm (t_{\alpha/2,v} * s/\sqrt{N}), i.e.$$
$$40 - (2.306 * 12/\sqrt{9}) \le \mu \le 40 + (2.306 * 12/\sqrt{9}), i.e.$$
$$30.776 \le \mu \le 49.224$$

b. σ = 9

For σ = 9:

$$\bar{x} \pm (z_{\alpha/2} * \sigma/\sqrt{N}), \text{ i.e.,}$$
$$40 - (1.96 * 9/\sqrt{9}) \le \mu \le 40 + (1.96 * 9/\sqrt{9}), \text{ i.e.,}$$
$$34.12 \le \mu \le 45.88$$

B. Here are the results from a previous cohort's first exam in statistics. Compute the mean and variance of the scores. There were 9 Sociology Graduate Students in the class. (The exam was obviously much too easy – hopefully this exam will be more of a challenge for you.)

Score
84
100
109
110
106
97
102
82
105

Expand the table as follows:

	X	*X²*	*(X – E[X])²*
	84	7056	238.53
	100	10000	0.31
	109	11881	91.31
	110	12100	111.42
	106	11236	42.98
	97	9409	5.98
	102	10404	6.53
	82	6724	304.31
	105	11025	30.86
Sum	895	89835	832.22
Mean	99.4444	9981.6667	92.47

So, $\mu = 895/9 = 99.444$. $\sigma^2 = E(X^2) - E(X)^2 = 9981.6667 - 99.4444^2 = 92.47$. $\sigma = 9.62$. Or, $\sigma^2 = E((X - E[X])^2) = 92.47$.

C. A company has developed a questionnaire that measures attitudes toward its product on a scale that ranges from a low of 0 to a high of 200. It now wants to administer this questionnaire to a random sample of consumers to find out how popular its product is. On the one hand, it wants results that are fairly precise, but on the other hand it only has a limited budget for collecting data. If the population standard deviation σ = 45, how many people will have to be surveyed so that the true standard error is no greater than 3? How many people would have to be surveyed so that the true standard error would be no greater than 1? Comment briefly on how a desire for increased precision would affect the cost of this survey.

The true SE = $\sigma / \sqrt{N}$. So, if σ = 45 and we want the true SE to be 3,

$$\sigma / \sqrt{N} = 45 / \sqrt{N} = 3$$
$$\Rightarrow 15 = \sqrt{N}$$
$$\Rightarrow N = 225$$

Similarly, for the true SE to be 1,

$$\sigma / \sqrt{N} = 45 / \sqrt{N} = 1$$
$$\Rightarrow 45 = \sqrt{N}$$
$$\Rightarrow N = 2025$$

So, we would have to collect a sample 9 times as large to get the higher level of precision. The company will have to decide whether the added costs are worth it, or whether it can get by with a lesser level of precision.

D. It is January 2004. After a disappointing 7-7 record in 2002, the Nebraska Cornhuskers have shocked the college football world by going 13-0 in 2003. Today, however, in the National Championship game, they face their greatest challenge of all: the undefeated and #1 ranked Miami Hurricanes. Miami humiliated the Cornhuskers in the title game just two years ago, and Coach of the Year Frank Solich from Nebraska has vowed that it will not happen again.

Nebraska's greatest strength all season long has been its impenetrable defense. The Huskers estimate that they have a 60% chance of holding Miami to 14 points or less, and if they do, there is a 70% chance Nebraska will win. But, if the defense can't contain Miami, then Nebraska's offense, which has struggled all season, will have to come through. The coaches estimate that there is a 15% chance that Miami will score more than 14 points and Nebraska will win the game.

What is the probability that Nebraska will win the game if Miami scores more than 14 points? What is the probability that Nebraska will win the game?

You are told:

P(Miami scoring 14 or less) = .60, implying P(Miami scoring more than 14) = .40.
P(Nebraska Winning **|** Miami scores 14 or less) = .70
P(Nebraska Winning ∩ Miami scores more than 14) = .15.

You are asked to find P(Nebraska winning **|** Miami scores more than 14) and P(Nebraska Winning). Along the way, you'll also need to find P(Nebraska Winning ∩ Miami scores 14 or less)

So,

P(Nebraska Winning ∩ Miami scores 14 or less) =
P(Nebraska Winning **|** Miami scores 14 or less) * P(Miami scoring 14 or less) =
.70 * .60 = .42

P(Nebraska winning **|** Miami scores more than 14) =
P(Nebraska Winning ∩ Miami scores more than 14)/ P(Miami scores more than 14) =
.15/.40 = $\boxed{.375}$
P(Nebraska Winning) =
P(Nebraska Winning ∩ Miami scores 14 or less) +
P(Nebraska Winning ∩ Miami scores more than 14) =
.42 + .15 = $\boxed{.57}$.

E. You have to take a true-false test on a subject you know absolutely nothing about. You must get 2/3 or more of the answers right in order to pass. Would you rather take a 3 question test, where you had to get at least 2 answers right, or a 36 question test, where you had to get at least 24 answers right? Or would it not make any difference to you how long the test was? Explain your reasoning.

In order to pass this exam, you will have to get lucky. You expect to get half the questions right (i.e. p = .5) and you need to get 2/3 right (i.e. $\hat{p}$ = .67 or better). You therefore want to choose whatever test length will maximize your chances of getting lucky.

As we saw in the class notes, the bigger N is, the closer $\hat{p}$ will tend to be to p, because the sampling variance gets smaller as sample size increases. So, the more questions you have, the less likely you are to get an exceptionally large number of questions right. (You are also less likely to get a really low score, but if you flunk, you flunk; it doesn't matter whether you flunk by a little or a lot.) Ergo, you want the shorter test.

To think of it another way, suppose you had to get *all* the questions rights. You would certainly prefer a shorter exam then. Well, 2/3 isn't perfect, but it is still better than you are naturally capable off, so take as few questions as possible and hope you get lucky.

If you want to compute the exact probabilities, you can easily determine that, when p = .5 and N = 3, P(2 right) = .375, P(3 right) = .125, so P(at least 2 right) = .50. So, even though you know nothing about the topic, you have a 50% chance of passing a 3 question true-false test just by guessing.

For N = 36 and p = .5, we want to find the probability of 24 or more successes, i.e. P(X ≥ 24). Perfectionistic (masochistic?) individuals can figure out the exact probability using the binomial distribution (i.e. for N = 36 and p = .50, compute P(24), P(25), P(26), …, P(36), and then add up the 13 probabilities), but everyone else will probably prefer to use the Normal Approximation to the Binomial. We compute

$$z = \frac{\#\text{ of successes} - .5 - Np}{\sqrt{Npq}} = \frac{24 - .5 - 18}{\sqrt{36 * .5 * .5}} = \frac{5.5}{3} = 1.83$$

So, finding P(X ≥ 24) is equivalent to finding P(Z ≥ 1.83). The probability of getting a Z score of 1.83 or higher is 1 – F(1.83) = 1 - .96637509 = .033, i.e, you only have about a 3% chance of passing a 36 question test if you are guessing and have to get 2/3 or more right.

So, if you have a choice between 3 questions or some multiple of 3, the fewer questions, the better. In fact, if you are given the option and want to get the pain over with as quickly as possible, you might as well stake everything on getting a single question right – you'll have a 50% chance of success.

I suspect many people's intuition leads them astray on this because (a) we think we know something about the topic, even though the question explicitly states that we don't, and (b) we are risk-adverse – we think it is better to drag things out instead of taking a chance that will settle things immediately. In this case though, you are better off just trying to get things resolved as quickly as possible.

There are some real-life applications to this. In games where both chance and skill are involved, if you are up against a stronger player you should try to maximize the role of luck. For example, if you were playing backgammon against the world champion, you'd prefer a 5 point match over a 25 point match. Conversely, if you are the stronger player, you'd like a longer match – you might get unlucky in a few games, but in a long match luck will tend to even out and skill will become more important.

Incidentally, if you were determined to get it exactly right, or if I was nice and let you use a computer on exams, you could enter the following formula into a cell in an Excel spreadsheet:

=1-BINOMDIST(23,36,0.5,1)

This computes 1-P(X ≤ 23) = P(X ≥ 24) for N = 36 and p = .5. The result Excel gives is 0.032623.

To have Excel do the Normal approximation for you, use the following formula:

=1-NORMDIST(23.5,18,3,1)

This computes 1 – F(23.5) = P(X > 23.5) where X has a mean of 18 and a standard deviation of 3. The result Excel gives is 0.033376. Or, to use the standardized normal distribution, you could also do something like *=1-NORMSDIST(1.83)* or to be really precise *=1-NORMSDIST(5.5/3).* Of course, you don't really need to use the approximation in this case since the computer can easily get the exact value for you.

Excel has help that explains these functions and others if you want to use them. Other programs also often contain something similar, e.g. in Stata 7, the command

bitesti 36 24 .5

will give you the probability of 24 or more successes when N = 36 and p = .5. The output is

```
        N   Observed k   Expected k   Assumed p   Observed p
------------------------------------------------------------
       36         24           18       0.50000      0.66667

  Pr(k >= 24)           = 0.032623  (one-sided test)
  Pr(k <= 24)           = 0.985592  (one-sided test)
  Pr(k <= 12 or k >= 24) = 0.065245  (two-sided test)
```

---

3. (25 points) A recent article on the WebMD web pages (http://content.health.msn.com/content/article/73/88984.htm) states the following:

> Married men are healthier men. But for women, the health benefit of marriage depends on the health of the marriage. Over and over again, studies show that marriage is good for men's health. For women, the picture has been less clear. Some studies suggest that women need marriage like a fish needs a bicycle. That's true, a new study finds -- but only for women who aren't highly satisfied. Women who say their marriages are very satisfying have better heart health, healthier lifestyles, and fewer emotional problems, report Linda C. Gallo, PhD, and colleagues.

Another study decides to investigate the relationship between gender, marital satisfaction, and health. A sample of 1000 men and 1000 women, each of whom has been married for at least 30 years, is drawn. 70% of the men but only 40% of the women report that they are happy in their marriage. Five hundred (500) men and 300 women report that they have happy marriages and are in good health. For those who do not have happy marriages, 30% of the men and 35% of the women report they are in good health.

a. (10 pts) Complete the following table. Remember, there were 1000 men and 1000 women in this study.

	Male			Female		
Health/Marital satisfaction	Happy Marriage	Unhappy marriage	Σ	Happy Marriage	Unhappy Marriagr	Σ
Good health						
Poor health						
Σ			1000			1000

It is fairly easy to fill in the numbers:

	Male			Female		
Health/Marital satisfaction	Happy Marriage	Unhappy marriage	Σ	Happy Marriage	Unhappy Marriage	Σ
Good health	500	90	590	300	210	510
Poor health	200	210	410	100	390	490
Σ	700	300	1000	400	600	1000

b. (5 pts) What percentage of those in unhappy marriages have good health? What percentage of those in happy marriages have good health?

There are 900 unhappy marriages (300 men, 600 women). 300 of those (90 men, 210 women) have good health, i.e. 33.33%. There are 1100 happy marriages (700 men, 400 women). 800 of those (500 men, 300 women) have good health, i.e. 72.7%. So, you are more than twice as lucky to have good health if you are in a happy marriage as compared to an unhappy marriage,

c. (10 pts) As these figures show, women tend to have worse health than men do. However, women are also less likely to be happy in their marriages. Suppose that just as many women had happy marriages as men did. Suppose further that women maintained their marital satisfaction-specific health rates. What percentage of women would then be in good health? Based on these results, do you think that differences in marital satisfaction explain much of the health differences between married men and women, or does it explain relatively little?

For men, P(Happy) = .70, P(Unhappy) = .30.
For women, P(Good health | Happy) = 300/400 = .75, P(Good Health | Unhappy) = .35.

So, if women were as happy in their marriages as men were while maintaining their satisfaction-specific health rates,

$$P(Good\ Health) =$$
$$P(Good\ Health \mid Happy)^W * P(Happy)^M + P(Good\ Health \mid Unhappy)^W * P(Unhappy)^M =$$
$$(.75 * .70) + (.35 * .30) = .63$$

Ergo, in this hypothetical situation, more women would be in good health than men are (i.e. 63% of women would be healthy compared to 59% of the men). This suggests that the main reason that married women are less healthy than married men is because they are less likely to be happy with their marriages than men are. A happy marriage is actually slightly more beneficial for women than it is for men, but unfortunately women are less likely to be happy in their marriages. Of course, an alternative explanation might be that, because women are less healthy, they are also less likely to be happy in their marriages.

4. (25 points) Although he announced his candidacy only days ago, General Wesley Clark is already being touted by many as the frontrunner for the Democratic Presidential nomination. There is universal agreement that on the day he announced, 40% of all Democrats wanted Clark to be the party's nominee. But, his rivals for the Democratic nomination claim that Clark's support has already declined. A random sample of 175 Democratic voters is drawn, 58 of whom state that Clark is their first choice to be the Party's nominee for President. Using the .05 level of significance, test whether Clark's support has declined. Be sure to indicate:

(a) The null and alternative hypotheses - and whether a one-tailed or two-tailed test is called for.

H_0: p = .40 (or E(X) = 70)

H_A: p < .40 (or E(X) < 70)

The alternative is one-tailed since Clark's rivals claim his support has declined. Obviously, they won't have much of a case if the sample results suggest that he has gained support instead.

(b) The appropriate test statistic

The appropriate test statistic is

$$z = \frac{\#\ of\ Clark\ Supporters \pm CC \text{ - } Np_0}{\sqrt{Np_0 q_0}} = \frac{x \pm CC \text{ - } (175*.4)}{\sqrt{175*.4*.6}} = \frac{x \pm CC \text{ - } 70}{\sqrt{42}}$$

For the correction for continuity, we will subtract .5 if there are more than 70 supporters, and we will add .5 if there are less than 70 supporters.

(c) The critical region

For the critical region, we will reject H_0 if $Z_c < -1.645$

(d) The computed value of the test statistic

$$z = \frac{\#\ of\ Clark\ Supporters \pm CC \text{ - } Np_0}{\sqrt{Np_0 q_0}} = \frac{x \pm CC \text{ - } (175*.4)}{\sqrt{175*.4*.6}} = \frac{x \pm CC \text{ - } 70}{\sqrt{42}} = \frac{58 + .5 - 70}{\sqrt{42}} = -1.77$$

(e) Your decision - should the null hypothesis be rejected or not be rejected? Why?

Reject. The test statistic falls in the critical region. It appears that Clark's support has indeed declined.

(f) Would your decision change if you used the .01 level of significance instead? Why or why not?

Do not reject. For the .01 level of significance, you reject if $Z_c < -2.33$. The computed test statistic of -1.77 is greater than that. Ergo, if we use a more stringent standard, we would not conclude that Clark's support has declined. Even so, if I was Clark, I might be a little nervous about these results and hope that any downturn is only temporary and not part of a trend.

Here are some ways to work this problem with a computer:

EXCEL. To have Excel compute the probability of 58 or fewer successes, enter

=BINOMDIST(58,175,0.4,1)

Excel gives you 0.0369, i.e. there is less than a 4% chance of getting this few successes by chance if p really does equal .5. Or, to use the Normal approximation (which really isn't necessary with a computer given that you can easily get the exact probability), enter

=NORMDIST(58.5,70,SQRT(42),1)

Excel gives 0.037991.

To have Excel give you the minimum number of successes you would need to *not* reject the null at the .05 level of significance, enter

=CRITBINOM(175,0.4,0.05)

Excel says you need at least 59 successes. For the .01 level of significance, enter

=CRITBINOM(175,0.4,0.01)

Excel says you need at least 55 successes to not reject the null.

STATA 7. To get Stata 7 to do it for you, enter

bitesti 175 58 .4

Stata gives you

```
        N   Observed k   Expected k   Assumed p   Observed p
------------------------------------------------------------
      175         58           70       0.40000      0.33143

  Pr(k >= 58)            = 0.974167  (one-sided test)
  Pr(k <= 58)            = 0.036903  (one-sided test)
  Pr(k <= 58 or k >= 83) = 0.064486  (two-sided test)
```

SPSS. As far as I know, SPSS requires you to get the raw data entered in order to work this problem. Luckily, in this case, it is easy to do. To get SPSS to work this problem for you, enter the syntax

```
Data List Free / Supports Freq.
Begin Data.
  0  117
  1   58
end data.
Weight by Freq.
Value Labels Supports 0 "Does not support Clark" 1 "Supports Clark".

NPAR TEST
  /BINOMIAL (.40)= supports(1,0)
  /MISSING ANALYSIS.
```

The first few commands tell SPSS that there are 117 people who don't support Clark and 58 who do (using the weight command saves you the trouble of entering all 175 cases separately). The NPAR tests command provides a 1-tailed test of whether p = .40. SPSS gives you

NPar Tests

Binomial Test

		Category	N	Observed Prop.	Test Prop.	Asymp. Sig. (1-tailed)
SUPPORTS	Group 1	1.00 Supports Clark	58	.3	.4	.037[a,b]
	Group 2	.00 Does not support Clark	117	.7		
	Total		175	1.0		

a. Alternative hypothesis states that the proportion of cases in the first group < .4.

b. Based on Z Approximation.