

Problem Set 1

Note: this problem set is primarily intended to get you used to manipulating and presenting data using a spreadsheet program. While subsequent problem sets will be useful indicators of the difficulty of exam questions, this one is not.


Descriptive Statistics

1. Go to the Bureau of Justice Statistics. Find the data for reported crime in Pennsylvania for 1960-2009. Put these data in a spreadsheet.


These data are available in convenient formats at

<http://www.ucrdatatool.gov/Search/Crime/State/StateCrime.cfm>

2. Graph the total violent crime by year.



3. Graph the total violent crime rate by year.



- If you were trying to give someone a sense of the evolution of crime over time, is the graph from #2 or #3 more useful? Why?

It turns out that population growth in PA is more or less linear, so the total and rate series give, more or less, the same perspective (especially since inter-Census years are linearly interpolated for population). In terms of which is better, it depends on the questions of interest. If one is interested in getting some sense of how likely an individual is to be a victim of violent crime, rates are more useful. If, instead, someone was interested in budgetary issues, totals may be more useful.

- Calculate the standard deviation of violent crime rates in PA over this period.

In excel, the standard deviation macro is “=STDEV.P(cell range)” [for the purposes of this class, ignore the difference between a standard deviation for a population and for a sample].

111.11

- Calculate the standard deviation of homicide rates in PA over this period.


1.18

- From the Bureau of Justice Statistics, pull crime totals and rates by all available categories of crime for the year 2009 into a spreadsheet.

<http://www.ucrdatatool.gov/Search/Crime/State/StateCrime.cfm>

- Create a histogram for violent crime rates.


You need to load the “Analysis Tool Pack.” To see how to do that, search in the help.



- Get state per capita GDP data from the Bureau of Economic Analysis for 2009. Create a scatterplot of violent crime rates on the vertical axis and per capita GDP on the horizontal axis. What does the relationship appear to be? Does this relationship change if DC is eliminated from the graph?


Data available at: <http://www.bea.gov/regional/gsp/>

Scatter with DC



With DC Excluded

Violent Crime rate



Problem Set 2

Market timing, while not strictly illegal, is a practice mutual funds wish to discourage. What market timing in this context means is the process of moving money within a mutual fund complex from funds containing US securities to funds containing, for example, British securities in search of what are essentially risk free returns. The existence of these returns arises from the fact that world markets tend to be positively correlated. When the British market closes for the day at 4:30 pm, US markets still have five hours of trading. If the US market turns up after 11:30 am, such an upward movement will not yet be capitalized into British securities, but due to the positive correlation, it likely will be once the London markets open again. Because of mutual fund pricing conventions (and, to some extent regulations), if I buy shares in the mutual fund with London securities, I get the price that the fund closed at the previous day even though the fund is likely to be priced higher at the end of the day my purchase is credited due to the uptick in US markets. Because this arbitrage essentially moves appreciation from existing owners of the fund to the market timer, mutual fund companies have an incentive to police it. What some companies do is disallow trades that look suspicious. However, investors often value liquidity and ease of movement between funds, viewing it as one of the primary values that mutual funds provide.

While the prospectuses of all of the funds FSU funds offers contain language suggesting FSU will police market timing and disallow some trades, it also uses language suggesting that investors will enjoy easy movement across funds. FSU freezes the account of one of its customers due to a suspicion that his trades exhibit the pattern of a market timer. The customer sues FSU claiming that he has a contractual right to move his money around easily and FSU only has a right to restrict his trades if it has clear evidence that he was engaged in market timing. He seeks an injunction against FSU freezing his funds and restricting his trading, as well as damages for lost gains he would have made if the trades had been permitted.

FSU has no direct evidence of the customer engaging in market timing. Its analysts, however, are suspicious of his trading patterns. Assume that the pricing of mutual funds (due to the efficient markets hypothesis) follows a random walk, meaning (for our purposes) that on any given day, the price of the fund is as likely to go up as it is to go down.

The customer in question engaged in 20 trades where he took money out of funds investing in US equities and put it in funds investing in British equities. On 15 of those trades, the British equities funds in question appreciated in value. The customer claims he was simply engaging in a rebalancing of his portfolio and was not engaged in any market timing. In support of this, he notes that on 5 of the trades, the British funds exhibited negative returns on the day in question.

You are asked to provide a statistical analysis of the likelihood that the individual was indeed engaged in trading that was random (at least with respect to market timing concerns).

This problem can be solved using a binomial distribution given the random walk assumption. That is, on any given day, a fund will appreciate with probability 0.5 and will depreciate with probability 0.5. $n = 20$ trades and the investor had 15 “successes.” Determine the probability of this happening by chance.

Recall

$$P(x) = \frac{n!}{(n-x)!x!} p^x (1-p)^{n-x}$$

So the likelihood of picking 15 correct days out of 20 is

$$\frac{20!}{5!15!} 0.5^{15} 0.5^5 = 15504 * .000001 = 0.015$$

But we presumably also need to factor in the possibility that he could have done better as well (i.e., the relevant question is probably what is the likelihood that he would have chosen at least 15 correct days to trade), so we also want to add in P(16), P(17), P(18), P(19), and P(20).

Each of those can be calculated the same way as above (replacing x with 16, 17, etc).

$$P(16) = .005$$

$$P(17) = .001$$

$$P(18) = .0002$$

$$P(19) = .00002$$

$$P(20) = .000001$$

So the total likelihood of picking at least 15 correct days to trade out of 20 is 0.021, so he had only slightly more than a 2 percent chance of this happening randomly.

Problem Set 3

Assume the same set up as in problem set 2. However, in this case, assume that the number of suspicious trades is 400. 400 trials are enough to use the normal distribution as an approximation for the binomial distribution.

Determine the likelihood of having at least x number of trades occurring on days where the British funds appreciate in value where:

$$x = 210$$

$$x = 220$$

$$x = 230$$

$$x = 240$$

The easiest way to handle this problem, because of the fact that 400 trials allows us to use the normal distribution as an approximation to the binomial distribution, is to set this up as a simple hypothesis test, where the hypothesis that the customer is trading at random (at least with respect to market timing issues).

We would expect, under random trading, that the customer would choose correctly 200 times out of 400. However, there will be variation around this.

We need to convert each x to standard deviation terms, so we need to calculate the standard deviation. Recall that in the binomial distribution, variance is simply $n \cdot p \cdot (1-p) \Rightarrow 400 \cdot .5 \cdot .5 = 100$, so the standard deviation is 10. We thus need to determine how likely one is to observe more than a 1 standard deviation departure from the mean, 2 standard deviations, 3 standard deviations, and 4 standard deviations.

If you consult the table of critical values for the standard normal distribution (back cover of text, or fold out chart, or a number of other sources), you will find that the probability mass to the left of z standard deviations is:

$$z = 1: 0.8413$$

$$z = 2: 0.9772$$

$$z = 3: 0.9987$$

$$z = 4: 0.999+$$

so the likelihood of observing at least 210 correctly timed trades at random is $1 - 0.8413 < 16\%$, 220 is $< 3\%$, 230 is $< 0.2\%$, and 240 is $< 0.1\%$.

Problem Set 4

You are asked to perform a preliminary analysis of some evidence in a discrimination suit. The claim asserts that FSU Corp. discriminates against black employees. While there is no direct evidence of discrimination, you are told that a random sample of 30 black FSU employees has an average wage of \$30,000, whereas the average wage for non-black employees is \$37,000. You ask human resources to calculate the variance of the salaries of non-black FSU employees. You are told that the variance in non-black wages is 508,650,519.

1. Is this evidence generally consistent with anti-black discrimination?

We need to set this up as a hypothesis test. The hypothesized mean is \$37,000. The test statistic is 30,000-37,000 normalized in standard deviation terms. So we need to calculate:

$$\left(\frac{30,000 - 37,000}{\frac{\sqrt{\text{var}}}{\sqrt{n}}} \right) = \frac{-7,000}{\frac{\sqrt{508,650,519}}{\sqrt{30}}} = \frac{-7,000}{\frac{22,553}{5.5}} = \frac{-7,000}{4,101} = -1.71$$

Because we are dealing with a relatively small sample (n = 30), we need to consult the t distribution which tells us that for 29 degrees of freedom (n - 1), the critical values associated with the 95% confidence interval are -2.045 and 2.045. Since our test statistic is inside of this range, we cannot reject the hypothesis that the black employees are paid comparably to the white employees.

2. At what type 1 error level is this evidence consistent with anti-black discrimination?

Consulting the t distribution table (again using the 29 degrees of freedom row), we need to find 1.71 as a critical value. The closest we can find in the table (though of course we could get more precision consulting electronic sources) is 1.699 which is associated with a p value of < 0.10 (two tailed test). This tells us that the black mean wage is consistent with the non-black mean wage at any type 1 error of less than 10 percent or so.

3. If you are the defendant, what kinds of arguments do you make on your behalf in light of this evidence?

If you are the defendant, you use the analysis above to suggest that it is not possible to distinguish the black pay from non-black pay; therefore, there can be no inference of discrimination.

4. If you are the plaintiff, what kinds of arguments do you make on your behalf in light of this evidence?

If you are the plaintiff, you should argue that we should only be using a one tailed test since there would not be any inference of discrimination from a legal standpoint if the black mean wage were above the non-black mean wage. Thus, none of the type 1 error should be apportioned to the right hand tail. Using a one tailed test, the critical value associated with a 5 % type 1 error is -1.699. Our

test statistic is less than this critical value, so the black mean wage is not consistent with the non-black mean wage.

Both sides may then argue that different attributes should be controlled for (in a regression framework) or that the comparison mean wage should be calculated on different bases (just whites vs non-blacks, etc) depending on which calculations are more favorable to their positions.

Problem Set 5


Use the excel file called PS5.xlsx for this problem set.

The wage variable is a random variable drawn from a normal distribution; treat it as the relevant population. Sample1 . . .sample6 are 6 independently drawn random samples from wage.


1. Create histograms for wage and each of the samples. How are they similar? How are they different?

Histograms are found in the data analysis add-in pack (as described in an earlier problem set).


Wage




Sample1




Sample 2




Sample 3




Sample 4



Sample 5



Sample 6



Similarity: They are all centered on something close to the population mean (although not exactly)
 Differences: The samples are generally more narrow in their ranges than the population (min for each sample > min for pop; max for each sample < max for pop). Generally, you can't tell with great confidence that the samples are actually coming from the distribution from a visual standpoint.

2. Calculate the mean, variance, and standard deviation (using excel commands) for wage.

=average(a2..a10001) which yields 149,998
 =var.p(a2..a10001) which yields 621,549,795
 =stdev.p(a2..a10001) which yields 24,931

3. Calculate the means for sample 1 through sample 6. For each one, test the hypothesis that the sample mean is equivalent to the population mean at the 20% type 1 error level

For each of these we need a test statistic and we need to compare that test statistic to the appropriate critical value from the t distribution (since n = 100 for the samples). The critical values for the t distribution with 99 degrees of freedom (n-1) at the 20% type 1 error are -1.290 and 1.290 (note I used degrees of freedom for 100 since that's what the book prints and the difference between dof 99 and dof 100 is small, but you could be more precise using electronic sources).

Now we need t statistics. The general form will be:

$$t = \frac{\bar{x} - 149,998}{sd(x) / \sqrt{100}}$$

Sample 1: average = 147,316; sd = 23,458; t = -1.14, so we can't reject

Sample 2: average = 144,608; sd = 23,676; t = -2.28, so we reject

Sample 3: average = 148,011; sd = 23,952; t = -0.83, so we can't reject

Sample 4: average = 155,729; sd = 25,290; t = 2.27, so we reject

Sample 5: average = 149,192; sd = 24,081; t = -0.33, so we can't reject

Sample 6: average = 150,281; sd = 25,720; t = 0.11, so we can't reject

So, as you can see, while we don't reject in 4/6 times, we do incorrectly reject one third of the time (more or less "close" to our chosen 20% type 1 error).

4. Wagesub is a random variable; test whether it comes from the same distribution as the wage variable at the 1% type one error level.

Using the same formula for the test statistic, we get:

$$t = \frac{96,026 - 149,998}{25,642 / \sqrt{100}} = -21.05$$

Which is well beyond the critical value of -2.63, so we can safely reject the hypothesis that wagesub comes from the same distribution.

5. Subsample1 through subsample6 are random draws from wagesub. Test whether each comes from the same distribution as wage providing the appropriate p value.

Subsample	1	2	3	4	5	6
Average	92349	97494	99687	98846	101656	98168
SD	22565	21095	28061	24198	24162	28109
t	-14.00	-13.64	-9.82	-11.58	-10.96	-10.10

All are well below the critical value of -2.77 which is the 29 dof critical value for the 1% type 1 error, so we know that $p < 0.01$ for each of these.

Problem Set 6

Use the dataset PS6.xlsx for this problem set. It contains state level data for 2009, including state population, violent crime total, and property crime total (from the Bureau of Justice Statistics). It also includes real per capita state GDP (rGDPpc; from the Bureau of Economic Analysis), the unemployment rate (from the Bureau of Labor Statistics), and an indicator for the fraction of people indicating they are heavy drinkers, have less than a high school education, and have a college degree or higher (from the Behavioral Risk Factor Surveillance Survey done by the Centers for Disease Control).

1. When analyzing the determinants of crime, does it make more sense to examine the levels or the rate of crime and why?

It depends on the question, in some sense. For example, if one were attempting to estimate a criminal enforcement budget for an entire state, looking at total crime in the state could be interesting. Generally speaking, however, for most policy purposes, we probably care about crime rates. For example, almost regardless of anything else, California will have more total crime than anywhere else solely because of its size, but that doesn't tell us much about the causes or predictors of crime in a policy relevant sense.

2. Regress the violent crime rate (crime per 100,000 people is the general convention) on the heavy drinker variable and a constant. Interpret your coefficient, including its statistical significance. Discuss what the p value represents.

Excel output:

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.080678905
R Square	0.006509086
Adjusted R Square	-0.013766239
Standard Error	208.4441406
Observations	51

ANOVA

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>
Regression	1	13948.62989	13948.62989	0.321034841
Residual	49	2128999.028	43448.95976	
Total	50	2142947.658		

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>
Intercept	331.3154506	126.2769593	2.62372053	0.011562928
Heavy Drinking	13.59171484	23.98822812	0.566599366	0.573572697

This suggests that violent crimes per 100,000 residents increases by about 14 additional crimes for every 1% increase in the heavy drinking rate. However, the standard errors are very large, leaving us with a t stat of 0.6 which is not statistically significant (i.e., it cannot be distinguished from random noise around a zero effect). We conclude from this that we are not very confident that there is a non-zero effect of heavy drinking incidence on the rate of violent crime for conventional choices of a type 1 error. The p value, in fact, tells us that we would not reject the null of zero effect unless we had a two tailed type 1 error of greater than 57 percent.

- Intuitively, discuss what is likely to happen to the coefficient for the heavy drinking variable if you added the unemployment rate to the regression in #2. Run the regression with the unemployment rate and discuss the interpretation of each coefficient, including a discussion of how the heavy drinking coefficient changes.

Presumably, both heavy drinking and unemployment are positively associated with violent crime. Also, presumably, drinking goes up during times of unemployment, so we might expect the magnitude of the heavy drinking effect to decline when we include unemployment since the heavy drinking effect in #2 was likely picking up some of the unemployment effect as well.

Excel output:

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.339509176
R Square	0.11526648
Adjusted R Square	0.078402584
Standard Error	198.7427999
Observations	51

ANOVA				
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>
Regression	2	247010.034	123505	3.126812
Residual	48	1895937.624	39498.7	
Total	50	2142947.658		

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>
Intercept	53.4032407	166.0896025	0.321533	0.749203
unemployment	34.69005386	14.28108482	2.429091	0.01893
Heavy Drinking	10.46748218	22.90791052	0.456937	0.649777

Our intuition was borne out wrt the heavy drinking since the magnitude of the effect declines once we control for unemployment. This suggests that the number of violent crimes per 100,000 population increases by almost 35 crimes when unemployment increases by 1%. This effect is statistically significant. We know this because the t stat is 2.4 which is larger than the critical value associated with a 5% type 1 error. In fact, the p value suggests that we would reject the null of no effect for any

type 1 error greater than 1.9% (two tailed). The heavy drinking interpretation is the same as above (with an effect of 10.5 crimes) and it is still not statistically significant.

4. Add the per capita income variable to the regression from #3. How do the coefficients change.

Excel Output

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.735405131
R Square	0.540820707
Adjusted R Square	0.511511391
Standard Error	144.6931596
Observations	51

ANOVA				
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>
Regression	3	1158950.468	386316.8	18.45217734
Residual	47	983997.1903	20936.11	
Total	50	2142947.658		

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>
Intercept	-96.59380904	123.0375412	-0.78508	0.436348693
rGDPpc	0.009036	0.001369118	6.599868	3.31051E-08
unemployment	40.82137304	10.4386549	3.910597	0.000294822
Heavy Drinking	-46.59052246	18.78549727	-2.48013	0.016772598

The interpretations of unemployment and heavy drinking are the same as above except now the effect of a 1% increase in the unemployment rate is to increase the violent crime rate by 41 crimes per 100,000 population (statistically significant at any type 1 error above 0.03%) and the heavy drinking effect flips signs such that now we would estimate that a 1 % increase in the heavy drinking rate would be associated with a reduction of almost 47 violent crimes per 100,000 population (and the effect is statistically significant at any type 1 error above 2%). The income effect implies that when real per capita GDP increase by \$1, crime goes up by 0.01 violent crimes per 100,000 population. While this seems silly, the effect is statistically significant at any type 1 error bigger than 0.000003%. Of course, there's some omitted variable bias at work here.

5. Run the kitchen sink regression (income, unemployment, drinking, both education variables). Discuss reasons why this model is superior to the model estimated in #4.

Excel Output:

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.756958003
R Square	0.572985419
Adjusted R Square	0.525539354
Standard Error	142.6004438
Observations	51

ANOVA				
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>
Regression	5	1227877.762	245575.6	12.07656
Residual	45	915069.8963	20334.89	
Total	50	2142947.658		

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>
Intercept	-241.4941655	182.7995166	-1.32109	0.193151
rGDPpc	0.009336265	0.001730708	5.394476	2.45E-06
unemployment	34.34016287	10.93607975	3.14008	0.002982
Heavy Drinking	-29.33366837	21.56903443	-1.35999	0.180608
BelowHS	13.62103714	8.005380954	1.701485	0.095753
CollegePlus	-0.63879844	4.98407484	-0.12817	0.898587

Presumably, education has some effect on crime and education levels are probably correlated with income, unemployment, and heavy drinking, leading to omitted variables problems. For example, income probably goes up with education and crime probably declines with both of these, so not controlling for education probably understates the negative effect of income on crime. However, when we compare the estimates, this does not seem to be borne out, but a similar intuition for unemployment does seem to be borne out. Until we face a degree of freedom problem (and with n=51, we probably don't face one until we have > 20 controls in the model), it's probably better to control for stuff. Unfortunately, we still have lots of stuff missing here, so we do not have any real confidence in any of these estimates.

Problem Set 7

For background on this analysis, see Klick and Sitkoff (2008)
<http://www.law.upenn.edu/fac/jklick/108CLR749.pdf>

1. Find daily return data for the Hershey company (ticker HSY) online (possible sources include Google Finance or Yahoo Finance).

You can find the HSY price data at <http://www.google.com/finance?q=hsy> and click on historical prices. From there, you can set a date range and you can download the data directly to an excel spreadsheet. 100 trading days before July 25, 2002 is March 4, 2002.

There are a couple of versions of the S&P 500 available on Google Finance (ETFs based on S&P 500; mutual funds based on it, etc). Here is one: <http://www.google.com/finance?q=NYSE%3ASPY>

You need to use the closing prices to generate returns. To do this, get excel to calculate $r(t) = [P(t) - P(t-1)]/P(t-1)$

Note that you will have one fewer return than you have price data. This is fine, but if you want to based your model on 100 returns, you would add March 3, 2002 in your data. Any differences should be tiny.

2. Perform an event study to determine the effect of the Hershey Trust's decision to sell its controlling interest in the firm which was announced on July 25, 2002, using data from 100 trading days before that date to estimate your prediction model and using the S&P 500 for your market return.
 - a. First estimate it using the multiple step (finance) approach discussed in class.

Regress HSY return on SP 500 return (and a constant) for the days before July 25, 2002 (do not include the event date). You will get the following excel output:

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.420031741
R Square	0.176426663
Adjusted R Square	0.167936217
Standard Error	0.012079138
Observations	99

ANOVA

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>
Regression	1	0.003031835	0.003032	20.77943
Residual	97	0.01415284	0.000146	
Total	98	0.017184675		

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>
Intercept	3.30799E-05	0.001237341	0.026735	0.978726
X Variable 1	0.35990722	0.078953924	4.558446	1.5E-05

X Variable 1 is the market return. To get our predicted return, we take this implied equation

.00003 + 0.36*market return = predicted HSY return on July 25.

The S&P 500 return on July 25 was -0.009, so our predicted HSY is 0.00003+(0.36*-0.009) = -0.003

The actual HSY return was 0.253, so the abnormal return is 0.253 – 0.003 = 0.25 (or 25%).

We now need to standardize this, so we need some estimate of the standard deviation of HSY returns (before July 25) which we can calculate in excel. If we use stdev.s, we get 0.013. Our standardized abnormal return then is 0.25/0.013 = 19.23 standard deviations away from zero. This is statistically significant at the 5% level (in fact, $p < 0.001$, so it's significant at pretty much any type 1 error level you would choose).

b. The estimate it using the 1 step dummy variable regression approach.

You need to create an event dummy which takes the value of 0 for all of the non July 25 dates, and then regress HSY on the S&P 500 indicator and the event dummy you created (including the observation for July 25). The excel output is:

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.908426427
R Square	0.825238573
Adjusted R Square	0.821635244
Standard Error	0.012079138
Observations	100

ANOVA

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>
Regression	2	0.066831	0.033415	229.0212
Residual	97	0.014153	0.000146	
Total	99	0.080984		

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>
Intercept	3.30799E-05	0.001237	0.026735	0.978726

SPr	0.35990722	0.078954	4.558446	1.5E-05
event	0.255825622	0.012148	21.05965	5.85E-38

- c. Verify that you get comparable results.

You can see that the intercept and market coefficient are exactly the same. The abnormal return or event coefficient differ very slightly (due to some rounding I did in 2a; the computer keeps more digits than I do) and the standard error differs a tiny bit as well, again due to rounding. The conclusions are the same regardless of the approach used.

- d. Interpret the event day abnormal return, including its statistical significance.

The sale announcement was associated with a 25% return above and beyond any change due to market volatility. This movement is MUCH larger than the volatility observed due to random fluctuations in HSY prices. This effect is statistically significant at even a 0.0000000001 % type 1 error.

- e. Interpret your market return coefficient, including its statistical significance.

On average, when the s&p 500 goes up by 10%, HSY goes up by 3.6% and this effect is also statistically significant for any type 1 error that would generally be chosen.

- f. It turns out that the workers at Hershey were on strike from April 26, 2002 to June 8, 2002. Discuss what problems that might introduce into your analysis and how you might address these problems.

Presumably, the strike led to some changes in HSY's returns. This period is in your estimation period but since the strike is over by July 25, your counterfactual performance would not include any strike effect. That is, your prediction model is wrong. How big of an effect this is depends on how much the strike changed HSY's returns and how much this is "averaged" out over your estimation period. Since the strike occupies about 1/3 of your estimation period, this could severely influence the quality of your model. You could drop those observations (and, presumably, go back further in time for additional observations) or you could "control" for them by creating a strike dummy and including it in your regression model.

- g. For the event study to "work" in isolating the causal effect of the event, what assumption(s) must hold?

Your model must be correct in order to generate a useful counterfactual. That is, you want to make sure your prediction model is estimated during a time period that serves as a useful comparison for your event day and you include the relevant explanatory variables in your regression. Second, it must be the case that your "event" is the only thing that is different about day t . If other stuff also differs on that day (except for the things you controlled for in your model), you will not be able to parse out how much of the effect is caused by the event in question and how much is caused by the other stuff.

Problem Set 7 answer correction:

Note in 2a, the abnormal return should have been $0.253 - (-0.003) = 0.256$. Thus, the standardized abnormal return should be $0.256/0.013 = 19.7$. All of the rest remains unchanged.

Sorry for the confusion.