

Chapter Seven: Correlation and Regression

Correlation and regression analysis (also called “least squares” or “ordinary least squares (OLS)” analysis) helps us examine relationships among interval or ratio variables. In this chapter, we’ll explore techniques for doing correlation and bivariate regression. Chapter 8 will include a look at multiple correlation and regression.

To illustrate these techniques, we’ll use the “COUNTRIES.sav” file, derived from several sources and containing data on the countries of the world. See Appendix B for a codebook with information on the variables included in this file. Open the file following the instructions in Chapter 1 under “Getting a Data File.”

We’ll begin by considering the relationship between perceived lack of government corruption (in other words, perceived honesty in government) and Internet freedom. Our hypothesis will be that in countries where Internet freedom is high, people will have a greater sense that they can hold government accountable (the technical term for this sense is called “political efficacy”) and they will tend to regard their system as less corrupt. We’ll also add a measure of political rights (rights to “participate freely in the political process”) to the mix, primarily for consideration in Chapter 8.

Perceived honesty in government (*honestgov*) is taken from a measure devised by [Transparency International](#). The Internet Freedom Index (*ifreedom*) and the Political Rights Index (*polrights*) are taken from measures devised by [Freedom House](#). The Civil Liberties Index (*civillib*), which is not used in this chapter but which you may choose to use in an exercise, is also from Freedom House. The *honestgov* and the *ifreedom* indexes range from 0 to 100; the *polrights* and *civillib* indexes are on scales that range from 1 to 7¹.

Correlation

How close are the relationships among Internet freedom, political rights, and perceived honesty in government? To find out, click on **Analyze, Correlate, and Bivariate**. A dialog box will appear on your screen. Click on *honestgov* and then click the arrow to move it into the box. Do the same with *ifreedom* and *polrights*. The dialog box should look like Figure 7–1.

¹ To make using them more intuitive, and more consistent with *honestgov*, *ifreedom*, *polrights*, and *civillib* have been recoded from the original sources so that the **higher** the number, the **higher** the level of perceived honesty in government, Internet freedom, political rights, and civil liberties.



Figure 7-1

The most widely used bivariate test is the Pearson’s r correlation coefficient. It is intended to be used when both variables are measured at either the interval or ratio level and each variable is normally distributed. However, sometimes we do violate these assumptions. If you do histograms of our three variables (see Chapters 4 and 9), you will notice that none are actually normally distributed. Furthermore, variables we are using could arguably be considered ordinal, not interval, measures. We’ll use the Pearson’s r , but will need to proceed with caution. IBM SPSS includes another correlation test, Spearman’s ρ , which is designed to analyze variables that are not normally distributed, or are ranked (i.e., ordinal rather than interval). We will conduct both tests to see how much the results differ depending on the test used—in other words, whether those who use Pearson’s r for these variables are seriously off base.

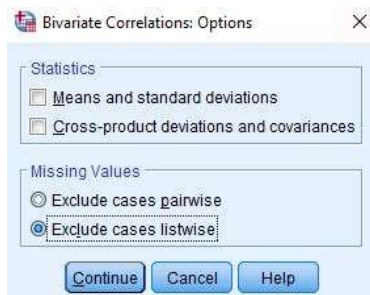


Figure 7-2

In the dialog box, click on **Options** and, in the resulting box, on **Exclude cases listwise**. The result should look like Figure 7–2). The reason for doing this is that, as we’ve noted, *ifreedom* is based on many fewer cases than the other two variables, and we want to be able to make “apples to apples” comparisons. Click on **Continue**.

The box next to Pearson is already checked, as this is the default. Click in the box next to Spearman. Click the button next to **One-tailed** test of significance. (This is because we will be testing “directional” hypotheses, that is, not just the idea that two variables are related but, for example, that the **higher** the value of the *ifreedom* index, the **higher** the value of the *honestgov* index.) Therefore, we would expect the correlation to be positive. Your dialog box should now

look like the one in Figure 7–3. Click OK to run the tests.

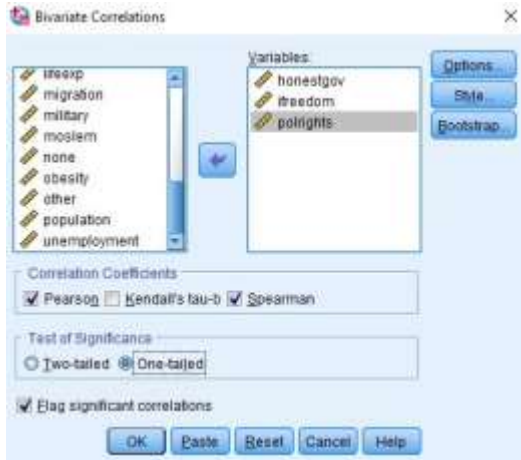


Figure 7-3

Your output screen will show two tables (called matrices): one for Pearson’s r and one for Spearman’s ρ . The Pearson’s correlation matrix should look like the one in Figure 7–4. The cells of the table show the Pearson’s r correlation between each variable and each other variable, the level of statistical significance of the relationship (that is, the likelihood that it could have occurred by chance), and the number of cases on which the correlation is based.

Correlations^b

		honestgov Lack of Perceived Corruption Index	ifreedom Internet Freedom Index	polrights Political Rights Index
honestgov Lack of Perceived Corruption Index	Pearson Correlation	1	-.467**	-.601**
	Sig. (1-tailed)		.001	.000
ifreedom Internet Freedom Index	Pearson Correlation	.467**	1	.830**
	Sig. (1-tailed)	.001		.000
polrights Political Rights Index	Pearson Correlation	.601**	.830**	1
	Sig. (1-tailed)	.000	.000	

** Correlation is significant at the 0.01 level (1-tailed).

b. Listwise N=46

Figure 7-4

The correlation coefficient may range from -1 to 1, where -1 or 1 indicates a “perfect” relationship. The further the coefficient is from 0, regardless of whether it is positive or negative, the stronger the relationship between the two variables. Thus, a coefficient of .467 is exactly as strong as a coefficient of -.467. Positive coefficients tell us there is a direct relationship: when one variable increases, the other increases. Negative coefficients tell us that there is an inverse relationship: when one variable increases, the other decreases. Notice that the Pearson’s r for the relationship between Internet freedom and perceived honesty in government is .467. This tells us that, just as we predicted, as Internet freedom increases, perceived honesty in government increases as well. But should we consider the relationship strong? We’ll revisit this question later in the chapter.

The correlation matrix also gives the probability that the relationship we have found could have occurred just by chance. (Labeled as Sig. [1-tailed]). The probability value is .001, which is well below the conventional threshold of $p \leq .05$. Thus, our hypothesis is supported. There is a relationship (the coefficient is not 0), it is in the predicted direction (positive), and is statistically significant.

Correlations^b

			Honestgov Lack of Perceived Corruption Index	Freedom Without Freedom Index	polrights Political Rights Index
Spearman's rho	Honestgov Lack of Perceived Corruption Index	Correlation Coefficient	1.000	.414**	.539**
		Sig. (1-tailed)		.002	.000
	Freedom Without Freedom Index	Correlation Coefficient	.414**	1.000	.851**
		Sig. (1-tailed)	.002		.000
	polrights Political Rights Index	Correlation Coefficient	.539**	.851**	1.000
		Sig. (1-tailed)	.000	.000	

** Correlation is significant at the 0.01 level (1-tailed).
b. Listwise N = 40

Figure 7-5

Recall that we had some concerns about using the Pearson's r coefficient. Figure 7–5 shows the results using Spearman's rho. Notice that the coefficient for the relationship between *ifreedom* and *honestgov* is .414, or about the same as the value of Pearson's r for this relationship. Similarly, the other values of Spearman's rho are similar to those for Pearson's r. This is reassuring.

Regression

Let's look more closely at the relationship between *ifreedom* and *corruption* graphically by creating a scatterplot. Click on **Graphs, Chart Builder**. This will open up the dialog box shown in Figure 7-6. (If you get a message telling you to be sure that the measurement levels of each variable have been set properly, click on **OK**, since this has already been done for you for the "COUNTRIES.sav" file.)

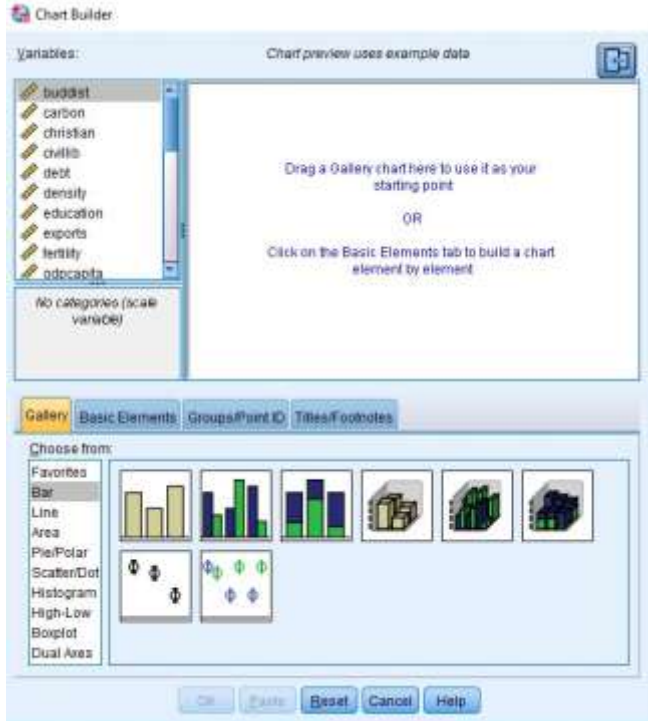


Figure 7-6

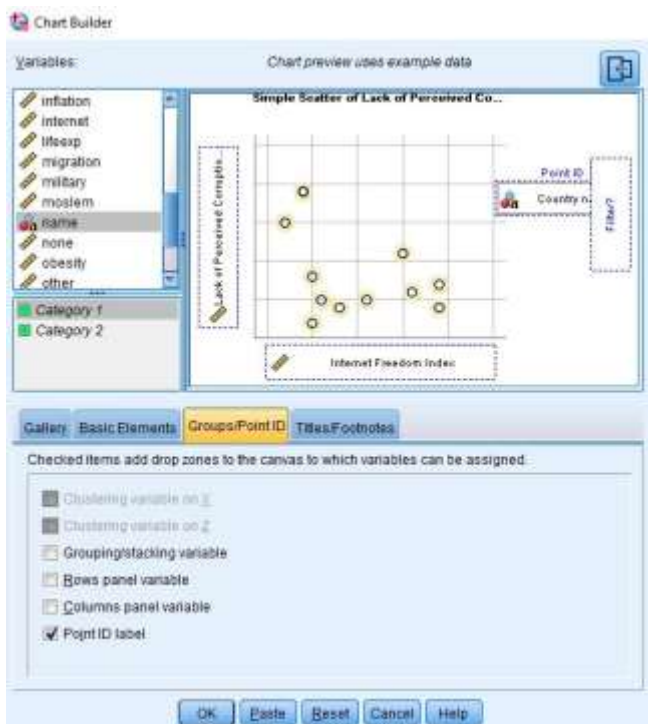


Figure 7-7

Next, in the “Choose from,” list at the lower left, click on Scatter/Dot. Then, shift your attention to the sample graph patterns, and click on the first one (upper left). Holding down the

mouse button, drag the sample to the large chart preview window. Then, add variables to the chart preview window. From the list of variables, click on *ifreedom* and drag it to the box located on the horizontal (X) axis (because it is the independent variable in our hypothesis and the independent variable belongs on the horizontal axis). Next, click on *honestgov* and drag it into the box located on the vertical (Y) axis. Finally, add data labels: from the menu in the middle of the Chart Builder, click on Groups/Point ID, select Point ID label and, from the list of variables, click on *name* and drag it to the box on the chart called “Point Label Variable?” (Note: Point ID labels aren’t a good idea if you have a large number of cases, but will work well here.) Your dialog box should now look like the one in Figure 7–7. Then, click OK. What you see is a plot of perceived honesty in government for each country included in the chart by each country’s level of Internet freedom. Your scatterplot should look like the one in Figure 7–8.




Figure 7-8

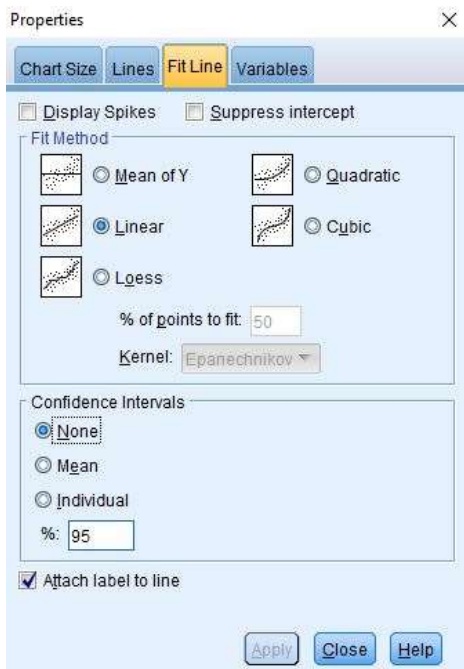


Figure 7-9

useful, since it tells us the “proportional reduction in error” we achieve in “predicting” the value of the dependent variable by knowing that of the independent variable.

How strong a relationship is this? There’s no firm answer to this question. One scholar (Karl Deutsch) once suggested that, if you can explain at least 10% of the variance of a variable, you have something worth talking about. If your r^2 exceeds .5 (that is, it explains over 50% of variance), then your knowledge exceeds your ignorance! We would probably consider anything between an r^2 of .1 and .5 (or an r between about $\pm.3$ and $\pm.7$) to be a moderately strong relationship.

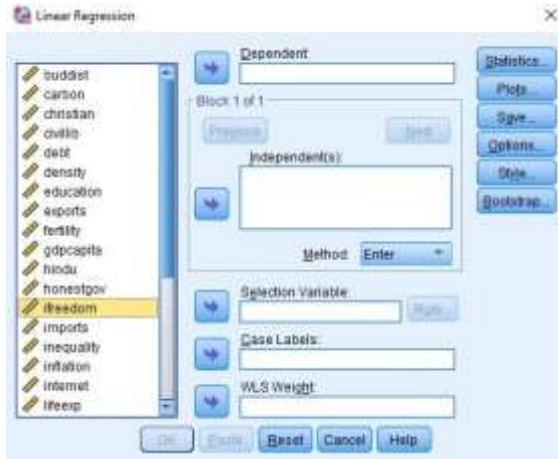


Figure 7-11

Doing a regression analysis can help us to understand the regression line in more detail. Close the IBM SPSS Chart Editor. Click on **Analyze, Regression, and Linear**. This opens up the dialog box shown in Figure 7-11. Move *honestgov* to the **Dependent** box, and *ifreedom* to the **Independent(s)** box. Click **OK**. The results should look like those shown in Figure 7-12.

The first table just shows the variables that have been included in the analysis. The second table, “Model Summary,” shows the R-square statistic, which is .218 (Where have you seen this before? What does it mean?) (Note: the “Adjusted R Square,” .200, is slightly lower because it takes into account the number of independent variables in the equation.) The third table, ANOVA, gives you information about the model as a whole. ANOVA is discussed briefly in Chapter 6. Note that if you take the Regression Sum of Squares (the variance explained by the relationship) and divide by the Total Sum of Squares, the result is equal to R^2 . The final table, Coefficients, gives results of the regression analysis that are not available using only correlation techniques. Look at the “Unstandardized Coefficients” column. Two statistics are reported: “B,” which is the regression coefficient, and the standard error. Notice that there are two statistics reported under B, one labeled as “(Constant),” the other labeled as “ifreedom Internet Freedom Index”. The latter is the regression coefficient, which is the slope of the line that you saw on the scatterplot. (Note that in scholarly reports, it is conventional to refer to the regression coefficient using the lower case, “b.”) The one labeled as (Constant) is not actually a regression coefficient, but is the Y-intercept (IBM SPSS reports it in this column for convenience only).

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.467 ^a	.218	.200	14.410

a. Predictors: (Constant), ifreedom Internet Freedom Index

ANOVA^a

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	2544.374	1	2544.374	12.253	.001 ^b
	Residual	9136.431	44	207.646		
	Total	11680.804	45			

a. Dependent Variable: honestgov Lack of Perceived Corruption Index

b. Predictors: (Constant), ifreedom Internet Freedom Index

Coefficients^a

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	21.164	5.768		3.669	.001
	ifreedom Internet Freedom Index	.348	.099	.467	3.500	.001

a. Dependent Variable: honestgov Lack of Perceived Corruption Index

Figure 7-12

What do these numbers mean? You may recall from your statistics course that the formula for a straight line is:

$$Y = a + bX$$

Y refers to the value of the dependent variable for a given case, a is the Y-intercept (the point where the line crosses the Y-axis, listed as Constant on your output), b is the slope of the line which describes the relationship between the independent and dependent variables, and X is the value of the independent variable for a given case.

We know that the linear relationship between X and Y (*ifreedom* and *honestgov*) is not perfect. The correlation coefficient was not 1 (or -1), and the scatterplot showed plenty of cases that did not fall directly on the line. Thus, it is clear to us that knowing a country's level of Internet freedom will not tell us without fail its level of perceived honesty in government. It is clear that there is some error built into our findings. This is the reason that the regression line is also called the "Best Fit Line." For these reasons, it is conventional to write the formula for the line as:

$\hat{Y} = a + bX + e$, where "e" refers to error. \hat{Y} ("Y hat") indicates the value of Y predicted by the equation for a given case. We could also write it as "Y'" (Y prime) or "Y^c" (the calculated value of Y).

What can we do with this formula? One thing we can do is make predictions about particular values of the dependent variable, using just a little arithmetic. All we have to do is plug the values from our output into the formula for a line. For now, we will ignore the error terms (“e”), but will come back to them shortly. Plugging the numbers from Figure 7–12 into the formula for a straight line, we obtain $\hat{Y}=21.164+.348*X$, the same equation we saw earlier in Figure 7–10, except that, here, numbers have been carried out to three decimal places. We can then plug in the value of X (*ifreedom*) for any given country, multiply by .348, and add that to 21.164. The result will be the predicted value of the *honestgov* variable for that country.

For example, looking at the file in **Data View** mode (see Chapter 2), we see that South Africa, the United Kingdom, and Ukraine all have similar *ifreedom* scores (74, 75, and 73 respectively). Plugging these values into the equation we obtain:

- For South Africa, $\hat{Y}=21.164+.348*74=46.916$.
- For the United Kingdom, $\hat{Y}=21.164+.348*75=47.264$.
- For Ukraine, $\hat{Y}=21.164+.348*73=46.568$.

These numbers represent the predicted values of *honestgov* for these three countries, that is, what the values would be if all three countries fell right on the regression line. In other words, we would predict that, since all three countries have similar *ifreedom* scores, they will also have similar *honestgov* scores. Going back to **Data View**, however, we see that the actual scores are 43, 74, and 26 respectively. If we subtract the predicted scores from the actual scores ($Y-\hat{Y}$), we obtain the “residual,” which is a measure of the error in our prediction for a given case. In this example, the residuals are:

- For South Africa, $Y-\hat{Y}=43-46.916=-3.916$.
- For the United Kingdom $Y-\hat{Y}=74-47.264=26.736$.
- For Ukraine, $Y-\hat{Y}=26-46.568=-20.568$.

In other words, as can be seen in Figure 7-10 above, perceived honesty in government in South Africa is about what we would expect, whereas it is much higher than predicted in the United Kingdom, and much lower than predicted in Ukraine.

We won’t go into it here, but you can, for all cases, add the predicted values of the dependent variable and the residuals as additional variables in the data file. To do this, click on **SAVE** in the regression dialog box, and select **Unstandardized Predicted Values** and **Unstandardized Residuals**.

Chapter Seven Exercises

1. Can you think of any other variables included in the codebook in Appendix B that might help explain levels of perceived government honesty among countries? Repeat the analysis presented in this chapter, but substitute your variable for *ifreedom*.
2. Pick another variable from the codebook (for example, adult obesity rate). Pick another

variable that you think might help explain why some countries have a much higher rate than others. Repeat the analysis presented in this chapter, but substitute your variables for *ifreedom* and *honestgov*.

3. The variables in the General Social Survey are mostly nominal or ordinal, but there are some exceptions. In this exercise, we'll use the data set GSS16A.sav and work with two of these variables, the number of hours per week a respondent reports watching television (*tvhours*), and the respondent's age (*age*).
 - a. It is likely that people of different ages watch different amounts of television. How do you think these may be related? Write a hypothesis that predicts the direction of the relationship between *age* and *tvhours*.
 - b. Do a Pearson correlation to test your hypothesis. Was your hypothesis supported? Explain. Remember that whether or not your hypothesis is supported depends on three things: whether or not the coefficient was 0, whether your prediction of the hypothesized direction of the relationship (+ or -) was correct, and the significance (the probability that you will be wrong if you generalize your finding to the population from which the sample was drawn). Be sure to discuss all three in your explanation.
 - c. Discuss the strength of the relationship between *age* and *tvhours*. Then, speculate about a second factor that might also influence the amount of television that people watch.
 - d. How much of the variance in *tvhours* is explained by *age*? Tell how you found out.
 - e. Do a regression analysis of the relationship between *age* and *tvhours*. Be sure to place your variables into their proper boxes (in other words, correctly identify the independent and dependent variable). If you were writing a scholarly report, how would you describe the relationship between *age* and *tvhours* based on your results? (Hint: If it is small, IBM SPSS may have expressed your regression coefficient in scientific notation in order to save space. If you see something like 2.035E-2 on your IBM SPSS output, that is scientific notation. The E-2 is telling you to move the decimal point two places to the left. Thus, 2.035E-2 becomes .02035. If you don't want to move the decimal yourself, click rapidly several times on the coefficient in the output screen and IBM SPSS will show you the actual value of the coefficient.)
 - f. Do the results of the regression analysis suggest that your hypothesis is supported? Be sure to discuss the magnitude of the regression coefficient, the direction (+ or -), and the probability.
 - g. How many hours of television does your model predict that people aged 21 tend to watch each day? People aged 42? Show how you calculated these predicted scores.
4. Repeat exercise 3, but this time use *income16* as the dependent variable, and *educ* as the independent variable