

2.5: An In-Depth Look At Measures of Association

Measures of association are statistics that tell analysts the strength of the relationship between two (or more) variables, as well as in some cases the direction of that relationship. There are a variety of measures of association; choosing the correct one for any given analysis requires understanding the nature of the variables being used for that analysis. This chapter will detail a number of measures of association that are used by quantitative analysts, though there are others that will not be covered here. While the chapter will not provide full instructions for calculating most measures of association, it aims to give those who are new to quantitative analysis a general understanding of how calculations of measures of association work, how to interpret and understand the results, and how to choose the correct measure of association for a given analysis.

To start, then, what do measures of association tell us? Remember that they do *not* tell us whether a result is statistically significant, as discussed in the chapter on statistical significance. Instead, they are designed to tell us about the nature and strength of the observed relationship between the variables, whether or not that relationship is likely to have occurred by chance. There are different ways of thinking about what association means: for instance, two variables that are strongly associated are those in which the values of one variable tend to co-occur with the values of the other variable. Or we might say that strongly associated variables are those in which variation in one variable can explain much of the variation in another variable. In addition, for analyses using only ordinal and/or continuous variables, some measures of association can tell us about the direction of the relationship—are we observing a direct (positive) relationship, where as the value of x goes up the value of y also goes up, or are we observing an inverse (indirect or negative) relationship, where as the value of x goes up the value of y goes down?

Keep in mind that it is possible for a relationship to appear to have a moderate or even strong association but for that association to not be meaningful in explaining the world. This can occur for a variety of reasons—the relationship may not be significant, and thus the likelihood that the observed pattern occurred by chance could be high. Note that even a $p < 0.001$ there is a one in one-thousand likelihood that the result occurred by chance! Or the relationship may be spurious, and thus while it *appears* that the two variables are associated, this apparent association is only a reflection of the fact that each variable is separately associated with some other variable. Or the strong association may be due to the fact that both variables are basically measuring the same underlying phenomena, rather than measuring separate but related phenomena (for instance, one would observe a very strong relationship between year of birth and age).

There is one other important difference between statistical significance and measures of association: while the computation of statistical significance assumes that data has been collected using a random sample, measures of association do not necessarily require that the data be from a random sample. Thus, for instance, measures of association can be computed for data from a census.

Preparing to Choose a Measure of Association

When choosing a measure of association, analysts must begin by ensuring that they understand how their variable is measured as well as the nature of the question they are asking about their data so that they can choose the measure of association that is best suited to these variables and this question. There are a number of relevant factors to consider.

First, the levels of measurement of the variables that are being used: different measures of association are appropriate for variables of different levels of measurement.

Second, whether information about the direction of the relationship is important to the research question. Some measures of association provide direction and others do not.

Third, whether a symmetric or an asymmetric measure is required. Symmetric measures consider the impact of each variable upon the other, while asymmetric measures are used in circumstances where the analyst wants to use an independent variable to explain or predict variation in a dependent variable. Note that when producing asymmetric measures of association in statistical software, the software will typically produce multiple versions, and the analyst must ensure that they use the one for the correct independent/dependent variable.

Fourth, the number of attributes of each variable (for non-continuous variables). Some measures of association are only appropriate for variables with few attributes—or for crosstabulations in which the resulting tables are relatively small—while others are appropriate for greater numbers of attributes and larger tables.

There are also specific circumstances that are especially suited to particular measures of association based on the nature of the research question or characteristics of the variables being used. And, as will be discussed below, it is essential to understand the way attributes are coded. It is especially important in the case of ordinal and continuous variables to understand whether increasing

numerical values of the variable represent an increase or a decrease in the underlying concept being measured. Finally, there are a variety of factors other than the actual relationship between the variables that can impact the strength of association, including the sample size, unreliable measurements, the presence of outliers, and data that are restricted in range.^[1] Analysts should explore their data using descriptive statistics to see if any of these issues might impact the analysis.

Keep in mind that while it is sometimes appropriate to produce more than one measure of association as part of an analysis, it is not appropriate to simply run all of them and select the one that provides the most desirable result. Instead, the analyst should carefully consider the variables, their question, and the options and choose the one or two most appropriate to the situation to produce and interpret.

General Interpretation of Measures of Association

When interpreting measures of association, there are two piece of information to look for: (1) strength and (2) direction.

Table 1. Strength of Association

Strength	Value
None	0
Weak/Uninteresting	$\pm 0.01-0.09$
Moderate	$\pm 0.10-0.29$
Strong	$\pm 0.30-0.59$
Very Strong	$\pm 0.60-0.99$
Perfect Identity	± 1

The strength of nearly all measures of association ranges from 0 to 1. Zero means there is no observed relationship at all between the two (or more) variables in question—in other words, their values are distributed completely randomly with respect to each other. One would represent what we call a complete identity—in other words, the two variables are measuring the exact same thing and all values line up perfectly. This would be the situation, for instance, if we looked at the association between height in inches and height in centimeters, which are after all just two different ways of measuring the same value. While different researchers do use different scales for assessing the strength of association, Table 1 provides one approach for doing so. Note that very strong values are quite rare in social science, as most social phenomena are too complex for the types of simple explanations where one variable explains most of the variation in another.

The direction of association, where applicable, is determined by whether the measure of association is a positive or negative number—whether the number is positive or negative does not tell us anything about strength (in other words, $+0.5$ is not bigger than -0.5 —they are the *same strength* but a *different direction*). Positive numbers mean a direct association, while negative numbers mean an inverse relationship. Direction cannot be determined when examining relationships involving nominal variables, since nominal variables themselves do not have direction. Keep in mind that it is essential to understand how a variable is coded in order to interpret the direction. For example, imagine we have a variable measuring self-perceived health status. That variable could be coded as 1:poor, 2:fair, 3:good, 4:excellent. Or it could be coded as 1:excellent, 2:good, 3:fair, 4:poor. If we looked at the relationship between the first version of our health variable and age, we might expect that it would be negative, as the numerical value of the health variable would decline as age increased. And if we looked at the relationship between the second version of our health variable and age, we might expect that it would be positive, as the numerical value of the health variable would increase as age increased. The actual health data could be exactly the same in both cases—but if we change the direction of how our variable is coded, this changes the direction of the relationship as well.

Details on Measures of Association

In this section, we will review a variety of measures of association. For each one, we will provide information about the circumstances in which it is most appropriately used and other information necessary to properly interpret it.

Phi

Phi is a measure of association that is used when examining the relationship between two binary variables. Cramer's V and Pearson's r , discussed below, will return values identical to Phi when computed for two binary variables, but it is still more

appropriate to use Phi. It is a symmetric measure, meaning it treats the two variables identically rather than assuming one variable is the independent variable and the other is the dependent variable. It can indicate direction, but given that binary variables are often assigned numerical codes somewhat at random (should yes be 0 and no 1, or should no be 0 and yes 1?), interpretation of the direction may not be of much use. The computation of Phi is the square root of the Chi square value divided by the sample size. While Phi is the most commonly used measure of association for relationships between two binary variables in social science data, there are other measures used in other fields (for instance, risk ratios in epidemiology) that are asymmetric. Yule's Q, discussed in several other chapters, is another example. These will not be discussed here.

Cramer's V

If there is any "default" measure of association, it is probably Cramer's V. Cramer's V is used in situations involving pairs of nominal, ordinal, or binary variables, though not in situations with two binary variables (then Phi is used) and it is less common in situations where both variables are ordinal. It is symmetric and non-directional. The size of the table/number of attributes of each variable does not matter. However, if there is a large difference between the number of columns and the number of rows, Cramer's V may overestimate the association between the variables. It is calculated by dividing the Chi square by the sample size multiplied by whichever is smaller, the number of rows in the table minus one or the number of columns in the table minus one, and then taking the square root of the resulting number.

Contingency Coefficient

The Contingency Coefficient is used for relationships in which at least one of the variables is nominal. It is symmetric and non-directional, and is especially appropriate for large tables (those 5×5 or larger—in other words, circumstances in which both variables have more than five attributes). This is because, for smaller tables, the Contingency Coefficient is not mathematically able to get close to one. It is computed by dividing the Chi square by the number of cases plus the Chi square, and then taking the square root of the result.

Lambda and Goodman & Kruskal's Tau

Lambda is a measure of association used when at least one variable is nominal. It is asymmetric and nondirectional. Some statisticians believe that Lambda is not appropriate for circumstances in which the dependent variable's distribution is skewed. Unlike measures based on the Chi square, Lambda is based on calculating what is called "the proportional reduction in error" (PRE) when one uses the values of the independent variable to predict the values of the dependent variable. The formula for doing this is quite complex, and involves the number of columns and rows in the table, the number of observations in a given row and column, the number of observations in the cell where that row and column intersect, and the total number of observations.

Goodman & Kruskal's Tau works according to similar principles as Lambda, but without consideration of the number of columns and rows. Thus, it is generally advised to use it only for fairly small tables. Like Lambda, it is asymmetric and non-directional. In some statistical software packages (including SPSS), Goodman & Kruskal's Tau is produced when Lambda is produced rather than it being possible to select it separately.

Uncertainty Coefficient

The Uncertainty Coefficient is also used when at least one variable is nominal. It is asymmetric and directional. Conceptually, it measures the reduction in prediction error (or uncertainty) that occurs when one variable is used to predict the other. Some analysts prefer it to Lambda because it better accounts for the entire distribution of the variable, though others find it harder to interpret. As you can imagine, this makes the formula even more complicated than the formula for Lambda; it relies on information about the total number of observations in each row, each column, and each cell.

Spearman

Spearman is used when both variables are ordinal. It is symmetric and directional and can be used for large tables. In SPSS, it can be found under "correlations." Computing Spearman requires converting values into ranks and using the difference in ranks and the sample size in the formula. Note that if there are tied values or if the data is truncated or reduced in range Spearman may not be appropriate.

Gamma and Kendall's Tau (b and c)

The two Kendall's Tau measures are both symmetric and directional and are used for relationships involving two ordinal variables. However, Kendall's Tau b is used when tables are square, meaning that they have the same number of rows and columns, while Kendall's Tau c is used when tables are not square. Like Spearman, Kendall's Tau is based on looking at the relationship between

ranks. After converting values to ranks, one counts pairs of values that are in agreement below a given rank (concordant pairs) and how many are not in agreement (discordant pairs). The formula, then, involves subtracting the number of discordant pairs from the number of concordant pairs, then dividing this number by the number of discordant pairs plus the number of concordant pairs.

Gamma is similar—also symmetric and directional and used for relationships involving two ordinal variables, and with a similar method of calculation, except using same-order and different-order (ranking high or low on both variables versus ranking high on one and low on the other) instead of concordant and discordant pairs. Gamma is preferred when many of the observations in an analysis are tied, as ties are discounted in the computation of Kendall's tau and thus Kendall's tau will produce a more conservative (in other words, lower) value in such cases. However, Gamma may overestimate association for larger tables.

Kappa

Kappa is a measure of association that is especially likely to be used for testing interrater reliability, as it is designed for use when both variables are ordinal with the same categories. It measures agreement between the two variables and is symmetric. Kappa is calculated by subtracting the degree of agreement between the variables that would be expected by chance from the degree of agreement that is observed; subtracting the degree of agreement that would be expected by chance from one, and dividing the former by the latter.

Somers' D

Somers' D is designed for use in examining relationship involving two ordinal variables and is directional, but unlike the other ordinal x ordinal measures of association discussed above, Somers' D is asymmetric. As such, it measures the extent to which our ability to predict values of the dependent variable is improved by knowing the value of the independent variable. It is a conservative measure, underestimating the actual extent to which two variables are associated, though this underestimation declines as table size increases.

Eta

Eta is a measure of association that is used when the independent variable is discrete and the dependent variable is continuous. It is asymmetric and non-directional, and is primarily used as part of a statistical test called ANOVA, which is beyond the scope of this text. In circumstances where independent variables are discrete but not binary, many analysts choose to recode those variables to create multiple dummy variables, as will be discussed in the chapter on multivariate regression, and then use Pearson's R as discussed below.

Pearson's r

Pearson's r is used when examining relationships between two (or more) continuous variables and can also be used in circumstances where an independent variable is binary and a dependent variable is continuous. It is symmetric and directional. The calculation of Pearson's r is quite complex, but conceptually, what this calculation involves is plotting the data on a graph and then finding the line through the graph that best fits this data, a topic that will be further explored in the chapter on Correlation and Regression.

Other Situations

Attentive readers will have noticed that not all possible variable combinations have been addressed above. In particular, circumstances in which the independent variable is continuous and the dependent variable is *not* continuous have not been addressed. For beginning analysts, the most straightforward approach to measuring the association in such relationships is to recode the continuous variable to create an ordinal variable and then proceed with crosstabulation. However, there are a variety of more advanced forms of regression that are beyond the scope of this book, such as logistic regression, that can also handle relationships between these sorts of variables, and there are various pseudo-R measures of association that can be used in such analyses.

Exercises

1. Determine the strength and direction for each of the following measure of association values:
 - o -0.06
 - o 0.54
 - o 0.13
 - o -0.27

2. Select the most appropriate measure of association for each of the following relationships, and explain why it is the most appropriate:
 - Age, measured in years, and weight, measured in pounds
 - Opinion about the local police on a 5-point agree/disagree scale and highest educational degree earned
 - Whether or not respondents have health insurance (yes/no) and whether or not they have been to a doctor in the past 12 months (yes/no)
 - Letter grade on Paper 1 and letter grade on Paper 2 in a first-year composition class
3. Explain, in your own words, the difference between association and significance.

-
1. For instance, a study looking at the relationship between age and health that only included people between the ages of 23 and 27 would be restricted in range in terms of age. ↩

This page titled [2.5: An In-Depth Look At Measures of Association](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by [Mikaila Mariel Lemonik Arthur](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.