I'm Less and Less Happy until Finally I'm Sad:
Estimating Correlations When Variables Divide a Construct into Parts

Kimberly A. Barchard
University of Nevada, Las Vegas

James A. Russell
Boston College

UNLV

BOSTON COLLEGE

Contact: Kimberly Barchard, kim.barchard@unlv.edu
Website: http://barchard.faculty.unlv.edu/examining-opposites/

Abstract

When exact opposites divide a construct into two parts, each variable is 0 when the other is positive: They do not correlate -1. We have created a method of estimating the correlation between the two constructs when each of their measures may not cover all parts of those constructs.

Introduction

Testing whether two variables – such as happiness and sadness – are best represented as separate dimensions or as opposite parts of a single bipolar dimension is controversial. Imagine that happiness and sadness are two variables that form a single bipolar dimension. When that bipolar dimension is measured as a whole (with response options from 1 = *very happy* to 5 = *very sad*), it correlates -1 with its reverse (which uses response options 1 = *very sad* to 5 = *very happy*). As the first variable increases, the second always decreases. See Figure 1. However, when the bipolar dimension is divided into two parts that are measured separately, their correlation is far from -1. Imagine we ask participants, "How happy are you?" (response options 0 = *Not at all* to 10 = *Extremely*), and, "How sad are you?" (same response options). If happiness and sadness are mutually exclusive parts of a single bipolar construct, then any degree of happiness implies a complete lack of sadness and vice versa: Each variable is 0 whenever the other is positive. See Figure 2. This makes it impossible for the variables to correlate -1. For example, if these two variables divide their bipolar dimension exactly in half, and if scores on the full bipolar dimension have a normal distribution, these variables correlate -.467 (Russell & Carroll, 1999).

Cacioppo and Bernstein (1994) noted that we cannot test if a variable is bipolar by using bipolar items: Measuring happy-sad as a single bipolar variable presupposes its bipolarity. Therefore, they argued that researchers need to divide potentially bipolar variables into two parts and measure each part separately. Still, researchers mistakenly assumed that the correlation between the two parts should be approximately -1. When they found correlations near -.467, they concluded the two were not parts of a single bipolar dimension.

To determine if two variables are opposite parts of a single bipolar dimension, researchers need a method of testing hypotheses about the correlation between underlying whole variables, X and Y, given empirical scores x and y covering only parts of those dimensions (see Symbols, below). We have invented a method of doing this. We refer to this method as *Inverse Inference*.

<div align="center">

Inverse Inference

</div>

The Inverse Inference method can be used for inference on a wide variety of statistics whose sampling distributions are difficult to calculate. The broad idea is that we calculate the statistic in a randomly generated sample that meets the requirements of the null hypothesis – a sample that is large enough to be a good stand-in for the population parameter. Then we evaluate how close our sample statistic is to the statistic that resulted when the null hypothesis was true. As far as we know, no one has proposed this approach to inference before. Here, we explain how inverse inference can be used to test hypotheses about the correlation between the whole concepts based upon data from measures that capture only parts of those constructs. See Table 1 for a summary of the steps.

The first step is to calculate the correlation between the two observed variables, x and y, and then calculate a 95% confidence interval for the population correlation, ρ_{xy}. A convenient confidence interval calculator is given at http://vassarstats.net/rho.html   For example, imagine that $r_{xy} = -.320$ in a sample of 400 participants. This website calculator shows that a 95% confidence interval for ρ_{xy} is given by the interval [-.405, -.230]. This calculator uses the Fisher r-to-z transformation. If you want to use a different type of confidence interval (e.g., bootstrap), you can do so.

In psychological research, our variables typically do not have perfect reliability, which reduces the correlation between the variables. Therefore, the second step is to correct the end points of the confidence interval for lack of reliability. This will create a confidence interval for the correlation between perfectly reliable variables. To correct the end points of the confidence interval, use the correction for attenuation formula (Spearman, 1904):

$$r_{xy\,corrected} = \frac{r_{xy}}{\sqrt{r_{xx'}r_{yy'}}}$$

where

$r_{xy\,corrected}$ is the estimate of what the correlation between x and y would have been, if they had both been perfectly reliable (or the estimate of one of the end points of the confidence interval for this correlation)

r_{xy} is the Pearson product-moment correlation between x and y (or one of the end points of the confidence interval for this correlation)

$r_{xx'}$ is an estimate of the reliability of x

$r_{yy'}$ is an estimate of the reliability of y

For example, we previously calculated a 95% confidence interval for ρ_{xy} of [-.405, -.230]. Imagine that the reliability of x is estimated as .70 and the reliability of y is estimated as .60. In this case, a 95% confidence interval for the correlation between perfectly reliable x and y, $\rho_{xy,corrected}$, is given by the interval [-.625, -.355].

This confidence interval takes into account sampling error (by forming a confidence interval rather than just providing a point estimate) and measurement error (by correcting the end points of the confidence interval for attenuation due to lack of reliability). However, it has not

corrected for the effect of censoring. Therefore, the third step is to use the Excel file to see what effect censoring had on the correlation.

Doing this takes three steps. First, we estimate the proportions of the constructs covered by each variable. For example, imagine that 30% of the scores on happiness are 0 and 70% of the scores are non-zero: We estimate that x covers 70% of X (and that the remaining 30% of scores were censored). Similarly, if 60% of the scores on sadness are 0, we estimate that y covers 40% of Y. We enter these proportions into our Excel file. Next, we select a null hypothesis to test. If we want to test whether happiness (X) and sadness (Y) are opposites, we hypothesize that $\rho_{XY} = -1$. We enter this number into our Excel file. Third, now that Excel has this information, it randomly generates 10,000 cases (or 500,000 cases) for which X and Y have the specified correlation and the specified proportions of censoring, and which are also normally distributed. Excel calculates r_{xy} for those 10,000 cases. For example, if we hypothesize $\rho_{XY} = -1$ and indicate that x covers 70% of X, and y covers 40% of Y, Excel calculates r_{xy} = -.538. We refer to this sample correlation as $r_{xy_{null}}$ to indicate that it was calculated from randomly generated data for which the null hypothesis was true.

Fourth, to test if our sample data are consistent with the hypothesized value of ρ_{XY}, we compare the Excel-generated $r_{xy_{null}}$ (for which the null hypothesis is true) with the 95% confidence interval for the correlation between perfectly reliable x and y (which we calculated above based upon the observed x and y data and using the correction for attenuation formula). If the Excel-generated value for $r_{xy_{null}}$ falls inside the confidence interval, we retain the null hypothesis. In our case, the Excel-generated value of -.538 does indeed fall within the 95% confidence interval of [-.625, -.355]. Therefore, we retain the null hypothesis that ρ_{XY} = -1. In summary, given the severe censoring in our data (60% on one variable and 30% on the other), our sample correlation (r_{xy}) of -.32 is consistent with the null hypothesis that the correlation between perfectly reliable uncensored variables (ρ_{XY}) is -1.

Fifth, in addition to conducting hypothesis tests about ρ_{XY}, we might also want to create a confidence interval for ρ_{XY}. We already have the confidence interval for the correlation between perfectly reliable censored variables, x and y: $\rho_{xy_{corrected}}$. We will use it to create a confidence interval for the correlation between perfectly reliable uncensored variables, X and Y: ρ_{XY}. To do this, we vary the hypothesized value of ρ_{XY} in the Excel file until we find Excel-generated $r_{xy_{null}}$ values that just barely fit within the 95% confidence interval for $r_{xy_{corrected}}$. For example, we previously found an empirical confidence interval for perfectly reliable x and y variables, $r_{xy_{corrected}}$, of [-.625, -.355]. We enter different hypothesized values for ρ_{XY} into the Excel file, until we find the values of ρ_{XY} that lead to Excel-generated values of $r_{xy_{null}}$ that are equal to -.625 and -.355. First I tried to find the upper limit on the confidence interval. I found that ρ_{XY} = -.55 resulted in r_{xy} = -.355. Therefore, the upper limit for the confidence interval for ρ_{XY} is -.355. Next I tried to find the lower limit on the confidence interval. I found that ρ_{XY} = -1 resulted in r_{xy} = -.544. Because ρ_{XY} cannot be less than -1, the lower limit for the confidence interval for ρ_{XY} is set to -1. Thus, the 95% confidence interval for ρ_{XY} is given by the interval [-1, -.55]. In summary, given the severe censoring in our data, our sample correlation (r_{xy}) of -.32 is consistent with the null hypothesis that the correlation between perfectly reliable uncensored variables (ρ_{XY}) is -1, but this sample correlation is also consistent with X and Y having only a moderate-to-small relationship.

Future research will be needed to determine how well the inverse inference method works for estimating ρ_{XY} based upon x and y variables that only cover parts of X and Y. In our example calculations, we imagined that we had a sample of 400 participants. However, the resulting confidence interval for ρ_{XY} ranged from a perfect negative relationship of -1 to a moderate-to-small relationship of -.355. Therefore, future research will be needed to determine if this method provides confidence intervals that accurately control coverage (i.e., do the 95% confidence intervals capture ρ_{XY} 95% of the time?) and to determine if this method is more or less powerful than alternative methods of dealing with this type of censored data (i.e., Do the hypothesis tests have higher or lower power for the same datasets? Equivalently, are the confidence intervals wider or narrower?).

Future research can also explore other uses of the inverse inference method. For example, if researchers are interested in some other measure of association (e.g., ICC(A,1), CCC(C,1)), they could follow the same steps but substitute the statistic of interest. Moreover, the general approach of inverse inference could be used for any statistic that has a difficult-to-derive sampling distribution (e.g., alpha-if-item-deleted).

References

Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin, 115*, 401-423. https://doi.org/10.1037/0033-2909.115.3.401

Russell, J. A., & Carroll, J. M. (1999). On the bipolarity of positive and negative affect. *Psychological Bulletin, 125*, 3-30. https://doi.org/0.1037/0033-2909.125.1.3

Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology*, *15*(1), 72–101. https://doi.org/10.2307/1412159

Symbols

Symbol	Meaning
x	One of the observed variables; a censored measure of X; x covers part of X; x is not perfectly reliable.
y	One of the observed variables; a censored measure of Y; y covers part of Y; y is not perfectly reliable.
X	A perfectly reliable theoretical variable, covering the whole of one of the constructs of interest.
Y	A perfectly reliable theoretical variable, covering the whole of one of the constructs of interest.
ρ_{XY}	The Pearson product-moment correlation between X and Y in the population. This is the parameter we want to estimate. To test if X and Y are opposites, we test the null hypothesis $\rho_{XY} = -1$.
r_{xy}	The Pearson product-moment correlation between x and y in the sample; an estimate of ρ_{xy}, the correlation in the population.
$r_{xx'}$	An estimate of the reliability of x.
$r_{yy'}$	An estimate of the reliability of y.
$r_{xy_{corrected}}$	The sample estimate for what the correlation between x and y would have been, if they had both been perfectly reliable. $r_{xy_{corrected}}$ is an estimate of $\rho_{xy_{corrected}}$, this same correlation in the population, for which we form a confidence interval in step 2.
$r_{xy_{null}}$	The sample value of r_{xy} that was generated by Excel in step 3, using random data from 10,000 (or 500,000 cases) for which the null hypothesis was true.

Table 1

Steps in Testing Hypothesis and Forming Confidence Intervals for ρ_{XY}

| Step	Action
1	Calculate r_{xy}. Correct for sampling error by calculating a confidence interval for ρ_{xy}.
2	Correct for attenuation due to lack of reliability by calculating a confidence interval for the correlation between perfectly reliable x and y.

3	Estimate the effect of censoring on the correlation by using Excel to generate a sample value of $r_{xy_{null}}$ when ρ_{XY} equals the hypothesized value.
	
	    a. Estimate the proportions of the constructs covered by each variable.
	    b. Select a null hypothesis to test, such as $\rho_{XY} = -1$.
	    c. Enter these values into our Excel file

4	To test hypotheses about ρ_{XY}, compare Excel-generated $r_{xy_{null}}$ to the confidence interval generated in step 2.
5	To create a confidence interval for ρ_{XY}, vary the values of ρ_{XY} in Excel to find $r_{xy_{null}}$ values that match the end points of the confidence interval generated in step 2.

Note. To obtain the necessary Excel files, see http://barchard.faculty.unlv.edu/examining-opposites/

Figure 1
Opposites that each measure the whole dimension
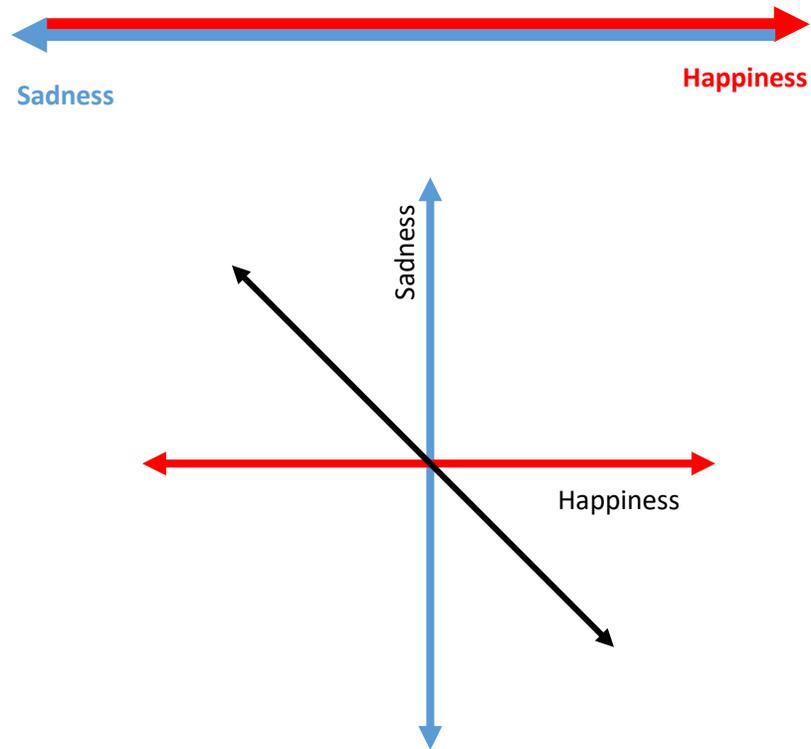
Note. If two variables, X and Y, each cover the **whole** of a bipolar concept, then to assert that they are opposites is to assert that the full continuum for one is the same as the continuum for the other, but that they go in opposite directions.

Figure 2

Opposites that each measure a part of the dimension
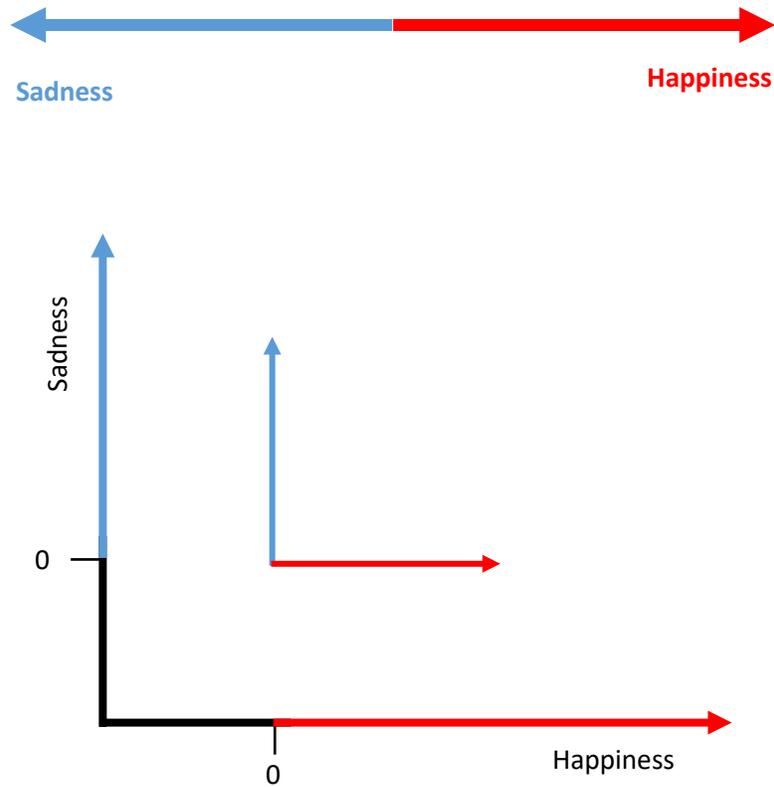
Note. If two concepts cover mutually exclusive **parts** of a bipolar concept, then to assert that they are opposites is to assert that the two parts together cover the entire continuum and are non-overlapping: Any degree of one implies the absence of the other.