Detecting a Lack of Association: An Equivalence Testing Approach

Jason Goertzen & Robert A. Cribbie

Department of Psychology

York University

Address correspondence concerning this article to Rob Cribbie (cribbie@yorku.ca).

Abstract

Researchers often test for a lack of association between variables. A lack of association is usually established by demonstrating a nonsignificant relationship with a traditional test (e.g., Pearson's *r*). However, for logical, as well as statistical, reasons, such conclusions are problematic. In this paper, we discuss and compare the empirical Type I error and power rates of three lack of association tests. The results indicate that large, sometimes very large, sample sizes are required for the test statistics to be appropriate. What is especially problematic is that the required sample sizes may exceed what is practically feasible for the conditions that are expected to be common among researchers in Psychology. This paper highlights the importance of using available lack of association tests, instead of traditional tests of association, for demonstrating the independence of variables, and qualifies the conditions under which these tests are appropriate.

Detecting a Lack of Association: An Equivalence Testing Approach

Researchers in psychology are frequently interested in testing for a lack of association between variables. For example, researchers running linear models with several covariates may be interested in limiting the number of covariates by removing those that are not related to the outcome variable. Other researchers test for a lack of association as the primary research hypothesis. For example, Wheadon et al. (1992) explored the potential lack of association between suicidality and the use of fluoxetine in the treatment of bulimia nervosa. However, the main problem with testing for a lack of association between variables is that the absence of a relationship is usually demonstrated by a *nonsignificant* test statistic (e.g., Pearson's *r*). This strategy is not appropriate because the research hypothesis (i.e., a lack of association) is aligned with the null hypothesis (e.g., H_o: $\rho = 0$) rather than with the alternative hypothesis. In other words, a traditional test of correlation was designed to test for the presence of an association, not a lack of association. In short, if researchers wish to appropriately detect a lack of association, they must utilize lack of association tests.

The objectives of the present paper are to: 1) outline some of the challenges researchers face when attempting to detect a lack of association; and 2) propose methods for detecting a lack of association. In addition, we hope to raise awareness regarding how lack of association decisions are currently being carried out in psychology and how these problematic tactics can be redressed. To the best of our knowledge, the present paper is the first study dealing with lack of association tests within the field of psychology.

Problems Related to the Current Approach to Detecting a Lack of Association

As outlined in the introduction, currently the most common method for detecting a lack of association between two variables is to look for a nonsignificant relationship with a traditional correlation or regression test statistic. However, this method is not appropriate for two main reasons: 1) not rejecting the null hypothesis of a lack of association (i.e., H_o: $\rho = 0$) does mean that we accept the null hypothesis; and 2) a nonsignificant test of correlation has poor asymptotic properties for demonstrating a lack of association because the probability of declaring a lack of association (i.e., not rejecting H_o: $\rho = 0$) decreases (instead of increases) as sample sizes increase. The first point simply reminds researchers that we cannot ‘accept’ the null hypothesis when our test statistic is not significant, as a relationship may be present but our test statistic may simply have not detected it (possibly due to a lack of power). The second point highlights that when a researcher conducts a traditional test of correlation, the likelihood of detecting a significant relationship increases as our sample sizes increase. In other words, more power for detecting an association in the population is attained as the size of the sample increases. Conversely then, the likelihood of detecting a nonsignificant relationship decreases as sample sizes increase, and therefore this is not a good method for detecting a lack of association. In other words, researchers attempting to find a lack of association among variables with a traditional test of association would increase the likelihood of detecting a lack of association (i.e., a nonsignificant test statistic) by using a smaller sample size, and researchers would declare equivalence almost 100% of the time with N = 3 (the lowest value of N for which the degrees of freedom are positive).

In sum, these two main problems represent challenges that need to be overcome through the development of new statistical tests or the adoption of alternative methods. With regard to alternative methods, one potential alternative is Bayesian analysis. Bayesian approaches address the probability of a hypothesis given the data, a position that could be useful for evaluating a lack of association by determining the probability of the null hypothesis in a traditional test of association (i.e., H_o: ρ = 0) being true. However, although Bayesian and other approaches may offer unique solutions to the problem of detecting a lack of association, the focus of this study is on the development of null hypothesis testing (i.e., frequentist) based lack of association procedures.

Another Significant Challenge: The Distribution of Pearson's r

When setting out to develop an appropriate test for demonstrating a lack of association, it is important to have an appreciation for the distribution of Pearson's *r* when there is no underlying relationship between the variables of interest in the population (i.e., ρ = 0). For example, we simulated 5000 Pearson's *r* statistics when ρ = 0 for each of seven different sample size conditions (N = 10, 15, 20, 25, 50, 100 and 200). Table 1 presents the proportion of sample correlations of each magnitude for each of the sample size conditions. It is important to point out that we are not focusing on statistical significance in this example, but instead on the magnitude of the correlation, as we expect that that would be the most important factor in trying to delineate whether a relationship exists among two variables. What is evident from the table is that it would be extremely difficult to 'prove' a lack of association when sample sizes are small to moderate. With a sample size of N = 10, approximately 75% of sample correlation values (in absolute value) exceed *r* = .1, 40% exceed *r* = .3, and even 15% exceed *r* = .5. With larger sample sizes there is still substantial variability in correlation coefficients. With a sample size of N = 20, approximately 40% of sample correlation values exceed *r* = .2, and 20% exceed *r* = .3. Even with a sample size of N = 100, more than 30% of sample correlation values exceed *r* = .1.

Why is this such a problem? Consider that you are trying to demonstrate a lack of association between two variables; from these results it would appear that more than 100 or 200 participants would be necessary in order to appropriately demonstrate the absence of any meaningful relationship given the extreme variability in sample *r* statistics. In other words, attempting to demonstrate a lack of association is tantamount to trying to determine that an effect size is negligible, and as one would expect this is very difficult to achieve. We expand on this topic below when we discuss establishing an interval of negligible association. Further, as a result of the difficulty in demonstrating a lack of association with small to moderate sample sizes, one of the purposes of the present paper is to determine the sample sizes that lack of association tests would require in order to produce acceptable results.

A Test of Lack of Association

The first lack of association test (equiv_r) that we propose is based on the 'two one-sided tests' approach that is common in mean equivalence testing (e.g., Westlake, 1981; Schuirmann, 1981, 1987) and is similar to regression based equivalence tests proposed by Robinson, Duursma and Marshall (2005) and Dixon and Pechmann (2005) (although the latter is focused on a log-linear model that would not be as widely applicable in psychology). The composite null hypotheses, H_{o1}: $\rho > \rho^*$ and H_{o2}: $\rho < -\rho^*$, are rejected if $t_1 \leq -t_{\alpha, N-2}$ and $t_2 \geq t_{\alpha, N-2}$, where:

ρ^* represents the lack of association interval (i.e., $-\rho^*$, ρ^*), N represents the sample size, $t_{\alpha, N-2}$ represents the α level critical value from the *t* distribution with N-2 degrees of freedom, and *r* represents the sample correlation value. Simultaneous rejection of H_{o1} and H_{o2} implies that the population correlation falls within the bounds $-\rho^*$ to ρ^*. As is evident from the equations above, lack of association tests make the potential presence of an association the null hypothesis, and the alternative hypothesis becomes H_1: $-\rho^* < \rho < \rho^*$.

An important advantage of lack of association tests, relative to using a traditional test designed to detect an association, is that researchers are forced to specify a lack of association interval ($-\rho^*$, ρ^*). A lack of association interval specifies the bounds for which the correlation is deemed meaningless (i.e., essentially zero). For example, a researcher may set ρ* = .1 (and hence an interval of -.1 to .1), where in this case any correlation less than .1 (in absolute value) would be considered insignificant within the framework of the study. It is important for researchers to consider the nature of the study when setting ρ*, as what is considered practically insignificant can vary considerably from study to study.

Potential Issues with the Lack of Association Test

It is fairly well known that the *t* formula for null hypothesis testing with Pearson's *r* is biased because the standard error of the test statistic is related to ρ (Fisher, 1915; Bond & Richardson, 2004). In other words, the *t* statistic:

is generally not appropriate for hypothesis testing with Pearson's *r*. However, this formula is the root of the lack of association test proposed above. The most common correction for the bias is Fisher's (1915) *z* transformation, which converts *r* to *z*, and the test statistic (F_z) is z / s_z , where:

With this test H_o: $\rho = 0$ is rejected if $|F_z| \geq z_\alpha$, where z_α is the α level critical value from the standard normal distribution. Therefore, a potentially improved lack of association test (equiv_fz) could result from utilizing Fisher's *z* transformation with the previously proposed lack of association test (i.e., equiv_r). The resulting test would reject H_{o1}: $\rho > \rho^*$ if $F_{Z\_LA1} \leq -z_\alpha$, and H_{o2}: $\rho < -\rho^*$ if $F_{Z\_LA2} \geq z_\alpha$, where:

.

A Resampling Based Approach to Lack of Association Tests

A third approach for generating an unbiased lack of association test is to resample (with replacement) N paired data points from the original N paired observations (equiv_rs). This approach is similar to the resampling based procedure developed by Robinson et al. (2005) for evaluating equivalence in regression based problems. This process is repeated many times, each time calculating the correlation coefficient between the paired data. An empirical sampling distribution of *r* is therefore generated from the sample data. The resulting test would reject H_{o1}: $\rho > \rho^*$ if $RS_{r, .95} - \rho^* \leq 0$ and H_{o2}: $\rho < -\rho^*$ if $RS_{r, .05} \geq 0$, where RS_r is the distribution of sample correlation values calculated from the resampled paired data. $RS_{r, .05}$ and $RS_{r, .95}$ represent the 5th and 95th quantiles from the empirically derived sampling distribution of *r*. We recommend that at least 10000 resamples be conducted in order to ensure an appropriate test statistic.

Current Study

The primary purpose of the present study is to evaluate the performance of the three aforementioned tests for detecting a lack of association (i.e., equiv_r, equiv_fz, and equiv_rs) in order to be able to make recommendations regarding best practices for conducting lack of association tests. Specifically, the authors are interested in recommendations regarding appropriate sample sizes for achieving appropriate results. Furthermore, the authors are interested in recommending which of the three tests can achieve appropriate results with the smallest sample size, which may be an important practical consideration for many researchers.

Method

A Monte Carlo study was used to evaluate the Type I error control and power of all three proposed tests (equiv_r, equiv_fz, and equiv_rs). Ten-thousand resamples were used in calculating the resampling based lack of association test (equiv_rs). Three primary variables were manipulated in this study: 1) lack of association interval; 2) population correlation; and 3) sample size. ρ^* was set at .05, .1, .15, .2, .25 or .3, with ρ set equal to ρ^* for Type I error conditions and at $\rho = 0$ or $\rho = .05$ for evaluating power (note that when assessing power with $\rho = .05$, the $\rho^* = .05$ condition was not investigated because it would replicate the Type I error results). Sample sizes were set at N = 50, 100, 500 and 1000.

Two random normal variates (X, Y) were generated for each simulation, where $X = a*b + e_1$, and $Y = a*b + e_2$. X and Y were generated to have population correlation ρ ($\rho > 0$), where a, e_1 and e_2 are random normal variates ($\mu = 0$, $\sigma = 1$) that were generated using the R generator "rnorm" (R Development Core Team, 2005), and:

Each of the conditions was crossed, and 5000 simulations were computed for each condition with a nominal Type I error rate of $\alpha = .05$. The simulation program was written in R (R Development Core Team, 2005).

Results

Type I error Rates

Type I error rates for the equiv_r, equiv_fz and equiv_rs procedure under each of the conditions investigated are presented in Table 2. The results indicate that the Type I error rates for the resampling based lack of association test (equiv_rs) were most accurate, followed by Fisher's *z* transformed procedure (equiv_fz), and the original lack of association test (equiv_r) . However, more generally the results indicate that none of the procedures had accurate Type I error rates when both the sample sizes and the lack of association intervals were small. More specifically, all of the procedures were extremely conservative when both sample sizes and lack of association intervals were small, with the resampling based procedure the least conservative. For example, with a sample size of N = 100, a lack of association interval of at least -.2 to .2 would be required in order to have an appropriate test of the hypothesis (i.e., an empirical Type I error rate of at least .04). Even with a sample size of N = 500, a lack of association interval of at least -.15 to .15 would be required in order to have an empirical Type I error rate that was close to the nominal rate. It is especially noteworthy that when the association interval was set at -.05 to .05, there were no false rejections (i.e., declarations of a lack of association) for any of the three tests, even when N = 1000.

Power

Power rates for the equiv_r, equiv_fz, and equiv_rs procedures, under each of the conditions investigated, are presented in Table 3.The results for $\rho = 0$ were very similar to those for $\rho = .05$ and therefore the following discussion will deal simultaneously with both sets of results. The power results very closely mirror the results for the Type I error rates, with all of the procedures extremely conservative when both sample sizes and lack of association intervals were small. More specifically, when sample sizes were moderate (e.g., N = 100), satisfactory power (e.g., > .8) can only be attained if a lack of association interval of at least -.3 to .3 is adopted. Even with a sample size of 500, researchers would need to use a lack of association interval of at least -.15 to .15 in order to have acceptable power. Again, the resampling based procedure (equiv_rs) was uniformly most powerful; however, the power advantage over the equiv_r and equiv_fz procedures was not large in most cases.

Discussion

As researchers become aware of the availability of lack of association tests, it is important that recommendations are available for how to conduct these tests in an appropriate manner. The primary goal of this study was to evaluate and compare the performance of three lack of association tests: equiv_r, equiv_fz, and equiv_rs. With regard to comparing the performance of the three lack of association tests, the resampling based procedure (equiv_rs) had more accurate Type I error rates and uniformly greater power. However, the result that stands out in this study is the finding that very large sample sizes (e.g., n > 500; see Table 2) are required in order to ensure that the lack of association tests have sufficient power with a lack of association interval narrower than -.25 to .25. This is very important because, theoretically, it becomes very difficult to state that two variables are unrelated with a lack of association interval as wide as -.25 to .25 (or wider).

Another way to frame these results is to consider that, for most psychological studies, a sample correlation approaching .3 would be considered practically, as well as statistically, significant in most cases. However, when attempting to conduct a test of lack of association with a lack of association interval of -.25 to .25 or narrower, the sample sizes required for appropriate results are very large (practically, for many psychology researchers). Since we expect common lack of association intervals within many psychological studies to be $\rho^* = .05$ to $\rho^* = .20$, researchers should be aware that the sample sizes required for producing appropriate results with these intervals may be prohibitively large (i.e., $N \geq 500$).

It is also important to highlight that, although available lack of association tests are not optimal (i.e., overly conservative), reverting back to traditional tests of association (e.g., Pearson's *r*) is not appropriate. For example, imagine that one wanted to determine if the association between two variables fell within the association interval of $\rho^* = .10$. We simulated Pearson's *r* statistics for $\rho = .2$ and found that, with N = 50, 71.4% of the test statistics were not significant. In other words, a researcher who incorrectly used a traditional association test to demonstrate a lack of association would declare the variables independent approximately 71% of the time, even though the population correlation falls outside of the association interval (note that the association interval does not even play a part in evaluating a lack of association with a traditional test of association). Even with N = 100, researchers would declare the variables independent approximately 50% of the time. Therefore, traditional tests of association are extremely inappropriate for assessing a lack of association, especially when sample sizes are small. In fact, Type I error rates for assessing a lack of association with traditional tests of association will approach unity as sample sizes decrease. In order to facilitate the use of lack of association tests we have included an R function in the Appendix that will generate test statistics for all the methods described in this paper.

As a result of the extreme conservativeness of the lack of association test statistics with small to moderate sample sizes, we suggest that the current study should represent a starting point for research on lack of association tests within psychology. The results of the present study represent current best practices for producing appropriate results, which importantly include a call to immediately halt lack of association conclusions that are reached from a failed test of association. However, in following these best practices, many researchers may find that they are unable to conduct their desired test because of the sample sizes required for the current procedures. Therefore, we recommend that additional research be carried out on lack of association tests; in particular the results of this study highlight the importance of developing lack of association tests that that can provide acceptable power with small to moderate sample sizes.

References

Bond, C. F., & Richardson, K. (2004). Seeing the Fisher *z*-transformation. Psychometrika, 69, 291-303.

Dixon, P. M., & Pechmann, J. H. K. (2005). A statistical test to show negligible trend. *Ecology, 86*, 1751-1756.

Fisher, R.A. (1915). Frequency distribution of the values of the correlation coefficient in samples of an indefinitely large population. *Biometrika, 10*, 507-521.

Robinson, A. P., Duursma, R. A., & Marshall, J. D. (2005). A regression-based equivalence test for model validation: shifting the burden of proof. *Tree Physiology, 25*, 903-913.

Schuirmann, D. J. (1981). On hypothesis testing to determine if the mean of a normal distribution is contained in a known interval. *Biometrics, 37*, 617.

Schuirmann, D. J. (1987). A comparison of the two one-sided tests procedure and the power approach for assessing equivalence of average bioavailability. *Journal of Pharmacokinetics and Biopharmaceutics, 15*, 657-680.

Wheadon, D. E., Rampey, A. H., Thompson, V. L., Potvin, J. H., Masica, D. N., & Beasley, C. M. (1992). Lack of association between fluoxetine and suicidality in bulimia nervosa. *Journal of Clinical Psychiatry, 53*, 235-241.

Westlake, W. J. (1981). Response to T.B.L. Kirkwood: bioequivalence testing - a need to rethink. *Biometrics, 37*, 589-594.

Table 1
Proportion of Sample Correlations that Exceed Designated Values of *r* as a Function of Sample Size when $\rho = 0$

	Sample Correlation (*r*) Magnitude (Absolute Value)									
N	>.9	>.8	>.7	>.6	>.5	>.4	>.3	>.2	>.1	>0
10	.001	.007	.025	.067	.140	.252	.398	.578	.780	1
15	0	<.001	.005	.019	.060	.139	.273	.465	.720	1
20	0	0	.001	.005	.025	.082	.199	.400	.687	1
25	0	0	<.001	.002	.012	.046	.145	.334	.628	1
50	0	0	0	0	<.001	.004	.036	.163	.492	1
100	0	0	0	0	0	0	.003	.048	.323	1
200	0	0	0	0	0	0	0	.005	.159	1

Table 2

Type I Error Rates for Lack of Association Tests as a Function of Sample Size and ρ.*

		N=50			N=100			N=500		N=1000		
ρ*	eq_r	eq_fz	eq_rs	eq_r	eq_fz	eq_rs	eq_r	eq_fz	eq_rs	eq_r	eq_fz	eq_rs
.05	0	0	0	0	0	0	0	0	0	0	0	0
.10	0	0	0	0	0	0	.049	.050	.050	.051	.052	.051
.15	0	0	0	0	0	.002	.047	.049	.049	.045	.047	.049
.20	0	0	.005	.030	.034	.042	.046	.051	.053	.049	.052	.054
.25	.009	.017	.035	.039	.048	.051	.043	.049	.050	.046	.050	.051
.3	.033	.041	.052	.043	.050	.052	.044	.050	.052	.040	.048	.047

Note: ρ* represents the lack of association interval; eq_r is the original lack of association test; eq_fz = modified lack of association test based on Fisher's *z* transformation; eq_rs = resampling based lack of association test.

Table 3

Power Rates for Lack of Association Tests as a Function of Sample Size, ρ, and ρ.*

	N=50			N=100			N=500			N=1000		
ρ*	eq_r	eq_fz	eq_rs	eq_r	eq_fz	eq_rs	eq_r	eq_fz	eq_rs	eq_r	eq_fz	eq_rs
							ρ=0					
.05	0	0	0	0	0	0	0	0	0	0	0	.003
.10	0	0	0	0	0	0	.425	.432	.434	.874	.874	.875
.15	0	0	0	0	0	0	.914	.917	.917	.998	.998	.999
.20	0	0	.008	.237	.262	.279	.995	.995	.995	1	1	1
.25	.044	.085	.137	.591	.621	.622	1	1	1	1	1	1
.3	.315	.365	.372	.820	.847	.849	1	1	1	1	1	1
							ρ=.05					
.10	0	0	0	0	0	0	.244	.247	.250	.475	.477	.477
.15	0	0	0	0	0	.008	.719	.724	.724	.934	.938	.938
.20	0	0	.007	.234	.257	.262	.957	.960	.959	.999	.999	.999
.25	.039	.075	.122	.540	.568	.570	.999	.999	.999	1	1	1
.3	.289	.340	.358	.756	.780	.785	1	1	1	1	1	1

Note: ρ* represents the lack of association interval; eq_r is the original lack of association test; eq_fz = modified lack of association test based on Fisher's *z* transformation; eq_rs = resampling based lack of association test.

Appendix

The following function was created to run a traditional Pearson correlation test, an equivalence based test of lack of association, an equivalence based test of lack of association using a Fisher's z transformation, and an equivalence based test of lack of association with resampling in the R software package. R is open source software that is available at www.r-project.org. To utilize the function, first run the entire function in R. In other words, copy and paste the syntax below into R. The program will run but you will not see any output at this time, you have just defined the function. At this point you will need to have opened your dataset (or otherwise defined your variables) in R. If you require assistance with opening a dataset or defining variables in R see the help files and manuals available at the R website listed above. Next, at the R prompt, run the equiv_corr function by typing (without the parentheses) "equiv_corr(v1,v2,equivint)". You would substitute the names of your two variables for 'v1' and 'v2', and your equivalence interval for 'equivint'. You can also change the alpha level by changing the fourth argument in the function. For example, after you have run the entire function in R, you might enter (without the parentheses) "equiv_corr(x,y,.2, .1)" at the prompt in R to run the statistical tests on the variables x and y with an equivalence interval of .2.

```
equiv_corr<-function (var1,var2, equivint, alpha=.05, na.rm=TRUE, ...) {
   if (na.rm)     x <- x[!is.na(var1)]
   if (na.rm)     y <- y[!is.na(var2)]
      corxy<-cor(var1,var2)
      n<-length(var1)
      nresamples<-10000
      #### Running a traditional t test to determine if the correlation is significant######
      t<-corxy/(sqrt((1-corxy^2)/(n-2)))
      pvalue_tradt<-1-pt(abs(t),n-2)
      ifelse (pvalue_tradt<=alpha,
         decis_tradt<-"The null hypothesis that there is no correlation between x and y can be rejected.",
         decis_tradt<-"The null hypothesis that there is no correlation between x and y cannot be rejected.")
      #### Running an original two t-test procedure for equivalence #######
      equivt1<-(corxy-equivint)/sqrt((1-corxy^2)/(n-2))
      pvalue1_equivt<-pt(equivt1,n-2)
      equivt2<-(corxy+equivint)/sqrt((1-corxy^2)/(n-2))
      pvalue2_equivt<-1-pt(equivt2,n-2)
      ifelse (pvalue1_equivt<=alpha & pvalue2_equivt<=alpha,
         decis_equivt<- "The null hypothesis that the correlation between var1 and var2 falls outside of the
equivalence interval can be rejected.",
        decis_equivt<-" The null hypothesis that the correlation between var1 and var2 falls outside of the equivalence
interval cannot be rejected.")
      ##### Run a two t-test procedure for equivlance with Fisher's z transformation ####
      zei<-log((1+equivint)/(1-equivint))/2
      zcorxy<-log((1+corxy)/(1-corxy))/2
      equivt1_fz<-(zcorxy-zei)/(1/sqrt(n-3))
      pvalue1_fz<-pnorm(equivt1_fz)
      equivt2_fz<-(zcorxy+zei)/(1/sqrt(n-3))
      pvalue2_fz<-1-pnorm(equivt2_fz)
      ifelse (pvalue1_fz<=alpha & pvalue2_fz<=alpha,
         decis_fz<-"The null hypothesis that the correlation between var1 and var2 falls outside of the equivalence
interval can be rejected.",
```

```
        decis_fz<-"The null hypothesis that the correlation between var1 and var2 falls outside of the equivalence
interval cannot be rejected.")
       #### Run the resampling version of the two t-test procedure for equivalence #####
       resamp <-function(x,m=10000,theta, conf.level= 0.95, ...)
         {
           n <- length(x)
           Data<- matrix(sample(x, size=n*m,replace=T), nrow=m)
           thetastar <- apply(Data, 1, theta, ...)
           M <- mean(thetastar)
           S <- sd(thetastar)
           alpha <- 1-conf.level
           CI <- quantile(thetastar, c(alpha/2, 1-alpha/2))
           return(list(ThetaStar=thetastar, Mean.ThetaStar=M, S.E.ThetaStar=S, Percentile.CI=CI))
           }
           matr<-cbind(var1,var2)
           mat<-as.matrix(matr)
           theta <- function(x,mat)
           {
           cor(mat[x,1], mat[x,2])
         }
       results<-resamp(x=1:n,m=nresamples, theta=theta, mat=mat)
       q1<-quantile(results$ThetaStar,alpha)
       q2<-quantile(results$ThetaStar,1-alpha)
       q1negei<-q1-equivint
       q2negei<-q2-equivint
       q1posei<-q1+equivint
       q2posei<-q2+equivint
       ifelse (q2negei<0 & q1posei>0,
          decis_rs<- "The null hypothesis that the correlation between var1 and var2 falls outside of the equivalence
interval can be rejected.",
        decis_rs<-"The null hypothesis that the correlation between var1 and var2 falls outside of the equivalence
interval cannot be rejected.")
       #### Summary #####
        title1<-"Traditional Test of Correlation, Ho: rho=0"
        title2<-"Equivalence Based Test of Lack of Association"
        title3<-"Equivalence Based Test of Lack of Association with Fisher's z transformation"
        title4<-"Equivalence Based Test of Lack of Association with Resampling"
        stats_tradt<-c(corxy,t,n-2,pvalue_tradt,decis_tradt)
        names(stats_tradt)<-c("Pearson r","t-statistic","df","p-value","Decision")
        stats_equivt<-c(corxy,equivint,equivt1,pvalue1_equivt,equivt2,pvalue2_equivt,n-2,decis_equivt)
        names(stats_equivt)<-c("Pearson r","Equivalence Interval","t-stat 1","pval_t1","t-stat 2", "pval_t2",
"df","Decision")
        stats_fz<-c(corxy,equivint,equivt1_fz,pvalue1_fz,equivt2_fz,pvalue2_fz,decis_fz)
        names(stats_fz)<-c("Pearson r","Equivalence Interval","z-stat 1","pval_z1","z-stat 2", "pval_z2", "Decision")
        stats_rs<-c(corxy,equivint,nresamples,q1,q2,decis_rs)
        names(stats_rs)<-c("Pearson r", "Equivalence Interval","# of Resamples","100(alpha)
Percentile","100(1-alpha) Percentile", "Decision")
        out<-list (title1,stats_tradt,title2,stats_equivt,title3,stats_fz,title4,stats_rs)
        out
}
```