Published in final edited form as:

Struct Equ Modeling. 2016 ; 23(6): 921–930. doi:10.1080/10705511.2016.1202771.

tetrad: A Set of Stata Commands for Confirmatory Tetrad Analysis

Shawn Bauldry

University of Alabama-Birmingham

Kenneth A. Bollen

University of North Carolina-Chapel Hill

Abstract

This article provides a brief overview of Confirmatory Tetrad Analysis (CTA) and presents a new set of Stata commands for conducting CTA with supporting examples. The Stata command, *tetrad*, allows researchers to use modelimplied vanishing tetrads to test the overall fit of structural equation models (SEMs) with continuous endogenous variables and the relative fit of two SEMs with continuous endogenous variables that are tetrad-nested. An extension of the command, *tetrad_matrix*, allows researchers to conduct CTA using a sample covariance matrix as input rather than relying on raw data. In addition, researchers can use the *tetrad_matrix* command to input a polychoric correlation matrix and conduct CTA for SEMs involving dichotomous, ordinal, or censored outcomes. An additional extension of the command, *tetrad_bootstrap*, provides a bootstrapped p-value for the chi-square test statistic. By drawing on Stata's recently developed suite of commands for structural equation modeling, researchers can integrate CTA with data preparation, likelihood ratio tests for model fit, and the estimation of model parameters in a single statistical software package.

Keywords

tetrad; Confirmatory Tetrad Analysis; Stata

Confirmatory Tetrad Analysis (CTA) is a complementary method of testing and comparing the fit of structural equation models (SEMs) with the commonly used likelihood ratio tests (Bollen, 1990; Bollen and Ting, 1993, 1998, 2000; Hipp and Bollen, 2003). CTA has a number of desirable properties supporting its use. First, some models that are not identified can still be assessed based on their implied vanishing tetrads (e.g., some models with causal indicators; Bollen and Ting, 2000). Second, some sets of models are nested in their vanishing tetrads even though they are not nested in the more traditional approaches necessary for likelihood ratio tests. Third, researchers can apply CTA testing to the whole model or to components of the model of particular interest.

*Corresponding author: Shawn Bauldry, Department of Sociology, University of Alabama-Birmingham, HHB 460, 1720 2nd Ave South, Birmingham, AL 35294-1152, (205) 934-0565, sbauldry@uab.edu.

Despite these desirable properties, CTA is underutilized in empirical studies. In part, this lack of use likely reflects the limited availability of software for CTA. SAS macros have been developed for conducting CTA (Hipp et al., 2005; Ting, 1997), but they have not been updated for the current version of SAS and they do not include the option of obtaining a bootstrapped p-value for the chi-square test statistic. The new set of Stata commands for CTA presented below coupled with Stata's recently developed suite of commands for structural equation models allow researchers to prepare data, estimate model parameters and obtain traditional SEM fit statistics, and conduct CTA in a single statistical software package.

The purpose of this paper is to give a brief overview of CTA and to describe and illustrate the new Stata commands for conducting CTA.

Confirmatory Tetrad Analysis

Confirmatory Tetrad Analysis uses tetrads to test the overall fit of a given structural equation model. Tetrads are formed from sets of four random variables as the differences between two product pairs of covariances of observed variables. Following Kelley's (1928) notation, population tetrads are defined as

$$\tau_{ijkl} = \sigma_{ij}\sigma_{kl} - \sigma_{ik}\sigma_{jl}, \quad (1)$$

where σ_{ij} is the covariance between the i^{th} and j^{th} variables of the population covariance matrix Σ . A hypothesized structural equation model with an implied covariance structure $\Sigma(\theta)$ given a vector of parameters θ will often imply that some of the tetrads will equal 0 (i.e., vanish). This fact can be used to derive a statistical test with the null hypothesis that the vanishing tetrads implied by the model are zero versus the alternative hypothesis that at least one of these vanishing tetrads is not zero. Because it is the structure of the model that implies vanishing tetrads, a correctly specified model should have these vanishing tetrads hold. Rejection of the null hypothesis thus implies that there is something wrong with the hypothesized model. In other words, rejecting the null hypothesis provides evidence that the hypothesized model is misspecified.

The chi square vanishing tetrad test statistic is given by

$$T = n\hat{\tau}'\mathbf{A}'\hat{\Sigma}_{\hat{\tau}}^{-1}\mathbf{A}\hat{\tau} \quad (2)$$

where T is the scalar test statistic, n is the sample size, $\hat{\tau}$ is the sample vector of vanishing tetrad values for a given model, $\mathbf{A}$ is a selection matrix that uses 0s and 1s to select from among the tetrads, and $\hat{\Sigma}_{\hat{\tau}}^{-1}$ is the inverse of the asymptotic covariance matrix of the sample estimates of the tetrads that correspond to a particular model. The selection matrix $\mathbf{A}$ has t^* rows and t columns where t^* is the number of nonredundant vanishing tetrads and t is the total number of vanishing tetrads implied by the model.

The new Stata tetrad commands provide options for two test statistics, the original vanishing tetrad test statistic (Bollen, 1990) and a modified tetrad test statistic (Johnson and Bodner, 2007). The original vanishing tetrad test statistic, T_0 , from Bollen (1990) uses,

$$T_0 = n\hat{\tau}'\mathbf{A}'\hat{\Sigma}_{\hat{\tau}(0)}^{-1}\mathbf{A}\hat{\tau} \quad (3)$$

with

$$\hat{\Sigma}_{\hat{\tau}(0)} = \mathbf{A}\mathbf{D}(\hat{\sigma})'\Sigma_s\mathbf{D}(\hat{\sigma})\mathbf{A}', \quad (4)$$

where $\mathbf{D}(\hat{\sigma}) = \partial\hat{\tau}/\partial\hat{\sigma}$ and Σ_s is the estimated asymptotic covariance matrix of the sample covariances.

Johnson and Bodner (2007) propose a modified test statistic, T_1 , which is

$$T_1 = n\hat{\tau}'\hat{\Sigma}_{\hat{\tau}(1)}^{-1}\hat{\tau} \quad (5)$$

with

$$\hat{\Sigma}_{\hat{\tau}(1)} = \text{diag}[\mathbf{D}(\hat{\sigma})'\Sigma_s\mathbf{D}(\hat{\sigma})], \quad (6)$$

where all vanishing tetrads are included but only the inverses of the variances of each tetrad appear in $\hat{\Sigma}_{\hat{\tau}(1)}$. The T_1 test statistic is computed on data transformed to be consistent with the null hypothesis that the model-implied vanishing tetrads are zero (Bollen and Ting, 1998). Redundant tetrads are not eliminated and this along with the diagonal weight matrix eases the computational burden as the number of variables increases (see example of computation times below).

CTA proceeds in three steps. The first step involves finding the tetrads that are implied to be vanishing based on a given model specification. The tetrad command identifies vanishing tetrads based on an empirical procedure outlined by Bollen and Ting (1993). The procedure involves selecting an arbitrary set of values for all of the parameters in a given model (the actual parameter estimates are a good choice), forming the implied covariance matrix with the set of parameter values, and then using the implied covariances to form all of the tetrads among the given set of variables. Those tetrads that are within rounding error of zero are then treated as vanishing tetrads. This procedure appears to be highly accurate as long as a sufficient number of significant digits are represented in the implied covariance matrix (i.e., typically 7–8 digits are recommended).

The second step is to identify a set of nonredundant vanishing tetrads (note that this step is not necessary for computing T_1). For many models some of the vanishing tetrads are

redundant (see Bollen and Ting, 1993, for examples). To compute the test statistic given in equation (3) one must select a set of nonredundant tetrads in order to be able to calculate the inverse of the asymptotic covariance matrix. Although asymptotically equivalent, it is possible that different sets of vanishing tetrads may yield different results in finite samples. Due to this possibility, Hipp and Bollen (2003) recommend randomly selecting sets of vanishing tetrads multiple times and assessing the sensitivity of the results to different selections. The tetrad command uses the sweep operator to identify sets of nonredundant vanishing tetrads and allows users to specify a desired number of replications that randomize the sets of nonredundant vanishing tetrads (Goodnight, 1979; Hipp and Bollen, 2003; Hipp et al., 2005). The third step is to compute the test statistic- T_0 or T_1 .

Nested Tetrad Tests

In addition to testing the fit of a single model, CTA can be used to test the comparative fit of two models that are nested in terms of the tetrads that each imply are vanishing. Suppose $\Sigma(\theta^{(1)})$ and $\Sigma(\theta^{(2)})$ are the implied covariance matrices for two hypothesized models and $\tau^{(1)}$ and $\tau^{(2)}$ are the respective tetrads for each model. If the vanishing tetrads implied by model 2 are a strict subset of the tetrads implied by model 1, then model 2 is said to be tetrad-nested in model 1. With two tetrad-nested models, the difference in chi-square test statistics itself follows a chi-square distribution with degrees of freedom equal to the difference in degrees of freedom between the two models. In other words, $T_{diff} = T_0^{(1)} - T_0^{(2)}$ with $df_{diff} = df^{(1)} - df^{(2)}$. Rejection of the null hypothesis provides evidence that model 2, the less restrictive model in terms of vanishing tetrads, has a better fit with the data than model 1, the more restrictive model in terms of vanishing tetrads. Examples 1 and 2 below illustrate the use of the new Stata commands for nested-tetrad tests.

Dichotomous, Ordinal, or Censored Outcomes

To conduct CTA for SEMs that include dichotomous, ordinal, or censored outcomes analysts must first obtain a polychoric correlation matrix for all of the observed variables and the associated asymptotic covariance matrix for the polychoric correlations.¹ The polychoric correlation matrix is then used in place of the sample covariance matrix in determining tetrads as in the sample analogue of equation (1). The asymptotic covariance matrix for the polychoric correlations is substituted for the asymptotic covariance matrix of the sample covariances in equation (4). With these two substitutions, CTA proceeds as described above (see Hipp and Bollen (2003) for an extended discussion).

Statistical Power

When conducting CTA a researcher may be interested in assessing the power of the statistical test. For instance, knowing the power of the statistical test may aid in the interpretation of a significant (or nonsignificant) test statistic. The key to evaluating the

¹Stata 14 does not currently provide an asymptotic covariance matrix for polychoric correlations. Researchers who plan to conduct CTA for a SEM that involves dichotomous, ordinal, or censored outcomes will need to obtain the polychoric correlations and associated asymptotic covariance matrix from another software package. Mplus 7 was used in the example below.

statistical power for a hypothesized structural equation model lies in formulating an alternative model. The statistical power thus reflects the ability to detect a significant difference between the hypothesized model and a specific alternative model.

Bollen and Ting (1993) outline five steps in calculating the power of a tetrad test. Step 1: determine a set of values for the parameters in the alternative model (θ_a). Step 2: generate the implied covariance matrix $\Sigma(\theta_a)$. Step 3: form the vector of nonredundant vanishing tetrads (τ_a) under H_0 using $\Sigma(\theta_a)$ in place of the sample covariance matrix. Step 4: determine a noncentrality value by substituting τ_a in place of τ in equation (2), i.e.

$$T = n\hat{\tau}_a' \mathbf{A}' \hat{\Sigma}_{\hat{\tau}}^{-1} \mathbf{A} \hat{\tau}_a. \quad (7)$$

Step 5: calculate the power of the tetrad test statistic based on the degrees of freedom (the number of nonredundant tetrads in τ_a), the desired Type I probability, and the noncentrality value calculated in Step 4. We illustrate how to calculate the statistical power of a tetrad test using the new Stata commands below.

Overview of Stata Commands

This section provides an overview of the three new Stata commands for conducting CTA (StataCorp, 2015b). The commands can be installed within Stata by typing (identifying link omitted) without the quotes in Stata's command window. Alternatively the commands can be accessed at (identifying link omitted) and downloaded for manual installation.

The tetrad Command

The syntax for tetrad is

```
tetrad varlist [if] [in], icm1(name) [icm2(name) reps(#) seed(#)]
tlist(1 = yes)]
```

The tetrad command requires a minimum of four numeric variables in the *varlist* and the name of an implied covariance matrix for a given model specification (*icm1(name)*).² The easiest way to obtain an implied covariance matrix is to estimate the model using Stata's *sem* command, request the display of estimation results using *estat framework, fitted*, and store the implied covariance matrix returned as *r(Sigma)* (StataCorp, 2015a). We provide examples of this procedure below. It is also possible to enter an implied covariance matrix using Stata's matrix commands.

The tetrad command has four optional arguments. The first option allows users to provide the name of an implied covariance matrix for a second model that is tetrad-nested in the first

²If an analyst enters an implied covariance matrix for a model with zero implied vanishing tetrads, the tetrad program returns a note that it finds no vanishing tetrads for the given model.

model.³ If such a model is provided, then the tetrad output will include the results of a nested tetrad test. The second option allows users to request a given number of replications of the tetrad test in which the order of the vanishing tetrads is randomly determined. The third option allows users to set a random number seed in order to ensure that the same results can be recovered from running the command at a later time. The fourth option includes a table in the output of all of the tetrads for the specified implied covariance matrix along with which tetrads are empirically determined to be vanishing.

The `tetrad_bootstrap` Command

The syntax for `tetrad_bootstrap` to obtain the T_1 adjusted test statistic described above is

```
tetrad_bootstrap varlist [if] [in], icm1(name) [icm2(name) reps(#)  
seed(#)]
```

For this command `reps` refers to the number of bootstrap replications with a default of 1000. Users also have the option of specifying a random number seed to recover the same estimates at a later time. If the user specifies a second implied covariance matrix that is tetrad-nested in the first implied covariance matrix, then the `tetrad_bootstrap` command returns T_1 for the nested-tetrad test.

The `tetrad_matrix` Command

In some cases, users may wish to perform CTA by entering matrices rather than working with raw data. One particular case is when a user specifies a SEM in which at least some of the outcomes are dichotomous, ordinal, or censored variables. The `tetrad_matrix` command is designed to facilitate CTA with dichotomous, ordinal, or censored outcomes, but it can also be used in place of the `tetrad` command for SEMs with continuous outcomes. The `tetrad_matrix` command requires users to input the sample size, a sample covariance (or correlation) matrix, and at least one implied covariance matrix for a given set of variables and a given model. The syntax for `tetrad_matrix` is

```
tetrad_matrix , obs(#) scm(name) icm1(name) [icm2(name) pcacm(name)  
reps(#) seed(#) tlist(1 = yes)]
```

The various matrices can be entered directly using Stata's `matrix` command. For a SEM with a mixture of dichotomous, ordinal, or censored outcomes, users should enter the polychoric correlation matrix with `scm(name)` and then provide the asymptotic covariance matrix for the polychoric correlations with the `pcacm(name)` option. All of the other options are the same as in the `tetrad` command. It is not possible to obtain a bootstrapped test statistic when working with matrices rather than raw data.

³If an analyst enters an implied covariance matrix for a second model that is not tetrad-nested in the first model, the `tetrad` program returns a note that the models are not tetrad-nested.

Examples

In this section we provide five examples of CTA using both simulated and real data. The first two examples feature the `tetrad` command with simulated data. The first illustrates the basic functionality of the command while the second illustrates the use of CTA for a model in which standard likelihood ratio tests may not be available. The third example features the `tetrad_matrix` command and illustrates its use for a model with a mixture of dichotomous and ordinal outcomes. The fourth example features the `tetrad_bootstrap` command and involves testing the fit of a multi-trait multi-method model. The third and fourth examples use real data. The fifth example illustrates the use of the `tetrad_matrix` command to calculate the statistical power of a tetrad test with respect to an alternative model.

Example 1: CFA vs. MIMIC

In this example the population model is given by a latent variable measured by five indicators. As an analyst we consider two models: (1) the true model and (2) a MIMIC model with two predictors of the latent variable and three indicators of the latent variable (see Figure 1). We simulate data for 500 cases based on the population model and perform a CTA comparing models 1 and 2. Given that model 1 is the true model, we expect a non-significant chi-square test statistic for this model and a non-significant chi-square test statistic for the nested model test comparing models 1 and 2.

The following Stata code estimates the structural equation models that correspond to models 1 and 2 and then uses the `tetrad` command to conduct a CTA for both models and a nested tetrad test.

```
. qui sem (Xi -> x1 x2 x3 x4 x5)
. qui estat framework, fitted
. mat sigma1 = r(Sigma)
. qui sem (Xi -> x3 x4 x5) (x1 x2 -> Xi)
. qui estat framework, fitted
. mat sigma2 = r(Sigma)
. tetrad x1 x2 x3 x4 x5, icml(sigma1) icm2(sigma2) reps(5) tlist(1)
```

Because the option `tlist` was selected, the first portion of the output from running `tetrad` provides a list of the tetrads implied by model 1 and which are vanishing. The first column provides the tetrad labels, the second column provides the tetrad residual, the third column provides the asymptotic variance for the given tetrad, the fourth column provides the Z-value used to determine whether the tetrad is vanishing, and the fifth column provides an indicator (1 = yes, 0 = no) for whether the tetrad is treated as vanishing. The structure of a CFA with five indicators, one latent variable, and no correlated errors implies that all of the tetrads vanish, which is what is indicated in the table of tetrads for Model 1.

Model-Implied Tetrads

Model 1				
tetrad	residual	AVar	Z-value	vanish
1234	0.000000	0.000380	0.0000	1
1342	0.000000	0.000380	0.0000	1
1423	0.000000	0.000381	0.0000	1
1235	-0.000000	0.000385	-0.0000	1
1352	0.000000	0.000385	0.0000	1
1523	0.000000	0.000386	0.0000	1
1245	-0.000000	0.000396	-0.0000	1
1452	0.000000	0.000396	0.0000	1
1524	0.000000	0.000396	0.0000	1
1345	0.000000	0.000391	0.0000	1
1453	0.000000	0.000391	0.0000	1
1534	0.000000	0.000390	0.0000	1
2345	0.000000	0.000390	0.0000	1
2453	0.000000	0.000390	0.0000	1
2534	0.000000	0.000388	0.0000	1

The second portion of the output lists the tetrads implied by model 2 and which are vanishing. The structure of a MIMIC implies that only some of the tetrad vanish, which is also seen in the table of tetrads for Model 2. For instance, the empirical check for vanishing tetrads indicates that τ_{1234} is not vanishing under the given model structure.

Model-Implied Tetrads

Model 2				
tetrad	residual	AVar	Z-value	vanish
1234	0.003192	0.000383	0.1632	0
1342	-0.000000	0.000383	-0.0000	1
1423	-0.003192	0.000377	-0.1645	0
1235	0.003157	0.000387	0.1605	0
1352	0.000000	0.000387	0.0000	1
1523	-0.003157	0.000382	-0.1616	0
1245	0.003072	0.000399	0.1538	0
1452	0.000000	0.000399	0.0000	1
1524	-0.003072	0.000391	-0.1554	0
1345	0.000000	0.000391	0.0000	1
1453	0.000000	0.000391	0.0000	1
1534	-0.000000	0.000390	-0.0000	1
2345	0.000000	0.000390	0.0000	1
2453	-0.000000	0.000390	-0.0000	1

Model 2				
tetrad	residual	AVar	Z-value	vanish
2534	-0.000000	0.000388	-0.0000	1

The final section of the output provides the CTA test statistics. The first column indicates the replication number, the next three provide the chi square test statistic (T_0 in this case), the degrees of freedom, and the p-value for the first model, the second set of three columns provide the same information for the second model, and the third set of three columns provides the results for the nested tetrad test. Because the option reps was specified with 5 replications, the output reports 5 test statistics based on randomly selected sets of vanishing tetrads. The results indicate that in all 5 replications the test statistics for Model 1, Model 2, and the nested test are all non-significant-as expected given that Model 1 is the population model.

Confirmatory Tetrad Analysis Results

Model 1				Model 2			M1 – M2		
rep	Chi-sq	df	p-val	Chi-sq	df	p-val	Chi-sq	df	p-val
1	6.4887	5	0.2615	6.4455	4	0.1683	0.0432	1	0.8353
2	6.5036	5	0.2602	6.4730	4	0.1665	0.0307	1	0.8609
3	6.5413	5	0.2570	6.4122	4	0.1704	0.1291	1	0.7193
4	6.5859	5	0.2533	6.3444	4	0.1749	0.2415	1	0.6231
5	6.5512	5	0.2562	6.4192	4	0.1700	0.1320	1	0.7164

Example 2: CTA Preferred to LRT

Our second example illustrates the use of CTA to compare two models for which the standard likelihood ratio test may not be available. In this example the population model is a two factor, seven indicator model that includes a cross-loading on the fourth indicator (see Figure 2).⁴ A researcher is interested in testing the dimensionality of the model-in particular whether a one factor or two factor model has a better fit with the data. This can be explored by constraining the covariance between the two factors to 1 and the factor loading from the second factor to the fourth indicator to 0. CTA may be preferred in this context because the constraint on the covariance between the two factors leads to a situation in which one is comparing models with different numbers of latent dimensions.

We simulate data for 500 cases based on the population model (with $Cov(\xi_1, \xi_2) = .75$, all $\lambda = .7$, all $Var(e) = .5$) and perform a CTA comparing the two models. Note that in this case the model with two factors and the cross-loading is tetrad-nested within the model with one factor. The following Stata code estimates the structural equation models that correspond to one- and two-factor models and then uses the tetrad command to conduct a CTA for both

⁴This simulation model is based on an empirical example presented in Hipp et al. (2005).

models and a nested tetrad test. For this example we request 5 replications and do not invoke the option for obtaining a list of all of the tetrads.

```
. qui sem (Xi -> x1 x2 x3 x4 x5 x6 x7)
. qui estat framework, fitted
. mat sigmal = r(Sigma)
. qui sem (Xi1 -> x1 x2 x3 x4) (Xi2 -> x4 x5 x6 x7)
. qui estat framework, fitted
. mat sigma2 = r(Sigma)
. tetrad x1 x2 x3 x4 x5 x6 x7, icml(sigmal) icm2(sigma2) reps(5)
```

As expected, the results indicate (1) that the model with one factor is not consistent with the data, (2) that the model with two factors is consistent with the data, and (3) the model with two factors represents a significant improvement over the model with one factor. This pattern of results is evident in all five replications.

Confirmatory Tetrad Analysis Results

rep	Model 1				Model 2				M1 – M2		
	Chi-sq	df	p-val		Chi-sq	df	p-val		Chi-sq	df	p-val
1	33.1278	14	0.0028		8.7041	12	0.7280		24.4237	2	0.0000
2	35.5956	14	0.0012		9.0042	12	0.7026		26.5915	2	0.0000
3	34.1335	14	0.0020		8.8248	12	0.7178		25.3087	2	0.0000
4	35.1844	14	0.0014		8.9307	12	0.7088		26.2537	2	0.0000
5	34.5335	14	0.0017		9.4909	12	0.6605		25.0425	2	0.0000

Example 3: Dichotomous and Ordinal Outcomes

Our third example draws on data from the 2014 wave of the General Social Survey (Smith et al., 2015) to illustrate CTA for a model involving a mixture of dichotomous and ordinal outcomes. The model is a measurement model with a latent variable representing traditional gender attitudes and four effect indicators. The first indicator is “Tell me if you agree or disagree with this statement: Most men are better suited emotionally for politics than are most women” with responses disagree or agree. The next three indicators are all ordinal measures with responses strongly agree, agree, disagree, or strongly disagree. The first is “A working mother can establish just as warm and secure a relationship with her children as a mother who does not work.” The second is “A preschool child is likely to suffer if his or her mother works.” The third is “It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family.” Data are available for 1,562 respondents.

The polychoric correlation matrix and asymptotic covariance matrix for the polychoric correlations were obtained in Mplus 7 (Muthén and Muthén, 2012) and entered into Stata using Stata’s matrix command. The implied covariance matrix was obtained in Stata by specifying the SEM as a four-indicator CFA with continuous covariates. Row and column

labels were added to the polychoric correlation matrix using Stata's `rownames` and `colnames` commands in order to ensure that the variables in the polychoric correlation match the variables in the implied covariance matrix. Users must take care that the asymptotic covariance matrix for the polychoric correlation matrix is in the proper order to correspond with the variables in the polychoric correlation matrix and the implied covariance matrix. The `tetrad_matrix` command assumes that the polychoric correlations are from below the main diagonal and have been read in column-wise. For instance, in this example the order of the polychoric correlations is assumed to be σ_{21} , σ_{31} , σ_{41} , σ_{32} , σ_{42} , σ_{43} and the asymptotic covariance matrix is input as a 6×6 matrix. In addition, if using `Mplus`, the entries of the asymptotic covariance matrix need to be multiplied by the sample size. Finally, when entering (or reading in) matrices into Stata it is important to maintain a high degree of precision (i.e., 7–8 significant digits).

The following Stata code invokes the command with the defined matrices *pcm* for the polychoric correlation matrix, *pcacm* for the asymptotic covariance matrix for the polychoric correlations, and *sigma* for the implied covariance matrix.

```
. tetrad_matrix , obs(1562) scm(pcm) icml(sigma) pcacm(pcacm)
```

Confirmatory Tetrad Analysis Results

Model 1			
rep	Chi-sq	df	p-val
1	22.1750	2	0.0000

The results of the CTA indicate that the four effect indicator model is not consistent with the data and thus suggests it is misspecified. Given the large sample size, however, the tetrad test has relatively high power to detect potentially minor model misspecifications (see Example 5 below for an illustration of how to calculate power).

Example 4: Bootstrapped Test Statistic

Our fourth example comes from a study of the measurement properties of blood pressure readings taken from a sample of respondents in Cebu, Philippines (Bauldry et al., 2015). The data contain three readings each for systolic and diastolic blood pressure at a given point in time. We are interested in decomposing the variance in the readings of blood pressure using multi-trait multi-method (MTMM) models. In this case, we treat systolic and diastolic blood pressure as two traits and the reading occasions as three methods (see Figure 3). For this model we constrain all of the factor loadings to equal one and the method factors to be uncorrelated with each other and the trait factors.

For this example we bootstrap the chi-square test statistic (T_1) with 1,000 bootstrap samples (see equation 5). The following Stata code specifies the MTMM model with Stata's `sem` command, saves the implied covariance matrix, and then uses the `tetrad_bootstrap` command to conduct a CTA.

```
. qui sem (SBP -> sbp13@1 sbp23@1 sbp33@1) ///
(DBP -> dbp13@1 dbp23@1 dbp33@1) ///
(R1 -> sbp13@1 dbp13@1) ///
(R2 -> sbp23@1 dbp23@1) ///
(R3 -> sbp33@1 dbp33@1), ///
cov(SBP*R1@0 DBP*R1@0 SBP*R2@0 DBP*R2@0 SBP*R3@0 DBP*R3@0) ///
cov(R1*R2@0 R1*R3@0 R2*R3@0)
. qui estat framework, fitted
. mat sigma1 = r(Sigma)
. tetrad_bootstrap sbp13 sbp23 sbp33 dbp13 dbp23 dbp33, icml(sigma1)
reps(1000)
```

Confirmatory Tetrad Analysis Results

Chi-sq	df	bootstrap
		p-value
0.9792	6	0.9540

We find a non-significant chi-square test statistic, T_1 , based on the bootstrapped chi-square test statistic, which is an indication that the model is consistent with the data. As a point of comparison, we also find a non-significant chi-square test statistic using the original tetrad test statistic ($T_0 = 5.0036$, $df = 5$, $p\text{-value} = .4154$).

Example 5: Power of Tetrad Test

Our final example illustrates the use of the `tetrad_matrix` command and two built-in Stata statistical functions to calculate the power of the tetrad test with respect to a specific alternative model. For this example we consider a two-factor model with three indicators per factor as the hypothesized model and an alternative model with a cross-loading from the first factor to the first indicator of the second factor. Following the steps outlined above, we first derive the implied covariance matrices for the hypothesized and the alternative models (labeled *sigma-h* and *sigma-a* respectively in the following code). In our derivation of the implied covariance matrix for the alternative model we treat the cross-loading as having a standardized factor loading of 0.5.

```
mat sigma_h = (1.000, 0.490, 0.490, 0.098, 0.098, 0.098 \ ///
0.490, 1.000, 0.490, 0.098, 0.098, 0.098 \ ///
0.490, 0.490, 1.000, 0.098, 0.098, 0.098 \ ///
0.098, 0.098, 0.098, 1.000, 0.490, 0.490 \ ///
0.098, 0.098, 0.098, 0.490, 1.000, 0.490 \ ///
0.098, 0.098, 0.098, 0.490, 0.490, 1.000)
mat colnames sigma_h = x1 x2 x3 x4 x5 x6
mat rownames sigma_h = x1 x2 x3 x4 x5 x6
mat sigma_a = (1.000, 0.490, 0.490, 0.448, 0.098, 0.098 \ ///
```

```

0.490, 1.000, 0.490, 0.448, 0.098, 0.098 \ ///
0.490, 0.490, 1.000, 0.448, 0.098, 0.098 \ ///
0.448, 0.448, 0.448, 1.000, 0.560, 0.560 \ ///
0.098, 0.098, 0.098, 0.560, 1.000, 0.490 \ ///
0.098, 0.098, 0.098, 0.560, 0.490, 1.000)
mat colnames sigma_a = x1 x2 x3 x4 x5 x6
mat rownames sigma_a = x1 x2 x3 x4 x5 x6

```

For the next step we use the `tetrad_matrix` command to enter the implied covariance matrix for the alternative model in place of the sample covariance matrix to generate (τ_a) and obtain the noncentrality parameter. In this step we set the sample size to 100. It is possible to calculate the power of the tetrad test across a range of sample sizes by simply changing the sample size with the `obs` option as we illustrate below.

```
. tetrad_matrix, scm(sigma_a) icml(sigma_h) obs(100)
```

Confirmatory Tetrad Analysis Results

Model 1			
rep	Chi-sq	df	p-val
1	14.3187	8	0.0738

The returned chi-square test statistic with the command specified using the implied covariance matrix for the alternative model in place of the sample covariance matrix provides the noncentrality parameter. The final step is to calculate the power of the tetrad test based on the *df*, a desired Type I probability, and the noncentrality parameter value. This step is accomplished with two calculations using Stata's built-in statistical functions. The first calculation finds the chi-square value associated with a given *df* and desired Type I probability. For this example we have 8 degrees of freedom and consider an alpha of 0.05. Using the inverse chi-square function (`invchi2`) we find a chi-square value of 15.507313 for *df* = 8 and $\alpha = 0.05$. The second calculation plugs the *df*, the noncentrality value returned from the `tetrad_matrix` command, and the chi-square value associated with the desired alpha into Stata's noncentral chi-square function (`nchi2`) and subtracts the result from 1 to obtain the power of the tetrad test to detect the difference between the hypothesized and the specific alternative model. The following code illustrates these calculations using Stata's display command.

```

. dis "chi-square value for alpha 0.05 = " invchi2(8, 1 - .05)
chi-square value for alpha 0.05 = 15.507313
. dis "power = " 1 - nchi2(8, 14.3187, 15.507313)
power = .77653918

```

Our calculations reveal that we have reasonable statistical power (0.78) at a sample size of 100 to detect the omission of a cross-loading with a standardized loading of .5 from the

hypothesized model. Similar calculations can be made across a range of sample sizes to determine the sample size needed for a desired level of power. Figure 4 illustrates the power of the tetrad test for this example in sample sizes ranging from 25 to 400. It is apparent that power reaches 0.8 at a sample size between 100 and 125 and that power is close to 1.0 at sample sizes 200 and higher.

Number of Variables and Computation Time

The number of vanishing tetrads and thus the computation time increases with the number of variables in a model. We conducted a simulation study to provide readers with an idea of the computation time involved for different sized models. The simulation study simply involves generating data based on a measurement model with one factor and an increasing number of indicators of the factor and logging the computation time for the original tetrad test statistic (Bollen, 1990; Bollen and Ting, 1993), T_0 , and the bootstrapped test statistic (Johnson and Bodner, 2007), T_1 . We set the sample size to 500 and ran 500 bootstrap replications for each T_1 . The simulation was run on a standard desktop computer with a 3.2 GHz Intel Core i5 processor and 16 GB of RAM.

Table 1 reports the computation time in seconds for models involving between 4 and 17 observed variables. We see that for T_0 models involving up to 11 observed variables take roughly up to a minute of computation time. The bootstrapped T_1 takes a bit longer, but is still less than 4 minutes. With more than 11 variables the computation times start to noticeably increase. At 14 observed variables the computational savings of T_1 balance the computational cost of bootstrapping-both T_0 and T_1 take about 21 minutes of computation time. From this point on, analysts are likely to be better off using T_1 , though it may depend on the specific computer configuration. At 17 observed variables T_0 took over 3.5 hours to run while T_1 took a little over two hours to run.

Conclusion

Confirmatory tetrad analysis provides a set of tools that researchers can add to their SEM toolbox. It permits test statistics for some underidentified models. CTA tests make possible the testing of parts of a model rather than only tests of the whole model. This permits the researcher to isolate the good fitting from the bad fitting parts of a model. CTA testing also allows some pairs of models that are not nested in the usual likelihood ratio sense to be compared in nested tetrads. The Stata commands we have introduced in this paper will enable analysts to take advantage of these features in the context of Stata, a general and widely used statistical package.

Acknowledgments

We thank Eli Lilly and Company and RTI Health Solutions for financial support for this project. Mark Boye's (Lilly) and Donald Stull's (RTI Health Solutions) encouragement and suggestions as well as David Braudt's research assistance are gratefully acknowledged.

References

- Bauldry S, Bollen KA, and Adair LS (2015). Evaluating measurement error in readings of blood pressure for adolescents and young adults. *Blood Pressure*, 24:96–102. [PubMed: 25548966]
- Bollen KA (1990). Outlier screening and a distribution-free test for vanishing tetrads. *Sociological Methods and Research*, 19:80–92.
- Bollen KA and Ting K-F (1993). Confirmatory tetrad analysis. *Sociological Methodology*, 23:147–175.
- Bollen KA and Ting K-F (1998). Bootstrapping a test statistic for vanishing tetrads. *Sociological Methods & Research*, 27:77–102.
- Bollen KA and Ting K-F (2000). A tetrad test for causal indicators. *Psychological Methods*, 5:3–22. [PubMed: 10937320]
- Goodnight JH (1979). A tutorial on the sweep operator. *American Statistician*, 33:149–158.
- Hipp JR, Bauer DJ, and Bollen KA (2005). Conducting tetrad tests of model fit and contrasts of tetrad-tested models: A new sas macro. *Structural Equation Modeling*, 12:76–93.
- Hipp JR and Bollen KA (2003). Model fit in structural equation models with censored, ordinal, and dichotomous variables: Testing vanishing tetrads. *Sociological Methodology*, 33:267–305.
- Johnson TR and Bodner TE (2007). A note on the use of the bootstrap tetrad tests for covariance structures. *Structural Equation Modeling*, 14:113–124.
- Kelley TL (1928). *Crossroads in the Mind of Man*. Stanford University Press, Stanford.
- Muthén LK and Muthén BO (1998–2012). *Mplus User's Guide*. Seventh Edition. Muthen & Muthen, Los Angeles, CA.
- Smith TW, Marsden P, Hout M, and Kim J (2015). General social surveys, 1972–2014 [machine-readable data file] Sponsored by National Science Foundation. –NORC ed.– Chicago: NORC at the University of Chicago [producer]; Storrs, CT: The Roper Center for Public Opinion Research, University of Connecticut [distributor].
- StataCorp (2015a). *Stata 14 Structural Equation Modeling Reference Manual*. StataCorp LP, College Station, TX.
- StataCorp (2015b). *Stata Statistical Software: Release 14*. StataCorp LP, College Station, TX.
- Ting K-F (1997). Confirmatory tetrad analysis in sas. *Structural Equation Modeling*, 2:163–171.

Figure 1:
Models for example 1.

Figure 2:
Model for example 2.

Figure 3:

Multi-trait multi-method model for example 4. $x_1 - x_3$ are three readings of systolic blood pressure, $y_1 - y_3$ are three reading of diastolic blood pressure. $R_1 - R_3$ are latent method factors for each reading occasion.

Figure 4:
Power of the tetrad test for the fourth example across a range of sample sizes.

Table 1:

Computation time (seconds) for CTA for models involving between 4 and 17 observed variables.

# vars	T_0	T_1	# vars	T_0	T_1
4	.02	2.29	11	76.69	210.04
5	.02	4.11	12	217.89	382.97
6	.07	8.19	13	551.56	695.91
7	.34	16.10	14	1302.14	1312.57
8	1.58	31.56	15	2899.44	2459.61
9	6.52	60.14	16	6439.51	4379.67
10	24.69	114.15	17	13119.60	7749.70

Notes: Computation based on a one-factor measurement model with 500 cases. T_1 calculated based on 500 bootstrap samples. Simulation run on a standard desktop computer with a 3.2 GHz Intel Core i5 processor and 16 GB of RAM.