

Testing Nonnested Structural Equation Models

Edgar C. Merkle and Dongjun You
University of Missouri

Kristopher J. Preacher
Vanderbilt University

In this article, we apply [Vuong's \(1989\)](#) likelihood ratio tests of nonnested models to the comparison of nonnested structural equation models (SEMs). Similar tests have been previously applied in SEM contexts (especially to mixture models), though the nonstandard output required to conduct the tests has limited their use and study. We review the theory underlying the tests and show how they can be used to construct interval estimates for differences in nonnested information criteria. Through both simulation and application, we then study the tests' performance in nonmixture SEMs and describe their general implementation via free R packages. The tests offer researchers a useful tool for nonnested SEM comparison, with barriers to test implementation now removed.

Keywords: model selection, structural equation models, nonnested models, likelihood ratio test, Vuong test

Researchers frequently rely on model comparisons to test competing theories. This is especially true when structural equation models (SEMs) are used, because the models are often able to accommodate a large variety of theories. When competing theories can be translated into nested SEMs, the comparison is relatively easy: One can compute the likelihood ratio statistic using the results of the fitted models (e.g., [Steiger, Shapiro, & Browne, 1985](#)). The test associated with this likelihood ratio statistic yields one of two conclusions: The two models fit equally well, so the simpler model is to be preferred, or the more complex model fits better, so it is to be preferred. As is well known, however, the likelihood ratio statistic does not immediately extend to situations in which models are nonnested.

In the nonnested case, researchers typically rely on information criteria for model comparison, including the Akaike information criterion (AIC; [Akaike, 1974](#)) and the Bayesian information criterion (BIC; [Schwarz, 1978](#)). One computes an AIC or BIC for the two models, then selects the model with the lower criterion as better. Thus, the applied conclusion differs slightly from the likelihood ratio test (LRT): We conclude from the information criteria that one or the other model is better, whereas we conclude from the LRT either that the complex model is better or that there is insufficient evidence to differentiate between model fits.

Although information criteria can be applied to nonnested models, the popular “select the model with the lower decision criterion” approach can be problematic. In particular, [Preacher and Merkle \(2012\)](#) showed that the BIC exhibits large variability at the sample sizes typically used in SEM contexts. Thus, the model that is preferred for a given sample often will not be preferred in new samples. Preacher and Merkle studied a series of nonparametric bootstrap procedures to estimate sampling variability in the BIC, but no procedure succeeded in fully characterizing this variability.

A problem with the “select the model with the lower criterion” approach involves the fact that one can never conclude that the models fit equally. There may often be situations in which the models exhibit close values of the information criteria, yet one of the models is still selected as better. To handle this issue, [Pornprasertmanit, Wu, and Little \(2013\)](#) developed a parametric bootstrap method that allows one to conclude that the two models are equally good (in addition to concluding that one or the other model fits better). Their results indicated that the procedure is promising, though it is also computationally expensive: One must draw a large number of bootstrap samples from each of the two fitted models, then refit each model to each bootstrap sample.

In this article, we study formal tests of nonnested models that allow us to conclude that one model fits better than the other, that the two models exhibit equal fit, or that the two models are indistinguishable in the population of interest. The tests are based on the theory of [Vuong \(1989\)](#), and one of the tests is popularly applied to the comparison of mixture models with different numbers of components (including count regression models and factor mixture models; [Greene, 1994](#); [Lo, Mendell, & Rubin, 2001](#); [Nylund, Muthén, & Asparouhov, 2007](#)). Although some authors have recently described problems with mixture model applications ([Jeffries, 2003](#); [Wilson, 2015](#)), the tests have the potential to be very useful in general SEM contexts. This is because nonnested SEMs are commonly observed throughout psychology (e.g., [Frese, Garst, & Fay, 2007](#); [Kim, Zane, & Hong, 2002](#); [Sani & Todman, 2002](#)).

This article was published Online First August 3, 2015.

Edgar C. Merkle and Dongjun You, Department of Psychological Sciences, University of Missouri; Kristopher J. Preacher, Department of Psychology and Human Development, Vanderbilt University.

This work was supported by National Science Foundation Grant SES-1061334. Portions of the work were presented at the 2014 Meeting of the Psychometric Society in Madison, Wisconsin, and at the 2015 International Workshop on Psychometric Computing in Amsterdam, The Netherlands. We thank Achim Zeileis for comments that improved the article. All remaining errors are solely the responsibility of the authors.

Correspondence concerning this article should be addressed to Edgar C. Merkle, Department of Psychological Sciences, University of Missouri, Columbia, MO 65211. E-mail: merkle@missouri.edu

Levy and Hancock (2007, 2011) studied the application of Vuong's (1989) theory to SEMs, describing relevant background and proposing steps by which researchers can carry out tests of nonnested models. Levy and Hancock (2007, 2011) bypassed an important step of the theory because of the nonstandard model output required, instead requiring researchers to algebraically examine the candidate models and, potentially, to carry out LRTs between each candidate model and a constrained version of the model. This procedure can accomplish the desired goal, but it also requires a considerable amount of analytic and computational work on the part of the user. We study the tests as Vuong originally proposed them, using the nonstandard model output that is required. This study is aided by our general implementation of the tests, available via the R package *nonnest2* (Merkle & You, 2014).

In the following pages, we first describe the relevant theoretical results from Vuong (1989). We also show how the theory can be used to obtain confidence intervals for differences in BIC (and other information criteria) associated with nonnested models. Next, we apply the tests to data on teacher burnout that were originally examined by Byrne (1994). Then, we describe the results of three simulations illustrating test properties in the context of structural equation modeling. Finally, we discuss recommendations, extensions, and practical issues.

Theoretical Background

In this section, we provide an overview of the theory underlying the test statistics. The overview is largely based on Vuong (1989), and the reader is referred to that article for further detail. For alternative overviews of the theory, see Golden (2000) and Levy and Hancock (2007). The theory is applicable to many models estimated via maximum likelihood (ML), though we focus on SEMs here.

We consider situations in which two candidate SEMs, M_A and M_B , are to be compared using a dataset X with n cases and p manifest variables. M_A may be represented by the equations

$$\mathbf{x}_i = \boldsymbol{\nu}_A + \boldsymbol{\Lambda}_A \boldsymbol{\eta}_{A,i} + \boldsymbol{\epsilon}_{A,i} \quad (1)$$

$$\boldsymbol{\eta}_{A,i} = \boldsymbol{\alpha}_A + \mathbf{B}_A \boldsymbol{\eta}_{A,i} + \boldsymbol{\zeta}_{A,i}, \quad (2)$$

where $\boldsymbol{\eta}_{A,i}$ is a vector containing the latent variables in M_A ; $\boldsymbol{\epsilon}_{A,i}$ and $\boldsymbol{\zeta}_{A,i}$ are zero-centered residual vectors, independent across values of i ; $\boldsymbol{\Lambda}_A$ is a matrix of factor loadings; and \mathbf{B}_A contains parameters that reflect directed paths between latent variables. The second model, M_B , is defined similarly, and we restrict ourselves to situations in which the residuals and latent variables are assumed to follow multivariate normal distributions (though this distributional assumption is not required to use the test statistics [see the General Discussion]).

The M_A parameter vector, $\boldsymbol{\theta}_A$, includes the parameters in $\boldsymbol{\nu}_A$, $\boldsymbol{\Lambda}_A$, $\boldsymbol{\alpha}_A$, and \mathbf{B}_A along with variance and covariance parameters related to the latent variables and residuals. These parameters imply a marginal multivariate normal distribution for \mathbf{x}_i with a specific mean vector (typically $\boldsymbol{\nu}_A$) and covariance matrix (see, e.g., Browne & Arminger, 1995) that allows us to estimate the model via ML. In particular, we choose $\hat{\boldsymbol{\theta}}_A$ to maximize the log-likelihood [$\ell(\cdot)$]:

$$\ell(\boldsymbol{\theta}_A; \mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_{i=1}^n \ell(\boldsymbol{\theta}_A; \mathbf{x}_i) = \sum_{i=1}^n \log f_A(\mathbf{x}_i; \boldsymbol{\theta}_A),$$

where $f_A(\mathbf{x}_i; \boldsymbol{\theta}_A)$ is the probability density function of the multivariate normal distribution, with the mean vector and covariance matrix implied by M_A and its parameter vector $\boldsymbol{\theta}_A$. Similarly, the parameters $\hat{\boldsymbol{\theta}}_B$ are chosen to maximize the log-likelihood $\ell(\boldsymbol{\theta}_B; \mathbf{x}_1, \dots, \mathbf{x}_n)$.

Instead of defining the ML estimates as in the foregoing, we could equivalently define them via the gradients

$$s(\hat{\boldsymbol{\theta}}_A; \mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_{i=1}^n s(\hat{\boldsymbol{\theta}}_A; \mathbf{x}_i) = 0$$

and

$$s(\hat{\boldsymbol{\theta}}_B; \mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_{i=1}^n s(\hat{\boldsymbol{\theta}}_B; \mathbf{x}_i) = 0,$$

where the gradients sum scores across individuals (where *scores* are defined as the casewise contributions to a gradient). This simply means that we choose parameters such that the gradient (first derivative of the likelihood function) equals zero. Assuming that M_A has k free parameters, the associated score function may be explicitly defined as

$$s(\boldsymbol{\theta}_A; \mathbf{x}_i) = \left(\frac{\partial \ell(\boldsymbol{\theta}_A; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_{A,1}}, \dots, \frac{\partial \ell(\boldsymbol{\theta}_A; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_{A,k}} \right)', \quad (3)$$

with the score function for M_B defined similarly (where the number of free parameters for M_B is q instead of k). We also define M_A 's expected information matrix as

$$\mathbf{I}(\boldsymbol{\theta}_A) = -E \frac{\partial^2 \ell(\boldsymbol{\theta}_A; \mathbf{x}_1, \dots, \mathbf{x}_n)}{\partial \boldsymbol{\theta}_A \partial \boldsymbol{\theta}_A'}, \quad (4)$$

where, again, the information matrix of M_B is defined similarly.

The statistics described here can be used in general model comparison situations, in which one is interested in which of two candidate models (M_A and M_B) is closest to the data-generating model in Kullback–Leibler (K-L) distance (Kullback & Leibler, 1951). Letting $g(\cdot)$ be the density of the data-generating model (which is generally unknown), the distances for M_A and M_B can be denoted $\text{KL}_{A,g}$ and $\text{KL}_{B,g}$, respectively. The distance $\text{KL}_{A,g}$ may be explicitly written as

$$\begin{aligned} \text{KL}_{A,g} &= \int \log \left(\frac{g(\mathbf{x})}{f_A(\mathbf{x}; \boldsymbol{\theta}_A^*)} \right) g(\mathbf{x}) d\mathbf{x} \\ &= E[\log(g(\mathbf{x}))] - E[\log(f_A(\mathbf{x}; \boldsymbol{\theta}_A^*))], \end{aligned} \quad (5)$$

where $\boldsymbol{\theta}_A^*$ is the M_A parameter vector that minimizes this distance (also called the *pseudo-true* parameter vector), and the expected values are taken with respect to $g(\cdot)$. The pseudo-true designation arises from the fact that $\boldsymbol{\theta}_A^*$ usually does not reflect the true parameter vector (because the candidate model M_A is usually incorrect). However, the parameter vector is pseudo-true in the sense that it allows M_A to most closely approximate the truth.

Relationships Between Models

Relationships between pairs of candidate models may be characterized in multiple manners. Familiarly, *nested* models are those for which the parameter space associated with one model is a

subset of the parameter space associated with the other model; every set of parameters from the less complex model can be translated into an equivalent set of parameters from the more complex model. Similarly, *nonnested* models are those whose parameter spaces each include some unique points. Aside from these two broad, familiar classifications, however, we may define other relationships between models. These include the concepts of equivalence, overlappingness, and strict nonnesting. The latter two concepts refer to specific types of nonnested models. All three are described in the following.

Many SEM researchers are familiar with the concept of model *equivalence* (e.g., Hershberger & Marcoulides, 2013; Lee & Hershberger, 1990; MacCallum, Wegener, Uchino, & Fabrigar, 1993): Two seemingly different SEMs yield exactly the same implied moments (mean vectors and covariance matrices) and fit statistics when fit to any dataset. SEM researchers are less familiar with the concept of *overlapping* models: These are models that yield exactly the same implied moments and fit statistics for some populations but not for others. Further, even if the models yield identical moments in the population of interest, they will be exactly identical only in the full population. Fits to sample data will generally yield different moments and fit statistics.

An example of overlapping models is displayed in Figure 1. Both M_A and M_B are two-factor models, where M_A has a free path from η_1 to X_4 , and M_B has a free path from η_2 to X_3 . These models are overlapping: Their predictions and fit statistics will usually differ, but they will be the same in populations in which the parameters labeled A and B both equal zero. Even if these two paths equal zero in the population, they generally will not equal zero when the models are fit to sample data. Therefore, the models will not have the same implied moments and fit statistics when fit to sample data. This means that, given sample data, we must test whether the models are distinguishable in the population of interest. Note that *overlappingness* is a general relationship between two models regardless of the focal population, whereas *distinguishability* is a specific relationship between models in the context of a single focal population.




Figure 1. Path diagram reflecting the models used in the simulation. Model A is the data-generating model, with the loading labeled A varying across conditions.

Overlapping models are one subclassification of nonnested models. The other subclassification is *strictly nonnested* models; these are models whose parameter spaces do not overlap at all. In other words, strictly nonnested models never yield the same implied moments and fit statistics for any population. Strictly nonnested models may have different functional forms (say, an exponential growth model vs. a logistic growth model) or may make different distributional assumptions.

It is often difficult to know whether candidate SEMs are overlapping. Consider the path models in Figure 2, which reflect four potential hypotheses about the relationships between nine observed variables. These models are obviously nonnested, but are they overlapping or strictly nonnested? Assuming that all models use the same data distribution (typically multivariate normal), the models will be indistinguishable in populations in which all observed variables are independent of one another. Therefore, SEMs that use the same form of data distribution will typically be overlapping. In other situations, it may be difficult to tell whether the candidate models are overlapping. Thus, it is important to test for distinguishability when doing model comparisons: If the observed data imply that the models are indistinguishable in the population of interest, then there is no point in further model comparison. This is especially relevant to the use of information criteria for nonnested model comparison, in which one is guaranteed to select a candidate model as better (at least, using the standard decision criteria). In addition, as shown later, when a pair of models is nonnested, the limiting distribution of the likelihood ratio statistic depends on whether the models are distinguishable.

The foregoing discussion suggests a sequence of tests for comparing two models. Assuming that the models are not equivalent to one another (regardless of population), we must establish that the models are distinguishable in the population of interest. Assuming that the models are distinguishable, we can then compare the models' fits and potentially select one as better. In the next section, we describe test statistics that can be used in this sequence.

Test Statistics

Vuong's (1989) tests of distinguishability and of model fit use the terms $\ell(\hat{\theta}_A; \mathbf{x}_i)$ and $\ell(\hat{\theta}_B; \mathbf{x}_i)$ for $i = 1, \dots, n$, which are the casewise likelihoods evaluated at the ML estimates. For the purpose of structural equation modeling, we focus on two separate statistics that Vuong proposed. One statistic tests whether models are distinguishable, and the other tests the fit of nonnested, distinguishable models. These tests proceed sequentially in situations in which we are unsure whether the two models are distinguishable in the population of interest. If we know in advance that two models are not overlapping, we can proceed directly to the second test.

To gain an intuitive feel for the tests, imagine that we fit both M_A and M_B to a dataset and then obtain $\ell(\hat{\theta}_A; \mathbf{x}_i)$ and $\ell(\hat{\theta}_B; \mathbf{x}_i)$ for $i = 1, \dots, n$. To test whether the models are distinguishable, we can calculate a likelihood ratio for each case i and examine the variability in these n ratios. If the models are indistinguishable, then these ratios should be similar for all individuals such that their variability is close to zero. If the models are distinguishable, then the variability in the casewise likelihood ratios characterizes general sampling variability in the likelihood ratio between M_A and M_B , allowing for a formal model comparison test that does not require the models to be nested.

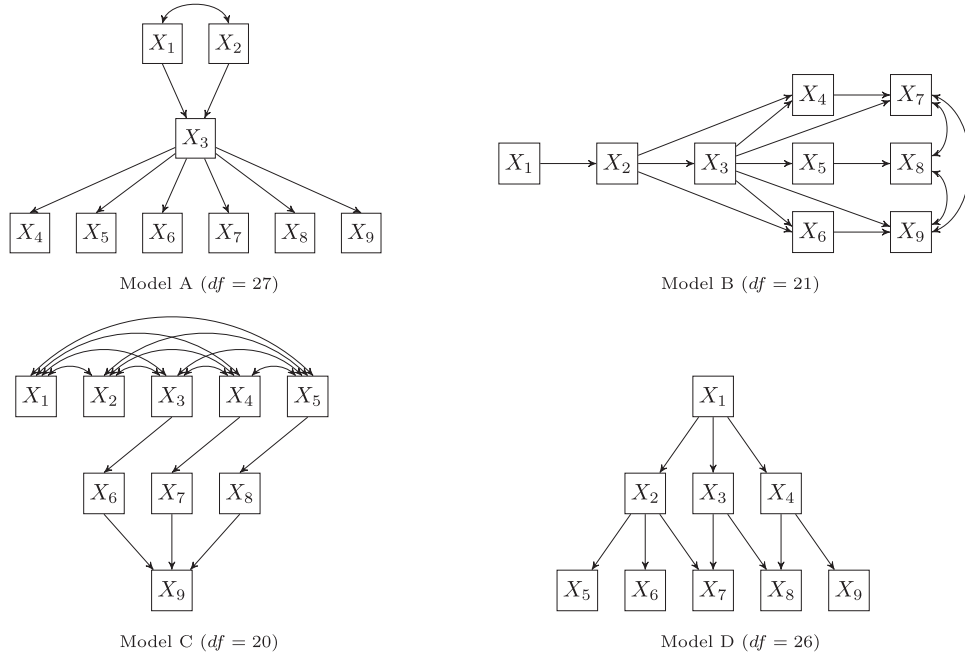


Figure 2. Path diagrams reflecting the models used in Simulation 2.

To formalize the ideas in the previous paragraph, we characterize the population variance in individual likelihood ratios of M_A versus M_B as

$$\omega_s^2 = \text{var} \left[\log \frac{f_A(\mathbf{x}_i; \boldsymbol{\theta}_A^*)}{f_B(\mathbf{x}_i; \boldsymbol{\theta}_B^*)} \right]$$

where this variance is taken with respect to the pseudotrue parameter vectors introduced in Equation 5. This means that we make no assumption that either candidate model is the true model. Using this equation, hypotheses for a test of model distinguishability may then be written as

$$H_0: \omega_s^2 = 0 \quad (7)$$

and

$$H_1: \omega_s^2 > 0, \quad (8)$$

with a sample estimate of ω_s^2 being

$$\hat{\omega}_s^2 = \frac{1}{n} \sum_{i=1}^n \left[\log \frac{f_A(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_A)}{f_B(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_B)} \right]^2 - \left[\frac{1}{n} \sum_{i=1}^n \log \frac{f_A(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_A)}{f_B(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_B)} \right]^2. \quad (9)$$

Division by $n - 1$ instead of n is also possible to reduce bias in the estimate, though this will have little impact at the sample sizes typically observed in SEM applications.

Vuong (1989) showed that, under Equation 7 and mild regularity conditions (ensuring that second derivatives of the likelihood function exist, observations are independent and identically distributed [IID], and the ML estimates are unique and not on the boundary), $n\hat{\omega}_s^2$ is asymptotically distributed as a particular weighted sum of chi-square distributions. Weighted sums of chi-square distributions arise when we sum the squares of normally distributed variables; the normally distributed variables involved

in the test statistics here are the ML estimates $\hat{\boldsymbol{\theta}}_A$ and $\hat{\boldsymbol{\theta}}_B$. The weights involved in this sum are obtained via the squared eigenvalues of a matrix \mathbf{W} that arises from the candidate models' scores (Equation 3) and information matrices (Equation 4); see the Appendix for details. This result immediately allows us to test Equation 7 using results from the two fitted models. If the null hypothesis is not rejected, we conclude that the models cannot be distinguished in the population of interest. In the case in which models are nested, this conclusion would lead us to prefer the model with fewer degrees of freedom. In the case in which models are not nested, the two candidate models may have the same degrees of freedom. Thus, depending on the specific models being compared, we might not prefer either model.

The software requirements for carrying out the test of Equation 7 are somewhat nonstandard. For each candidate model, we need to obtain the scores from Equation 3 and the information matrix from Equation 4. We then need to arrange this output in matrices, do some multiplications to obtain a new matrix, and obtain eigenvalues of this new matrix. Further, we need the ability to evaluate quantiles of weighted sum of chi-square distributions. The difficulty in obtaining these results led Levy and Hancock (2007) to bypass the test of model distinguishability and instead conduct an algebraic model comparison that provides evidence about whether the models are distinguishable. Whereas this algebraic comparison is reasonable, it does not always indicate whether models are distinguishable. Further, it is more complicated for the applied researcher who wishes to use these tests (assuming that an implementation of the statistical test is available).

Assuming that the null hypothesis from Equation 7 is rejected (i.e., that the models are distinguishable), we may compare the models via a nonnested LRT. We can write the hypotheses associated with this test as

$$H_0: E[\ell(\hat{\boldsymbol{\theta}}_A; \mathbf{x}_i)] = E[\ell(\hat{\boldsymbol{\theta}}_B; \mathbf{x}_i)], \quad (10)$$

$$H_{1A}: E[\ell(\hat{\boldsymbol{\theta}}_A; \mathbf{x}_i)] > E[\ell(\hat{\boldsymbol{\theta}}_B; \mathbf{x}_i)], \quad (11)$$

$$H_{1B}: E[\ell(\hat{\boldsymbol{\theta}}_A; \mathbf{x}_i)] < E[\ell(\hat{\boldsymbol{\theta}}_B; \mathbf{x}_i)], \quad (12)$$

where the expectations arise from the K-L distance in Equation 6 (note that we do not need to consider the expected value of $\log[g\{\mathbf{x}\}]$ because it is constant across candidate models). The hypothesis H_0 in the foregoing states that the K-L distance between M_A and the truth equals the K-L distance between M_B and the truth; this is like stating that the two models have equal population discrepancies. The test is written to be directional, so one chooses either H_{1A} or H_{1B} prior to carrying out the test. In practice, however, a two-tailed test is often carried out, with a single model being preferred in the situation in which H_0 is rejected. This follows the framework of Jones and Tukey (2001), whereby there are three possible conclusions available to researchers: M_A is closer to the truth (in K-L distance) than M_B , M_B is closer to the truth than M_A , or there is insufficient evidence to conclude that either model is closer to the truth than the other.

For nonnested, distinguishable models, Vuong (1989) showed that

$$\text{LR}_{AB} = n^{-1/2} \sum_{i=1}^n \log \frac{f_A(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_A)}{f_B(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_B)} \xrightarrow{d} N(0, \omega_*^2) \quad (13)$$

under Equation 10 and the regularity conditions noted earlier. Thus, we obtain critical values and p values by comparing the nonnested LRT statistic with the standard normal distribution. Assuming that the desired Type I error rate is α_2 for this test and that the desired Type I error rate is α_1 for the test of Equation 7, Vuong showed that the sequence of tests has an overall Type I error rate that is bounded from above by $\max(\alpha_1, \alpha_2)$. In practice, it is customary to set $\alpha_1 = \alpha_2$.

Assuming that the null hypothesis from Equation 7 is not rejected (i.e., that the models are indistinguishable), Vuong (1989) showed that $2n^{1/2}\text{LR}_{AB}$ follows a weighted sum of chi-square distributions, where the weights are obtained from the *unsquared* eigenvalues of the same matrix \mathbf{W} that arose in the test of Equation 7. This result is not typically used in the case of nonnested models: We first need the test of ω_*^2 to determine the limiting distribution that we should use for the likelihood ratio (either the normal distribution from Equation 13 or the weighted sum of chi-square distribution described here). If the test of ω_*^2 indicates indistinguishable models, however, there is no point in further testing the models via $2n^{1/2}\text{LR}_{AB}$. If the test of ω_*^2 indicates distinguishable models, then we rely on the limiting distribution from Equation 13 instead of the weighted sum of chi-squares. The result described in this paragraph can be used to compare nested models, however, and we return to this topic in the next subsection.

The test statistics just described are implemented for general multivariate normal SEMs (and other models) in the free R package *nonnest2* (Merkle & You, 2014). Models are first estimated via *lavaan* (Rosseel, 2012), then *nonnest2* computes the test statistics on the basis of the fitted models' output. In the following sections, we describe ways in which these ideas can be extended to test nested models and to test information criteria.

Testing Nested Models

In the situation in which M_B is nested within M_A , the likelihood ratio and the variance statistic $n\omega_*^2$ can each be used to construct a unique test of M_A versus M_B . For nested models, the null hypotheses from Equations 7 and 10 can be shown to be the same as the traditional null hypothesis:

$$H_0: \boldsymbol{\theta}_A \in h(\boldsymbol{\theta}_B), \quad (14)$$

and

$$H_1: \boldsymbol{\theta}_A \notin h(\boldsymbol{\theta}_B), \quad (15)$$

where $h()$ is a function translating the M_B parameter vector to an equivalent M_A parameter vector. In the general case, where M_A is not assumed to be correctly specified, the statistics $n\omega_*^2$ and $2n^{1/2}\text{LR}_{AB}$ (see Equations 9 and 13, respectively) both strongly converge (i.e., almost surely; Casella & Berger, 2002) to weighted sums of chi-square distributions under the null hypothesis from Equation 14. The specific weights differ between the two statistics; the weights associated with $n\omega_*^2$ are the squared eigenvalues of the \mathbf{W} matrix that is defined in the Appendix, whereas the weights associated with $2n^{1/2}\text{LR}_{AB}$ are the unsquared eigenvalues of the same \mathbf{W} matrix. This result differs from the usual multivariate normal SEM derivations (e.g., Amemiya & Anderson, 1990; Steiger et al., 1985), which use either an assumption that M_A is correctly specified or that the population parameters drift toward a point that is contained in M_A 's parameter space (for further discussion of this point and fit assessment of single models, see Chun & Shapiro, 2009). Under the assumption that M_A is correctly specified, the statistics $n\omega_*^2$ and $2n^{1/2}\text{LR}_{AB}$ weakly converge (i.e., in distribution) to the usual $\chi_{df=q-k}^2$ distribution under Equation 14. Hence, the framework here provides a more general characterization of the nested LRT than do traditional derivations.

Testing Information Criteria

Model selection with AIC or BIC (i.e., selecting the model with the lower criterion) involves adjustment of the likelihood ratio by a constant term that penalizes the two models for complexity. Thus, as Vuong (1989) originally described, the foregoing results can be extended to test differences in AIC or BIC. To show this formally, we focus on BIC and write the BIC difference between two models as

$$\text{BIC}_A - \text{BIC}_B = (k \log n - q \log n) - 2 \sum_{i=1}^n \log \frac{f_A(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_A)}{f_B(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_B)},$$

where k and q are the number of free parameters for M_A and M_B , respectively. This shows that we are simply taking a linear transformation of the usual likelihood ratio, so the test of Equation 7 and the result from Equation 13 apply here. In particular, if models are indistinguishable, then one could select the model that BIC penalizes less (i.e., the model with fewer parameters). If models are distinguishable, then we may formulate a hypothesis that $\text{BIC}_A = \text{BIC}_B$. Under this hypothesis, the result from Equation 13 can be used to show that

$$n^{-1/2} \left[((k - q) \log n) - 2 \sum_{i=1}^n \log \frac{f_A(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_A)}{f_B(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_B)} \right] \xrightarrow{d} N(0, 4\omega_*^2).$$

This result can be used to obtain an adjusted test statistic that

accounts for the models' relative complexity. Alternatively, we prefer to use this result to construct a $100 \times (1 - \alpha)\%$ confidence interval associated with the BIC difference. This is obtained via

$$(\text{BIC}_A - \text{BIC}_B) \pm z_{1-\alpha/2} \sqrt{4n\omega_s^2}, \quad (16)$$

where $z_{1-\alpha/2}$ is the variate at which the cumulative distribution function of the standard normal distribution equals $1 - \alpha/2$. To coincide with the tests described previously, α should be the same as the α level used for the test of distinguishability.

To our knowledge, this is the first analytic confidence interval for a nonnested difference in information criteria that has been presented in the SEM literature. This confidence interval is simpler to calculate than bootstrap intervals (for a discussion of bootstrap procedures, see, e.g., Preacher & Merkle, 2012), and, as shown later, its coverage is often comparable. The bootstrap intervals may still be advantageous if regularity conditions are violated.

Relation to the Nesting and Equivalence Test

Bentler and Satorra (2010) describe a nesting and equivalence test (NET) that assesses whether two models provide exactly the same fit to sample data, relying on the fact that equivalent models can perfectly fit one another's implied mean vectors and covariance matrices. As previously described, this differs from the indistinguishability characteristic that is relevant to the tests in this article. Indistinguishable models provide exactly the same fit in the population but not necessarily to sample data.

The NET procedure is convenient and computationally simple, and it is generally suited to examining whether two models are globally nested or equivalent across large sets of covariance matrices (though, for some pathological cases, see Bentler & Satorra, 2010). It cannot, however, inform us about whether two models are distinguishable on the basis of the sample data. In other words, equivalent models are indistinguishable, but indistinguishable models are not necessarily equivalent. Thus, the two methods are complementary: We can use the NET to determine whether two models can possibly be distinguished from each another, and we can use the test of Equation 7 to determine whether two models can be distinguished in the population of interest. In the following sections, we study the tests' applications to structural equation modeling using both simulation and real data.

Application: Teacher Burnout

Background

Byrne (1994) tested the impact of organizational and personality variables on three dimensions of teacher burnout. The application here is limited to the sample of elementary teachers only ($n = 599$ [intermediate and secondary school teachers were also observed]) and uses models that are related to those presented in Chapter 6 of Byrne (2009).

Method

The candidate models that we consider are illustrated in Figure 3, which shows only the latent variables included in the model (and

not the indicators of the latent variables or variance parameters). Each pair of models is nonnested, so the BIC (or some other information criterion) would typically be used to select a model from the set. Alternatively, we can use the statistics described in this article to study the models' distinguishability and fit in greater detail.

To expand on the model comparison procedure, we first use the NET (Bentler & Satorra, 2010) to determine whether pairs of models are equivalent to one another. For models that are not equivalent, we then use Vuong's (1989) distinguishability test to make inferences about whether each pair of models is distinguishable on the basis of the focal population of elementary teachers. Finally, we use Vuong's nonnested LRT to study the candidate models' relative fit. If desired, the latter test can be accompanied by BIC statistics and interval estimates of BIC differences.

Results

To mimic a traditional comparison of nonnested models, we first examine the three candidate models' BICs. We find that the BIC decreases as we move from M_1 to M_3 ($\text{BIC}_1 = 40,040.7$, $\text{BIC}_2 = 39,994.1$, $\text{BIC}_3 = 39,978.9$), which would lead us to prefer M_3 . In addition, the BIC difference between M_3 and its closest competitor, M_2 , is about 15. Using the Raftery (1995) grades of evidence for BIC differences, we would conclude that there is very strong evidence for M_3 over the other models.

We now undertake a larger model comparison via the NET and the Vuong (1989) tests, comparing each candidate model with M_3 . In applying the NET procedure, we find that neither M_1 nor M_2 is equivalent to M_3 ($\hat{F}_{13} = 57.9$ and $\hat{F}_{23} = 13.4$, respectively). Next, we test whether each pair of models is distinguishable, using the test of Equation 7. We find that both M_1 and M_2 are distinguishable from M_3 ($\hat{\omega}_{13}^2 = 0.12$ [$p < .01$] and $\hat{\omega}_{23}^2 = 0.05$ [$p < .01$], respectively). This yields evidence that the models can be distin-




Figure 3. Path diagram of candidate Models 1, 2, and 3 (M_1 , M_2 , and M_3 , respectively), which are related to those originally specified by Byrne (2009). Candidate model M_1 is displayed, with dotted lines reflecting parameters that are added (+) or removed (-) in models M_2 and M_3 . Only latent variables are displayed, and covariances between exogenous variables are always estimated. For specific detail on the variables used in this model, see Byrne (2009).

guished from one another on the basis of the population of interest, so it makes sense to further compare the models' fits.

To compare model fits, we use [Vuong's \(1989\)](#) nonnested LRT of [Equation 10](#). We find that M_3 fits better than M_1 ($z_{13} = 2.86$, $p = .002$), reinforcing the BIC results described earlier. The test of M_2 versus M_3 differs from the BIC results, however. Here, we do not reject the hypothesis that the two models' fits are equal in the population ($z_{23} = 0.84$, $p = .20$). Further, the 95% confidence interval associated with the BIC difference of M_2 versus M_3 is $[-5.3, 35.7]$. This overlaps with zero, providing evidence that we cannot prefer either model after adjusting for differences in model complexity.

These results imply that, despite having very strong evidence for M_3 via traditional BIC comparisons, the fits of M_2 and M_3 are sufficiently close that we cannot prefer either model over the other. [Vuong's \(1989\)](#) methodology provided us with the ability to draw these conclusions in a straightforward manner that can be generally applied to structural equation modeling. A reviewer noted that, for the example here, similar conclusions could be drawn by specifying a larger model within which all candidate models are nested. This larger model would include free parameters associated with every path (solid or dotted) in [Figure 3](#), and we could compare each candidate model with this larger model via traditional LRTs. This strategy is useful, though it is not as general as the Vuong methodology. For example, the larger model strategy could not be used if we compared models with differing distributional assumptions. Further, estimation of the larger model may be difficult or impossible in some situations (resulting in, e.g., the estimation algorithm failing to converge to the ML estimates).

In the following sections, we further study the tests' abilities via simulations.

Simulation 1: Overlapping Models

In Simulation 1, our data-generating model is sometimes a special case of both candidate models, so the models are sometimes indistinguishable. We study the test of [Equation 7](#)'s ability to pick up the indistinguishable models, and we also study the non-nested LRT's (of [Equation 10](#)) ability to compare models that are judged to be distinguishable. Finally, we compare the results obtained with these two novel tests with the use of (a) the [Bentler and Satorra \(2010\)](#) NET and (b) the BIC. Although the true model is included in the set of estimated models for simplicity, this need not be assumed for [Vuong's \(1989\)](#) tests to be valid.

Method

The two candidate models are displayed in [Figure 1](#): Both are two-factor models, and they differ in which loadings are estimated. The data-generating model, M_A , has an extra loading from the first factor to the fourth indicator (labeled *A* in the figure). The second model, M_B , instead has an extra loading from the second factor to the third indicator (labeled *B* in the figure).

To study the tests described in this article, we set the data-generating model's parameter values equal to the parameter estimates obtained from a two-factor model fit to the [Holzinger and Swineford \(1939\)](#) data (using the scales that load on the textual and speed factors). In addition, we manipulated the magnitude of the *A* loading during data generation: This loading could take values of

$d = 0, 0.1, \dots, 0.5$. In the condition in which $d = 0$, the data-generating model is a special case of both candidate models. In other conditions, M_A is preferable to M_B . However, when d is close to zero, the tests may still indicate that the models are indistinguishable from one another.

Simulation conditions were defined by $d = 0, 0.1, \dots, 0.5$ and by $n = 200, 500, \text{ and } 1,000$. In each condition, we generated 3,000 datasets and fit both M_A and M_B to the data. We then computed five statistics: the NET ([Bentler & Satorra, 2010](#)), the distinguishability test of [Equation 7](#), the LRT of [Equation 10](#), and each model's BIC. The BIC is not required here (the [Vuong, 1989](#), tests can be used in place of the BIC), but we included it for comparison. Using each statistic, we recorded whether M_A was favored for each dataset. To be specific, we counted each statistic as favoring M_A if (a) the NET implied that models were not equivalent, (b) the test of [Equation 7](#) implied that models were distinguishable at $\alpha = .05$, (c) the test of [Equation 10](#) was significant in the direction of M_A at a one-tailed $\alpha = .05$, and (d) the M_A BIC was lower than the M_B BIC. Of course, the fact that models are distinguishable does not necessarily imply that M_A should be preferred to M_B . However, these definitions allowed us to put the tests on a common scale for the purpose of displaying results.

In addition to these statistics, we computed two types of 90% confidence intervals of BIC differences (which are actually chi-square differences here, because the models have the same number of parameters). The first type of interval was based on the result from [Equation 16](#), whereas the second type of interval was based on the nonparametric bootstrap (based on 1,000 bootstrap samples per replication). Summaries of interest included interval coverage, mean interval width, and interval variability. The latter statistic is defined as the pooled standard deviation of the lower and upper confidence limits; for a given sample size and interval type, the statistic is computed via

$$s_{\text{int}} = \sqrt{\frac{(n_{\text{rep}} - 1) \times (s_L^2 + s_U^2)}{2 \times n_{\text{rep}} - 2}}, \quad (17)$$

where s_L^2 is the variance of the lower limit, s_U^2 is the variance of the upper limit, and n_{rep} is the number of replications within one simulation condition (3,000 for this simulation).

Results

Overall simulation results are displayed in [Figure 4](#). It can be seen that the NET procedure almost never declares the two candidate models to be equivalent, even in the condition in which $d = 0$. This is because the NET is generally a test for global equivalence, and the free paths that are unique to M_A and M_B result in model fits that are not equal to one another. The BIC, conversely, increasingly prefers the true model, M_A , as d and n increase. The problem, as mentioned earlier, involves the fact that the BIC provides no mechanism for declaring models to be indistinguishable or to fit equally. For example, in the $d = 0.1$, $n = 200$ condition, the BIC prefers M_A about 70% of the time. The other 30% of the time, the BIC incorrectly prefers M_B . In contrast, the [Vuong \(1989\)](#) tests provide a formal mechanism for concluding that neither model should be preferred (either because they are indistinguishable or because their fits are equal).

Focusing on the distinguishability test results in [Figure 4](#), we observe "true" Type I error rates in the $d = 0$ condition: Models




Figure 4. Power associated with test statistics in Simulation 1. LRT = likelihood ratio test; CondLRT = Likelihood ratio test conditioned on a significant distinguishability test; BIC = Bayesian information criterion; NET = nesting and equivalence test.

are incorrectly declared to be distinguishable approximately 5% of the time. In addition, the hypothesis that models are indistinguishable is increasingly rejected with both d and n . Finally, focusing on the [Vuong \(1989\)](#) LRT results, we see near-zero Type I error rates in the $d = 0$ condition. This partially reflects the fact that the LRT should be used only when models are distinguishable (i.e., when H_0 for the distinguishability test is rejected). To expand on this point, the likelihood ratio test conditioned on a significant distinguishability test line displays power for the LRT conditioned on a rejected H_0 for the distinguishability test. These lines show Type I error rates that are slightly closer to .05, though they are still low. This result matches the observation by [Clark and McCracken \(2014\)](#) that Vuong’s sequential testing procedure can be conservative. In particular, Vuong proved an upper bound for the sequential tests’ Type I error that may actually be lower in practice.

Aside from Type I error, the power of the LRT approaches 1 more slowly than the power of the distinguishability test. The space between the distinguishability and LRT curves (i.e., between the solid line and the lower dashed lines in [Figure 4](#)) is related to the proportion of time that the hypothesis of indistinguishability is rejected but the hypothesis of equal-fitting models is not rejected. This should not be taken as evidence that the distinguishability test can be used without the [Vuong \(1989\)](#) LRT. If we reject the hypothesis of distinguishability, we simply conclude that the models can potentially be differentiated on the basis of fit. We can draw no conclusions about one model fitting better than the other.

Results associated with 90% confidence intervals of BIC differences are shown in [Table 1](#). It can be seen that the two types of interval estimates exhibit similar widths and endpoint variabilities. Coverage is somewhat different, however. When $d = 0$, models are indistinguishable and the intervals invalid. This results in very high coverage rates across both types of intervals. As d initially moves away from zero, both types of intervals exhibit coverage that is too low. Finally, as d becomes larger (and as n increases), the intervals converge toward the nominal coverage rate. The bootstrap intervals have a slight advantage here, moving toward a coverage of 0.9 faster than the [Vuong \(1989\)](#) intervals.

In practice, when models are more complex and generally easier to distinguish, the [Vuong \(1989\)](#) intervals may exhibit coverage

that is more comparable to the bootstrap intervals regardless of n . In the following simulation, we studied this conjecture.

Simulation 2: BIC Intervals

The previous simulation showed that when candidate models are nearly indistinguishable from one another, interval estimates associated with BIC differences (or with the likelihood ratio) generally stray from the nominal coverage rate. As models become more distinguishable, the intervals initially exhibit coverage that is too high, followed by coverage that is too low, followed by coverage that is just right. In this simulation, we further study the properties of these interval estimates in more complex models that

Table 1

Average Interval Widths, Variability in Endpoints, and Coverage of Differences in Nonnested Bayesian Information Criteria in Simulation 1

n	Path A	Width		Endpoint SD		Coverage	
		Vuong	Boot	Vuong	Boot	Vuong	Boot
200	0	8.351	10.907	2.952	3.159	0.999	0.995
	0.1	11.174	13.051	4.277	4.496	0.940	0.979
	0.2	17.052	17.779	6.092	6.113	0.844	0.883
	0.3	23.631	23.359	7.846	7.526	0.856	0.884
	0.4	29.427	28.452	9.401	8.867	0.868	0.885
500	0	33.607	32.772	10.731	10.352	0.857	0.874
	0.1	8.256	10.697	2.889	3.127	0.999	0.995
	0.2	14.460	15.762	5.357	5.692	0.846	0.896
	0.3	25.645	25.939	8.548	8.607	0.861	0.888
	0.4	37.184	36.880	11.690	11.438	0.888	0.899
1,000	0	47.055	45.960	14.536	13.906	0.885	0.895
	0.1	53.824	52.641	16.581	16.004	0.890	0.896
	0.2	8.257	10.630	2.960	3.193	0.999	0.993
	0.3	19.282	19.991	6.701	6.986	0.860	0.887
	0.4	35.735	35.835	11.497	11.489	0.876	0.891
	0.5	52.155	51.920	16.295	16.142	0.895	0.899
	0.1	66.860	65.839	20.507	20.017	0.893	0.896
	0.2	76.426	75.339	23.596	23.052	0.893	0.900
	0.3						
	0.4						

Note. Analytic intervals based on [Vuong \(1989\)](#); Boot = Nonparametric bootstrap intervals.

are generally distinguishable from one another. In this situation, the intervals' coverages should be closer to their advertised coverages.

Method

The simulation was set up in a manner similar to the simulation from Preacher and Merkle (2012), using the models from Figure 2. These models reflect four unique hypotheses about the relationships between nine observed variables. One thousand datasets were first generated from Model D, with unstandardized path coefficients being fixed to 0.2, residual variances being fixed to 0.8, and the exogenous variance associated with variable X_1 being fixed to 1.0. We then fit Models A–C to each dataset and obtained interval estimates of BIC differences. We examined sample sizes of $n = 200, 500, \text{ and } 1,000$ and compared 90% interval estimates from Vuong's (1989) theory to 90% interval estimates from the nonparametric bootstrap. Statistics of interest were those used in Simulation 1: average interval width, interval variability, and interval coverage.

Results

Results are displayed in Table 2. The two rightmost columns show that coverage is generally good for both methods; the coverages are all close to 0.9. The other columns show that the Vuong (1989) intervals tend to be slightly better than the bootstrap intervals: The Vuong widths are slightly smaller, and there is slightly less variability in the endpoints. These small advantages may not be meaningful in many situations, but the results at least show that the bootstrap intervals and Vuong intervals are comparable here. The Vuong intervals have a clear computational advantage, requiring only output from the two fitted models (and no extra data sampling or model fitting). As mentioned previously, the bootstrap intervals may still exhibit better performance when regularity conditions are violated. In the following section, we examine the Vuong statistics' application to nested model comparison.

Simulation 3: Tests of Nested Models

When we first introduced the Vuong (1989) test statistics, we mentioned that the statistics $n\hat{\omega}_s^2$ and $2n^{1/2}LR_{AB}$ provide unique

Table 2
Average Interval Widths, Variability in Endpoints, and Coverage of Differences in Nonnested Bayesian Information Criteria in Simulation 2

Models	n	Width		Endpoint SD		Coverage	
		Vuong	Boot	Vuong	Boot	Vuong	Boot
A and B	200	39.350	42.241	12.137	12.521	0.919	0.873
	500	57.937	59.620	18.173	18.516	0.901	0.875
	1,000	79.614	80.682	23.484	23.798	0.916	0.907
B and C	200	45.469	47.805	14.508	14.885	0.899	0.912
	500	69.064	70.317	20.832	21.036	0.901	0.905
	1,000	95.699	96.403	29.665	29.953	0.893	0.891
C and A	200	44.843	48.082	13.704	14.107	0.919	0.894
	500	65.386	67.320	19.633	19.853	0.910	0.891
	1,000	89.381	90.753	27.655	27.977	0.899	0.882

Note. Analytic intervals based on Vuong (1989); Boot = nonparametric bootstrap intervals.

tests of nested models. Unlike the classical LRT (also known as the *chi-square difference test*), tests involving these statistics make no assumption related to the correctness of the full model. In this simulation, we compare the performance of the classical LRT of nested models with the Vuong tests involving $n\hat{\omega}_s^2$ and $2n^{1/2}LR_{AB}$.

Method

The data-generating model was the path model displayed in Figure 5, where some parameters are represented by solid lines, some parameters are represented by dashed lines, and variance parameters are omitted. Path coefficients (unstandardized) were set to 0.2, residual variances were set to 0.8, and the single exogenous variance was set to 1.0. The dashed covariance parameters, further described later, were manipulated across conditions. After generating data from this model, we fit two candidate models to the data: the data-generating model (M_A) and a model with the dashed covariance paths omitted (M_B). The two candidate models are therefore nested, and M_A should be preferred when the dashed paths are nonzero. This is a best-case scenario for the traditional LRT because the full model (M_A) is correct.

Simulation conditions were defined by n —which assumed values of 200, 500, and 1,000—and by the value of the dashed covariance parameters, d . Within a condition, these covariance parameters simultaneously took values of 0, 0.025, 0.050, 0.075, 0.100, or 0.125. In each condition, we generated 1,000 datasets and fit both M_A and M_B to the data. We then computed four statistics: the distinguishability test of Equation 7 (which is actually a test of nested models here), the Vuong (1989) LRT involving $2n^{1/2}LR_{AB}$, the classical LRT based on chi-square differences, and the BIC difference. For each statistic, we recorded whether M_A was favored over M_B using the same criteria that were used in Simulation 1.

Results

Results are displayed in Figure 6. It can be seen that power associated with the two Vuong (1989) tests is very similar to power associated with the classical LRT (labeled *Chidiff* in the figure), with the distinguishability test exhibiting slightly smaller power at $n = 200$. In the $n = 500$ and $n = 1,000$ conditions, the two Vuong statistics are nearly equivalent to the classical LRT. The BIC, conversely, is slower to prefer M_A as a result of the fact that M_A has three extra parameters. These extra parameters must be sufficiently far from zero before they are “worth it,” as judged by the BIC.

This simulation demonstrates that, within the controlled environment examined here, little is lost if one uses Vuong (1989) tests to compare nested models. In some conditions, the Vuong statistics even exhibit slightly higher power than the classical LRT. Vuong showed that the limiting distributions of his statistics matched the distribution of the classical LRT (when the full model was correctly specified), and our simulation provides evidence that his statistics approach the limiting distribution at a similar rate as the classical LRT. This is relevant to sample size considerations: The Vuong statistics exhibit similar power to the classical LRT, so we can potentially use classical LRT results to guide us on sample sizes that lead to sufficient power for the Vuong tests. Simulation studies can also be carried out to study the power of the Vuong tests in specific situations.




Figure 5. Path diagrams reflecting the models used in Simulation 3.

Future work could compare the [Vuong \(1989\)](#) statistics to both the classical LRT and other robust statistics (see [Satorra & Bentler, 1994](#); [Savalei, 2014](#)) when the candidate models are both incorrect. This is trickier than it sounds, because it is difficult to design a framework that allows assessment of the statistics' Type I error rates. Such a framework would require that (a) the null hypothesis from [Equation 14](#) hold and (b) the data-generating model differ from the full model. This framework is necessary so that we know whether differences in statistics' power are paired with differences in statistics' Type I error rates.

General Discussion

The framework described in this article provides researchers with a general means to test pairs of SEMs for differences in fit. Researchers also gain a means to test whether pairs of SEMs are distinguishable in a population of interest and to test for differences between models' information criteria. In the following discussion, we provide details on extension of the tests to comparing multiple (>2) models. We also address regularity conditions, extension to other types of models, and a recommended strategy for applied researchers.

Comparing Multiple Models

The reader may wonder whether the foregoing theory extends to simultaneously testing multiple models. [Katayama \(2008\)](#) derived the joint distribution of likelihood ratio statistics comparing $m - 1$

models with a baseline model (i.e., when we have a unique likelihood ratio statistic comparing each of $m - 1$ models to the baseline model) and obtained a test statistic based on the sum of the $m - 1$ squared likelihood ratio statistics. We do not describe all the details here but, instead, supply some informal intuition underlying the tests.

In a situation in which we wish to compare multiple models, we can obtain the casewise likelihoods $\ell_j(\theta; \mathbf{x}_i), i = 1, \dots, n, j = 1, \dots, m$. We could then subject these casewise likelihoods to an analysis of variance (ANOVA), in which case (i) is a between-subjects factor and model (j) is a within-subject factor. We should expect a main effect of case, because the models will naturally fit some individuals better than others. The main effect of model, however, serves as a test of whether all m model fits are equal. In addition, the error variance informs us of the extent to which models are distinguishable: If the error variance is close to zero, we have evidence for indistinguishable models. The ANOVA framework could also be useful for post hoc tests, whereby one wishes to specifically know which model fits differ from which others.

We have not implemented the tests derived by [Katayama \(2008\)](#), and the ANOVA just described is not the same as Katayama's tests. For example, the ANOVA assumes sphericity for the within-subjects model factor, whereas Katayama explicitly estimated covariances between model likelihoods. Simultaneous tests of multiple SEMs generally provide interesting directions for future research.

Regularity Conditions

We previously outlined the regularity conditions associated with the [Vuong \(1989\)](#) tests, and we expand here on the conditions' relevance to applied researchers. The conditions under which the Vuong tests hold are fairly general, including existence of second-order derivatives of the log-likelihood, invertibility of the models' information matrices, and independence and identical distribution of the data vectors \mathbf{x}_i . The invertibility requirement can be violated when we apply the distinguishability test to compare mixture models with different numbers of components. This violation is related to the fact that a model with $c - 1$ components lies on the




Figure 6. Power associated with test statistics in Simulation 3. Chidiff = chi-square difference; Vuong_LRT = [Vuong \(1989\)](#) likelihood ratio test; BIC = Bayesian information criterion.

boundary of the parameter space of a model with c components. Jeffries (2003) presented some evidence that this can cause inflated Type I error rates, and Wilson (2015) further described problems with the test's application to mixture models. As of this writing, researchers have generally ignored these issues. Similar issues may arise if one model can be obtained from the other by setting some variance parameters equal to zero. The IID requirement can be violated by certain time-series models. This implies that the tests described here may not be applicable to some dynamic SEMs (where n reflects observed time points), though Rivers and Vuong (2002) extended the tests to handle these types of models.

Recommended Use

Using the NET and the methods described here (and also in Levy & Hancock, 2007, 2011), one gains a fuller comparison of candidate models. These procedures give researchers the ability to routinely test nonnested models for global equivalence, distinguishability, and differences in fit or information criteria. We recommend the following nonnested model comparison sequence:

1. Using the NET, evaluate models for global equivalence (this can be done prior to data collection). If models are not found to be equivalent, proceed to 2.
2. Test whether models are distinguishable, using the ω^2_* statistic (data must have already been collected). If models are found to be distinguishable, proceed to 3.
3. Compare models via the nonnested LRT or a confidence interval of BIC differences.

If one makes it to the third step, then the test or interval estimate may allow for the preference of one model. Otherwise, one cannot prefer either model to the other. In the case that models are found to be indistinguishable or to have equal fit, follow-up modeling can often be performed to further study the data. For example, one can often specify a larger model that encompasses both candidate models as special cases. This can provide information about important model parameters, which may lead to an alternative model that is a cross between the original candidate models. Overfitting would be a concern associated with this strategy, and it may often be sufficient to simply report results of the encompassing model. Aside from this strategy, however, there is nothing inherently wrong with having indistinguishable or equal-fitting models. Sometimes, there is simply not enough information in the data to differentiate two theories. The suggested sequence provides more information about the models' relative standings than do traditional comparisons via the BIC, which should help researchers to favor a model only when the data truly favor that model.

Distributional Assumptions

Although our current implementation allows researchers to carry out the recommended steps using SEMs estimated under multivariate normality, extensions to other assumed distributions are immediate. For example, if our observed variables are ordinal, our model may be based on a multinomial distribution. If the model is estimated via ML, then the tests can be carried out as

described here; we just need to obtain the nonstandard model output (casewise contributions to the likelihood and scores) for that model. Researchers often elect to use alternative discrepancy functions that do not correspond to a specific probability distribution, however, especially when the observed data are discrete. Examples of estimation methods that do not rely on well-defined probability distributions include weighted least squares and pairwise ML (e.g., Katsikatsou, Moustaki, Yang-Wallentin, & Jöreskog, 2012; Muthén, 1984). Work by Golden (2003) has implied that Vuong's (1989) theory can also be applied to models estimated via these alternative discrepancy functions. Further work is needed to obtain the necessary output from models estimated under these functions and to study the tests' applications.

Finally, we note that the ideas described throughout this article generally apply to situations in which one's goal is to declare a single model as the best. One may instead wish to average over the set of candidate models, drawing general conclusions across the set (e.g., Hoeting, Madigan, Raftery, & Volinsky, 1999). Though it is computationally more difficult, the model-averaging strategy allows the researcher to explicitly acknowledge that all of the models in the set are ultimately incorrect.

Computational Details

All results were obtained using Version 3.1.3 of the R system for statistical computing (R Development Core Team, 2014), using the add-on packages *lavaan* 0.5–17 (Rosseel, 2012) for fitting of the models and score computation, *nonnest2* 0.2 (Merkle & You, 2014) for carrying out the Vuong (1989) tests, and *simsem* 0.5–3 (Pornprasertmanit, Miller, & Schoemann, 2013) for simulation convenience functions. R and the packages *lavaan*, *nonnest2*, and *simsem* are freely available under the General Public License Version 2 from the Comprehensive R Archive Network at <http://CRAN.R-project.org/>. R code for replication of our results is available at <http://semtools.R-Forge.R-project.org/>.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723.
- Amemiya, Y., & Anderson, T. W. (1990). Asymptotic chi-square tests for a large class of factor analysis models. *Annals of Statistics*, *18*, 1453–1463.
- Bentler, P. M., & Satorra, A. (2010). Testing model nesting and equivalence. *Psychological Methods*, *15*, 111–123.
- Browne, M. W., & Arminger, G. (1995). Specification and estimation of mean- and covariance-structure models. In G. Arminger, C. C. Clogg, & M. E. Sobel (Eds.), *Handbook of statistical modeling for the social and behavioral sciences* (pp. 185–249). New York: Plenum Press.
- Byrne, B. M. (1994). Burnout: Testing for the validity, replication, and invariance of causal structure across elementary, intermediate, and secondary teachers. *American Educational Research Journal*, *31*, 645–673.
- Byrne, B. M. (2009). *Structural equation modeling with AMOS: Basic concepts, applications, and programming*. Boca Raton, FL: CRC Press.
- Casella, G., & Berger, R. L. (2002). *Statistical inference* (2nd ed.). Pacific Grove, CA: Thomson Learning.
- Chun, S. Y., & Shapiro, A. (2009). Normal versus noncentral chi-square asymptotics of misspecified models. *Multivariate Behavioral Research*, *44*, 803–827.
- Clark, T. E., & McCracken, M. W. (2014). Tests of equal forecast accuracy for overlapping models. *Journal of Applied Econometrics*, *29*, 415–430.

- Frese, M., Garst, H., & Fay, D. (2007). Making things happen: Reciprocal relationships between work characteristics and personal initiative in a four-wave longitudinal structural equation model. *Journal of Applied Psychology, 92*, 1084–1102.
- Golden, R. M. (2000). Statistical tests for comparing possibly misspecified and nonnested models. *Journal of Mathematical Psychology, 44*, 153–170.
- Golden, R. M. (2003). Discrepancy risk model selection test theory for comparing possibly misspecified or nonnested models. *Psychometrika, 68*, 229–249.
- Greene, W. H. (1994). *Accounting for excess zeros and sample selection in Poisson and negative binomial regression models* (NYU Working Paper No. EC94-10). Retrieved from <http://ssrn.com/abstract=1293115>
- Hershberger, S. L., & Marcoulides, G. A. (2013). The problem of equivalent structural models. In G. R. Hancock & R. O. Mueller (Eds.), *Structural equation modeling: A second course* (pp. 3–39). Charlotte, NC: Information Age.
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science, 14*, 382–417.
- Holzinger, K. J., & Swineford, F. A. (1939). *A study of factor analysis: The stability of a bifactor solution*. Chicago: University of Chicago Press.
- Jeffries, N. O. (2003). A note on 'Testing the Number of Components in a Normal Mixture.' *Biometrika, 90*, 991–994.
- Jones, L. V., & Tukey, J. W. (2001). A sensible formulation of the significance test. *Psychological Methods, 5*, 411–414.
- Katayama, N. (2008). *Portmanteau likelihood ratio tests for model selection*. Retrieved from http://www.economics.smu.edu.sg/femes/2008/CS_Info/papers/169.pdf
- Katsikatsou, M., Moustaki, I., Yang-Wallentin, F., & Jöreskog, K. G. (2012). Pairwise likelihood estimation for factor analysis models with ordinal data. *Computational Statistics and Data Analysis, 56*, 4243–4258.
- Kim, I. J., Zane, N. W. S., & Hong, S. (2002). Protective factors against substance use among Asian American youth: A test of the peer cluster theory. *Journal of Community Psychology, 30*, 565–584.
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *Annals of Mathematical Statistics, 22*, 79–86.
- Lee, S., & Hershberger, S. (1990). A simple rule for generating equivalent models in covariance structure modeling. *Multivariate Behavioral Research, 25*, 313–334.
- Levy, R., & Hancock, G. R. (2007). A framework of statistical tests for comparing mean and covariance structure models. *Multivariate Behavioral Research, 42*, 33–66.
- Levy, R., & Hancock, G. R. (2011). An extended model comparison framework for covariance and mean structure models, accommodating multiple groups and latent mixtures. *Sociological Methods & Research, 40*, 256–278.
- Lo, Y., Mendell, N. R., & Rubin, D. B. (2001). Testing the number of components in a normal mixture. *Biometrika, 88*, 767–778.
- MacCallum, R. C., Wegener, D. T., Uchino, B. N., & Fabrigar, L. R. (1993). The problem of equivalent models in applications of covariance structure analysis. *Psychological Bulletin, 114*, 185–199.
- Merkle, E. C., & You, D. (2014). nonnest2: Tests of non-nested models [Computer software manual]. Retrieved from <http://cran.r-project.org/package=nonnest2>
- Muthén, B. (1984). A general structural equation model with dichotomous, ordered categorical, and continuous latent variable indicators. *Psychometrika, 49*, 115–132.
- Nylund, K. L., Muthén, B. O., & Asparouhov, T. (2007). Deciding on the number of classes in latent class analysis: A Monte Carlo simulation study. *Structural Equation Modeling, 14*, 535–569.
- Pornprasertmanit, S., Miller, P., & Schoemann, A. (2013). simsem: SIMulated structural equation modeling [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=simsem>
- Pornprasertmanit, S., Wu, W., & Little, T. D. (2013). Abstract: Taking into account sampling variability of model selection indices: A parametric bootstrap approach. *Multivariate Behavioral Research, 48*, 168–169.
- Preacher, K. J., & Merkle, E. C. (2012). The problem of model selection uncertainty in structural equation modeling. *Psychological Methods, 17*, 1–14.
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology, 25*, 111–163.
- R Development Core Team. (2014). R: A language and environment for statistical computing [Computer software manual]. Retrieved from <http://www.r-project.org/>
- Rivers, D., & Vuong, Q. (2002). Model selection tests for nonlinear dynamic models. *Econometrics Journal, 5*, 1–39.
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software, 48*(2), 1–36. Retrieved from <http://www.jstatsoft.org/v48/i02/>
- Sani, F., & Todman, J. (2002). Should we stay or should we go? A social psychological model of schisms in groups. *Personality and Social Psychology Bulletin, 28*, 1647–1655.
- Satorra, A., & Bentler, P. M. (1994). Corrections to test statistics and standard errors in covariance structure analysis. In A. von Eye & C. C. Clogg (Eds.), *Latent variables analysis: Applications for developmental research* (pp. 399–419). Thousand Oaks, CA: Sage.
- Savalei, V. (2014). Understanding robust corrections in structural equation modeling. *Structural Equation Modeling, 21*, 149–160.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics, 6*, 461–464.
- Steiger, J. H., Shapiro, A., & Browne, M. W. (1985). On the multivariate asymptotic distribution of sequential chi-square statistics. *Psychometrika, 50*, 253–264.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica, 57*, 307–333.
- Wilson, P. (2015). The misuse of the Vuong test for non-nested models to test for zero-inflation. *Economics Letters, 127*, 51–53.

Appendix

Tests Involving Weighted Sums of Chi-Square Distributions

Here, we describe technical details underlying the tests whose limiting distributions are weighted sums of chi-square statistics. As stated in the main text, under the hypothesis that $\omega_*^2 = 0, n\hat{\omega}_*^2$ converges in distribution to a weighted sum of $k + q$ chi-square distributions (with one degree of freedom each), where k and q are the number of free parameters in M_A and M_B , respectively. The weights are defined as the squared eigenvalues of a matrix \mathbf{W} , which is described in the following paragraphs. In addition, for nested or indistinguishable models, $2n^{1/2}LR_{AB}$ converges in distribution to a similar weighted sum of chi-square distributions, with weights defined as the unsquared eigenvalues of the same matrix \mathbf{W} .

To obtain \mathbf{W} , let the matrices $\mathbf{U}_A(\boldsymbol{\theta}_A)$ and $\mathbf{V}_A(\boldsymbol{\theta}_A)$ be defined as

$$\mathbf{U}_A(\boldsymbol{\theta}_A) = E \left[\frac{\partial^2 \ell(\boldsymbol{\theta}_A; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_A \partial \boldsymbol{\theta}_A'} \right]$$

and

$$\mathbf{V}_A(\boldsymbol{\theta}_A) = E \left[\frac{\partial \ell(\boldsymbol{\theta}_A; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_A} \cdot \frac{\partial \ell(\boldsymbol{\theta}_A; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_A'} \right],$$

which can be obtained from a fitted M_A 's information matrix and the cross-product of scores (see [Equations 3 and 4](#)), respectively.

The matrices $\mathbf{U}_B(\boldsymbol{\theta}_B)$ and $\mathbf{V}_B(\boldsymbol{\theta}_B)$ are defined similarly. Further, define $\mathbf{V}_{AB}(\boldsymbol{\theta}_A, \boldsymbol{\theta}_B)$ as

$$\mathbf{V}_{AB}(\boldsymbol{\theta}_A, \boldsymbol{\theta}_B) = E \left[\frac{\partial \ell(\boldsymbol{\theta}_A; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_A} \cdot \frac{\partial \ell(\boldsymbol{\theta}_B; \mathbf{x}_i)}{\partial \boldsymbol{\theta}_B'} \right],$$

which can be obtained by taking products of $s(\hat{\boldsymbol{\theta}}_A; \mathbf{x}_i)$ and $s(\hat{\boldsymbol{\theta}}_B; \mathbf{x}_i)$. The matrix \mathbf{W} is then defined as

$$\mathbf{W} = \begin{bmatrix} -\mathbf{V}_A(\boldsymbol{\theta}_A)\mathbf{U}_A^{-1}(\boldsymbol{\theta}_A) & -\mathbf{V}_{AB}(\boldsymbol{\theta}_A, \boldsymbol{\theta}_B)\mathbf{U}_B^{-1}(\boldsymbol{\theta}_B) \\ \mathbf{V}_{AB}'(\boldsymbol{\theta}_A, \boldsymbol{\theta}_B)\mathbf{U}_A^{-1}(\boldsymbol{\theta}_A) & \mathbf{V}_B(\boldsymbol{\theta}_B)\mathbf{U}_B^{-1}(\boldsymbol{\theta}_B) \end{bmatrix}.$$

The eigenvalues of \mathbf{W} then determine the weights involved in the limiting sum of chi-square distributions (see [Vuong, 1989](#), for the proof and further details).

Received February 26, 2014

Revision received February 27, 2015

Accepted March 6, 2015 ■