

Empirical Likelihood for Linear Structural Equation Models with Dependent Errors

Y. Samuel Wang^{a*}, Mathias Drton^a

Received 00 Month 2012; Accepted 00 Month 2012

We consider linear structural equation models that are associated with mixed graphs. The structural equations in these models only involve observed variables, but their idiosyncratic error terms are allowed to be correlated and non-Gaussian. We propose empirical likelihood (EL) procedures for inference, and suggest several modifications, including a profile likelihood, in order to improve tractability and performance of the resulting methods. Through simulations, we show that when the error distributions are non-Gaussian, the use of EL and the proposed modifications may increase statistical efficiency and improve assessment of significance. Copyright © 2017 John Wiley & Sons, Ltd.

Keywords: Empirical likelihood, causal inference, graphical model, structural equation model

1. Introduction

Structural equation models (SEMs) are multivariate statistical models in which each considered variable is a function of other variables and a stochastic error term. Often, some of these other variables are latent (Bollen, 1989). This paper, however, focuses on SEMs in which the effects of latent variables are summarized. Adopting the dominant linear paradigm, we will thus be concerned with models in which linear functions relate only observed variables, but error terms may be dependent. Such models are sometimes referred to as semi-Markovian (Shpitser & Pearl, 2006). Avoiding any explicit specification of latent confounding, the models play an important role in exploration of cause-effect structures (Colombo et al., 2012; Pearl, 2009; Richardson & Spirtes, 2002; Spirtes et al., 2000; Wermuth, 2011). Much insight about the models can be gained from a natural graphical representation by mixed graphs/path diagrams that originates in work of Wright (1921).

^aDepartment of Statistics
University of Washington
Seattle, WA 98103 USA
*Email: ysamwang@uw.edu

Formally, let Y_1, \dots, Y_n be a multivariate sample with each observation indexed by a set V . So, $Y_i = (Y_{vi})_{v \in V}$ with each Y_{vi} real-valued. Now consider the system of structural equations

$$Y_{vi} = \mu_v + \sum_{u \in V \setminus v} \beta_{vu} Y_{ui} + \epsilon_{vi}, \quad v \in V, \quad i = 1, \dots, n, \quad (1)$$

where the μ_v and β_{vu} are unknown parameters and the ϵ_{vi} are random errors. Define vectors $\epsilon_i = (\epsilon_{vi})_{v \in V}$ and $\mu = (\mu_v)_{v \in V}$, and a matrix $B = (\beta_{vu})_{v,u \in V}$ with $\beta_{vu} = 0$ if $v = u$. We assume that the error vectors ϵ_i are independent and identically distributed, have zero means, and have covariance matrix $\mathbb{E}(\epsilon_i \epsilon_i^t) = \Omega = (\omega_{vu})$. However, we do not specify any parametric form for their distribution. For each i , the equations in (1) can be written as $Y_i = \mu + B Y_i + \epsilon_i$. If $(I - B)$ is non-singular, then this system is solved uniquely by $Y_i = (I - B)^{-1}(\mu + \epsilon_i)$. This solution has mean vector $(I - B)^{-1}\mu$ and covariance matrix

$$\Sigma(B, \Omega) := (I - B)^{-1}\Omega(I - B)^{-t}. \quad (2)$$

Specific models are now obtained by hypothesizing that a particular collection of coefficients β_{vu} and error covariances ω_{vu} is zero.

An SEM can be represented conveniently by a path diagram/mixed graph $G = (V, E_{\rightarrow}, E_{\leftrightarrow})$. Here, the vertex set V yields a correspondence between the nodes of the graph and the observed variables. The set E_{\rightarrow} is a set of *directed edges* $u \rightarrow v$, which encode that variable u may have a direct effect on variable v . The set E_{\leftrightarrow} comprises *bidirected edges* $u \leftrightarrow v$ that indicate that the errors ϵ_{ui} and ϵ_{vi} may be correlated. Define the set of *parents* of node v as $\text{pa}(v) = \{u \in V : u \rightarrow v \in E_{\rightarrow}\}$. Similarly, define a set of *siblings* as $\text{sib}(v) = \{u \in V : u \leftrightarrow v \in E_{\leftrightarrow}\}$. Bidirected edges have no orientation, and $v \in \text{sib}(u)$ if and only if $u \in \text{sib}(v)$. Now, the graph G induces a model through the requirement that

$$B \in \mathcal{B}(G) := \{B \in \mathbb{R}^{V \times V} : \det(I - B) \neq 0, \beta_{vu} = 0 \text{ if } u \notin \text{pa}(v)\}, \quad (3)$$

$$\Omega \in \mathcal{W}(G) := \{\Omega \in \mathbb{R}^{V \times V} : \Omega \text{ pos. def., } \omega_{vu} = \omega_{uv} = 0 \text{ if } v \notin \text{sib}(u)\}. \quad (4)$$

We emphasize that our treatment allows the model to have feedback loops, that is, G may have directed cycles.

1.1. Related Work

Frequently, the errors in a SEM, and consequently also the observations Y_i , are assumed to be multivariate Gaussian which yield *maximum likelihood estimates* (MLEs). The Gaussian likelihood is often maximized using generic optimization methods; as done in the popular packages `sem` (Fox et al., 2017) and `lavaan` (Rosseel, 2012) for R (R Core Team, 2017). The coordinate-descent methods proposed by Drton et al. (2009) and Drton et al. (2017) can be a useful computational alternative that largely avoids convergence issues.

As a less parametric method, *generalized least squares* (GLS) minimizes a discrepancy between the sample covariance and the covariance implied by the parameters. Although the estimates are slightly more robust to misspecification, they are still asymptotically equivalent to the Gaussian MLEs (Olsson et al., 2000). When multivariate Gaussianity is inappropriate, MLEs and GLS generally lose statistical efficiency and yield incorrectly calibrated confidence intervals. *Weighted least squares* methods (WLS)—also called *asymptotically distribution free*—weight the discrepancy between the observed and hypothesized covariance structure by explicitly estimated fourth moments. Although WLS estimates are consistent and produce asymptotically correct confidence intervals even with non-Gaussian errors, the estimation of higher order moments may come at a loss of statistical efficiency and cause convergence issues, which has limited their use (Muthén & Kaplan, 1992).

Chaudhuri et al. (2007) propose using the *empirical likelihood* (EL) of Owen (2001) to estimate a covariance matrix with structural zeros. In our setup, this corresponds to the special case of a mixed graph with no directed edges. Kolenikov & Yuan (2009) use EL to estimate the parameters of a linear SEM. In contrast to the mixed graph formulation, Kolenikov & Yuan (2009) consider the case where the latent variable structure is explicitly modeled and all errors are independent. The EL approach is appealing as it gives consistent estimates and asymptotically correct confidence intervals even when the errors are not multivariate Gaussian. However, EL can present numerous practical difficulties when the sample size is small relative to the number of parameters or estimating equations used. Moreover, standard implementation of EL methods is computationally feasible only for systems with a handful of variables. We believe that these issues have prevented application of EL to linear SEMs beyond what was done by Kolenikov & Yuan (2009).

1.2. Contribution

In this article, we apply the empirical likelihood framework to SEMs represented by mixed graphs and propose several modifications to a naive approach which address the most salient practical concerns:

- (i) We show that in the mixed graph setting, the covariance parameters Ω can be profiled out. This greatly reduces the computational burden by reducing the number of estimating equations imposed and parameters directly estimated. It also naturally encodes the positive definite constraint on Ω and yields a positive definite estimate of Ω for any point B with a well defined empirical likelihood.
- (ii) When maximizing the empirical likelihood, we leverage a recent insight and directly incorporate gradient information in a quasi-Newton procedure instead of the typical derivative-free approaches to empirical likelihood optimization. This again yields substantial computational savings.
- (iii) We use the *adjusted empirical likelihood* (AEL), first proposed by Chen et al. (2008). This adjustment ensures that an empirical likelihood and corresponding gradient is well defined for every value in the parameter space.
- (iv) We apply the idea of *extended empirical likelihood* (EEL), which furnishes drastically improved coverage of confidence intervals at small sample sizes (Tsao & Wu, 2014).

Our simulations show that with these proposed modifications, empirical likelihood becomes an attractive alternative for practitioners concerned with non-Gaussianity in structural equation modeling.

2. Background on Empirical Likelihood

Let $Y = (Y_1, \dots, Y_n)$ be a sample from an m -variate distribution P belonging to a non-/semiparametric statistical model \mathcal{M} . Let \mathcal{P}_n be the $n - 1$ dimensional probability simplex. For $p = (p_1, \dots, p_n) \in \mathcal{P}_n$, define the log-empirical likelihood $\ell(p; Y) = \sum_{i=1}^n \log(p_i)$. This is the log-likelihood of the sample under the discrete distribution with mass p_i at each point Y_i . Suppose we are interested in a parameter $\theta = \theta(P)$ taking values in $\Theta \subseteq \mathbb{R}^d$ such that for a map $G : \mathbb{R}^m \times \mathbb{R}^d \mapsto \mathbb{R}^q$ we have $\mathbb{E}_P G(Y_i, \theta(P)) = 0$ for all $P \in \mathcal{M}$. The log-empirical likelihood at a given parameter value θ is then

$$\ell(\theta; Y) = \max_{p \in \mathcal{P}_\theta} \ell(p; Y) = \max_{p \in \mathcal{P}_\theta} \sum_{i=1}^n \log(p_i), \quad (5)$$

where the feasible set

$$\mathcal{P}_\theta = \left\{ p \in \mathcal{P}_n : \sum_{i=1}^n p_i G(Y_i, \theta) = 0 \right\} \quad (6)$$

reflects that the expectation of $G(\cdot; \theta)$ vanishes for distributions compatible with θ .

The empirical likelihood (EL) from (5) provides a basis for statistical inference. Maximizing it over $\theta \in \Theta$ yields the *maximum empirical likelihood estimator*

$$\check{\theta} = \arg \max_{\theta} \ell(\theta; Y) \quad (7)$$

that we refer to as MELE. Ratios of the EL yield *empirical likelihood ratio* statistics. [Owen \(1988\)](#) derives an EL analogue of Wilk's Theorem, and the result was expanded to the general estimating equation framework by [Qin & Lawless \(1994\)](#). The specific regularity conditions needed are discussed in Section 3.3, and the results imply under very general conditions that the MELE is consistent and asymptotically normal. In addition, EL ratio statistics have limiting χ^2 distributions that can be used to calibrate statistical tests and create confidence intervals or regions. For a detailed exposition of these ideas, we refer readers to [Owen \(2001\)](#).

The nice theoretical properties for EL, however, come at a high practical cost. The practical issues become particularly pressing for applications to linear SEMs, for which the number of parameters and estimating equations generally grow on the order of m^2 , where $m = |V|$ is the number of variables considered. We describe three difficulties that complicate the direct use of EL for SEMs:

- (i) For some values θ , the origin may be outside the convex hull of $\{G(Y_i, \theta) : i = 1, \dots, n\}$, in which case the feasible set \mathcal{P}_θ from (6) is empty and the EL at θ is zero. This “convex hull problem” occurs more often when the sample size is small relative to the number of estimating equations or when the data is skewed. As discussed by [Grendár & Judge \(2009\)](#), it is possible that $\mathcal{P}_\theta = \emptyset$ for all parameter vectors θ , which is known as the “empty set problem”. In addition, the log-EL is typically not a convex function ([Chaudhuri et al., 2017](#)), and finding an initial point that has well-defined EL and is in the basin of attraction of the MELE can be difficult.
- (ii) The optimization problem defining the log-EL $\ell(\theta; Y)$ from (5) is typically solved iteratively through its dual. Although this problem is convex, it can be computationally burdensome when the number of estimating equations, which corresponds to the number of dual variables, is large.
- (iii) Confidence intervals based on the asymptotic normal variance and χ^2 likelihood ratio calibration have been shown to often undercover at small sample sizes ([Tsao & Wu, 2014](#)).

3. Empirical Likelihood for SEMs

We now turn to the application of EL to SEMs. For expository simplicity, we assume throughout that our observations are centered. In other words, the intercept parameter vector μ for (1) is zero, so that $\mathbb{E}(Y_i) = 0$. However, our ideas extend straightforwardly to the case where we also make inference about $\mu \neq 0$.

3.1. Profiled Formulation

Consider the linear SEM given by a mixed graph $\mathcal{G} = (V, E_{\rightarrow}, E_{\leftrightarrow})$, as defined in the Introduction. The general framework laid out in Section 2 can be applied directly to such a model by taking the covariance matrix of the observations Y_i as the general parameter θ . We may then define an EL at a pair of parameter matrices (B, Ω)

as the EL at the covariance matrix $\Sigma(B, \Omega)$ from (2). In such a direct application to the linear SEM, the log-EL function $\ell(B, \Omega; Y)$ is the maximum of the log-EL $\ell(p; Y)$ over the set

$$\mathcal{P}_{\Sigma(B, \Omega)} = \left\{ p \in \mathcal{P}_n : \sum_{i=1}^n p_i Y_i = 0, \sum_{i=1}^n p_i [\text{vech} (Y_i Y_i^T) - \text{vech} \Sigma(B, \Omega)] = 0 \right\}. \quad (8)$$

Here, vech is the half-vectorization operator for symmetric matrices. Under this formulation, there are m constraints for the mean and $m(m+1)/2$ covariance constraints, and the MELE is computed by optimization with respect to the pair of $m \times m$ matrices (B, Ω) , with Ω restricted to be positive definite.

Inspection of the covariance constraints reveals that a great simplification is possible by profiling out Ω . Indeed, the covariance constraint yields an explicit solution for Ω given B , $Y = (Y_1, \dots, Y_n)$, and p . Specifically, with $\Pi = \text{diag}(p_1, \dots, p_n)$, we have

$$Y \Pi Y^T = \Sigma(B, \Omega) = (I - B)^{-1} \Omega (I - B)^{-t} \iff (I - B) Y \Pi Y (I - B)^t = \Omega. \quad (9)$$

The entries of Ω are either constrained to be zero or freely varying. No constraints arise from the freely varying entries, and we may base estimation of B on only the structural zeros in Ω , that is,

$$\{(I - B) Y \Pi Y (I - B)^t\}_{uv} = 0 \quad \forall \{u, v\} \notin E_{\leftrightarrow}.$$

Once a solution for B is found, we may simply compute $\Omega = \Omega(B)$ by setting $\omega_{uv} = \{(I - B) Y \Pi Y (I - B)^t\}_{uv}$ for $u = v$ or $\{u, v\} \in E_{\leftrightarrow}$. The profile log-EL in this approach is the function

$$\ell(B; Y) = \max_{p \in \mathcal{P}_B} \ell(p; Y) \quad (10)$$

obtained from the set of weight vectors

$$\mathcal{P}_B = \left\{ p \in \mathcal{P}_n : \sum_{i=1}^n p_i Y_i = 0, \sum_{i=1}^n p_i \left(Y_{vi} - \sum_{s \in \text{pa}(v)} \beta_{vs} Y_{si} \right) \left(Y_{ui} - \sum_{t \in \text{pa}(u)} \beta_{ut} Y_{ti} \right) = 0 \quad \forall \{v, u\} \notin E_{\leftrightarrow} \right\}. \quad (11)$$

The MELE \check{B} is found by maximizing $\ell(B; Y)$ over the set $\mathcal{B}(G)$ from (3), and then $\check{\Omega} = \Omega(\check{B})$. We emphasize that there are now only $m(m-1)/2 - |E_{\leftrightarrow}|$ covariance constraints, and only the matrix B needs to be optimized.

Following a standard strategy, we evaluate $\ell(B; Y)$, that is, solve the “inner maximization” in (10) at a fixed B , through the dual problem. Strong duality holds because the constraints in (11) are linear in the weights p_i . Let $G(Y_i, B)$ be the map with coordinates $G_v(Y_i, B) = Y_{iv}$ for $v \in V$ and $G_{uv}(Y_i, B) = g_u(Y_i, B)g_v(Y_i, B)$ for each nonedge $\{u, v\} \notin E_{\leftrightarrow}$, where $g_v(Y_i, B) = Y_{vi} - \sum_{s \in \text{pa}(v)} \beta_{vs} Y_{si}$. With dual variables $\alpha \in \mathbb{R}$ and $\lambda \in \mathbb{R}^{m+m(m-1)/2-|E_{\leftrightarrow}|}$, the Lagrangian for the inner optimization over \mathcal{P}_B is

$$L_B(p, \alpha, \lambda) = - \sum_{i=1}^n \log(p_i) + \alpha \left(\sum_{i=1}^n p_i - 1 \right) + n \sum_{i=1}^n p_i \left(\sum_{v \in V} \lambda_v G_v(Y_i, B) + \sum_{\{u, v\} \notin E_{\leftrightarrow}} \lambda_{uv} G_{uv}(Y_i, B) \right). \quad (12)$$

Maximizing over the weights, with $\alpha = n$, we find

$$\check{p}_i = \frac{1}{n} \frac{1}{1 + \sum_{v \in V} \lambda_v G_v(Y_i, B) + \sum_{\{u, v\} \notin E_{\leftrightarrow}} \lambda_{uv} G_{uv}(Y_i, B)}, \quad (13)$$

and substitution into L_B yields a convex dual function of λ . We optimize it via Newton-Raphson with a backtracking line search to ensure $0 \leq p_i \leq 1$.

In the “outer maximization”, we optimize $\ell(B; Y)$ with respect to B using a gradient based quasi-Newton method. Although we can only evaluate $\ell(B; Y)$ numerically, once we have the optimal dual variables λ and the corresponding weights from (13), we can analytically compute the gradient of $\ell(B; Y)$ as

$$\nabla \ell(B; Y) = -\lambda^T \sum_{i=1}^n \check{p}_i \nabla G(Y_i, B); \quad (14)$$

see Chaudhuri et al. (2017). The Hessian, however, cannot be computed in closed form, so we use BFGS which builds an approximate Hessian via the gradient.

Although both formulations yield the same MELE, the profile approach from (10) and (11) drastically eases difficulties (i) and (ii) discussed in Section 2 as the number of estimating equations for the covariance is reduced to $m + m(m - 1)/2 - |E_{\leftrightarrow}|$. This reduces the number of dual variables to optimize in the inner maximization. Moreover, when profiled, the outer maximization searches over only $B \in \mathcal{B}(G)$ while the naive direct formulation from (8) requires a search over both $B \in \mathcal{B}(G)$ and $\Omega \in \mathcal{W}(G)$; in particular, positive definiteness of Ω needs to be respected in the naive optimization. Finally, satisfying the convex hull condition for the error covariances typically requires a simultaneous good choice of B and Ω . The directed edge weights can be easily initialized with regression estimates, but the covariance parameters are typically more difficult to specify. In Section 4, we show that the computational advantages produce substantial gains in computation time and converge to a valid stationary point at a much higher proportion of the time even when the sample size is small.

3.2. Small Sample Improvements

In addition to reformulating the optimization problem, we make two modifications to improve the performance of EL for SEMs. We apply *adjusted empirical likelihood* (AEL) to improve the search for a MELE and use *extended empirical likelihood* (EEL) to improve the coverage of confidence intervals.

Chen et al. (2008) proposed AEL to alleviate the convex hull problem mentioned in difficulty (i) above. The adjustment amounts to adding a pseudo-observation whose contribution to the estimating equations is $G_{n+1}(B) = -a_n \bar{G}(B) = -a_n \frac{1}{n} \sum_{i=1}^n G(Y_i, B)$ for a choice of $a_n > 0$. Adding this term ensures that no matter the value of B , the set of feasible weight vectors, now in \mathcal{P}_{n+1} , is non-empty. Hence, the log-AEL $\ell^a(B; Y)$ and its gradient

$$\nabla \ell^a(B; Y) = -\lambda^T \sum_{i=1}^n \left[\check{p}_i + \left(-\frac{a_n}{n} \right) \check{p}_{n+1} \right] \nabla G(Y_i, B) \quad (15)$$

are well defined across the entire parameter space. Chen et al. (2008) show that AEL retains the asymptotic properties of the original EL when $a_n = o(n^{2/3})$, and suggest $a_n = \log(n)/2$. We adopt this choice.

The terms in our covariance constraint are products, $G_{uv}(Y_i, B) = g_v(Y_i, B)g_u(Y_i, B)$. This is generally not true for the added term $G_{n+1}(B)$ and is not clear how to define an appropriately sparse and positive definite matrix $\Omega(B)$ using AEL weights. Thus, we propose finding an estimate \check{B} that maximizes the AEL and computing $\check{\Omega} = \Omega(\check{B})$ based on weights from recalculating the original EL at \check{B} . As demonstrated in our numerical experiments, this approach alleviates some convergence issues but, of course, the original EL may be zero at the AEL maximizer \check{B} , in which case we do not have an estimate of Ω and say that the AEL procedure has not converged.

To address undercoverage of confidence regions for smaller samples, as described in difficulty (iii), we adopt the EEL of Tsao & Wu (2014) who show that their χ^2 -calibrated EEL confidence regions outperform those from the

original EL. Assuming the MELE \check{B} exists, a positive EEL may be defined for any matrix $B \in \mathcal{B}(G)$ by taking the original EL at a convex combination of B and \check{B} . Specifically, the log-EEL suggested by [Tsao & Wu \(2014\)](#) is

$$\ell^e(B; Y) = \ell(h^{-1}(B, Y), Y) \quad (16)$$

for $h(B, Y) = \check{B} + \gamma(n, \ell(\theta; Y))(B - \check{B})$ with $\gamma(n, \ell(B; Y)) = \left(1 + \frac{2(-n \log(n) - \ell(B; Y))}{2n}\right)$.

3.3. Asymptotic Distribution of Empirical Likelihood Estimators

It follows from [Qin & Lawless \(1994, Theorem 1\)](#) that under the following assumptions, MELEs are asymptotically normal and empirical likelihood ratios converge to χ^2 limits. The same is true for the modifications from Section 3.2.

Proposition 1

Let $G = (V, E_{\rightarrow}, E_{\leftrightarrow})$ be a mixed graph, let $B_0 \in \mathcal{B}(G)$ and $\Omega_0 \in \mathcal{W}(G)$. Let ϵ be a zero-mean random vector with covariance matrix Ω_0 . Assume that:

- (a) The Jacobian of the parametrization $(B, \Omega) \mapsto \Sigma(B, \Omega)$ defined on $\mathcal{B}(G) \times \mathcal{W}(G)$ has full rank at (B_0, Ω_0) .
- (b) The joint distribution of ϵ and $\epsilon^{(2)} = (\epsilon_v \epsilon_u : v, u \in V)$ is non-degenerate and has finite third moments.

If Y_1, \dots, Y_n is an i.i.d. sample from the distribution determined by $(B_0, \Omega_0, \epsilon)$, i.e., the distribution of $(I - B_0)^{-1}\epsilon$, then the MELE $\check{\theta} = (\text{vech}[\check{B}], \text{vech}[\Omega(\check{B})])$ is asymptotically normal with

$$\sqrt{n}(\check{\theta} - \theta_0) \rightarrow N(0, V), \quad V^{-1} = \mathbb{E}\left(\frac{\partial G(Y, \theta_0)}{\partial \theta}\right)^t \mathbb{E}[G(Y, \theta_0)G(Y, \theta_0)^t]^{-1} \mathbb{E}\left(\frac{\partial G(Y, \theta_0)}{\partial \theta}\right). \quad (17)$$

Here, G is given by the estimating equations corresponding to the naive formulation in (8). Furthermore, EL ratio statistics have χ^2 limits. In particular, for $q = m + m(m + 1)/2$ and $d = |E_{\rightarrow}| + |E_{\leftrightarrow}| + m$, we have

$$2(-n \log(n) - \ell(\check{\theta}; Y)) \rightarrow \chi^2_{(q-d)}, \quad 2[\ell(\check{\theta}; Y) - \ell(\theta_0; Y)] \rightarrow \chi^2_d. \quad (18)$$

We sketch the proof of the proposition in the appendix.

If the rank condition from (a) holds, then the rational map $(B, \Omega) \mapsto \Sigma(B, \Omega)$ has full rank Jacobian at almost all choices of (B, Ω) , and the map is generically finite-to-one. There is thus a connection to local/finite identifiability of (B, Ω) from the covariance matrix. For state-of-the-art methods for determining identifiability see [Foygel et al. \(2012\)](#); [Chen \(2016\)](#); [Drton & Weihs \(2016\)](#).

4. Numerical Simulations

We now show a series of numerical experiments to evaluate the effectiveness of the proposed methods and compare the results to existing methods.

4.1. Convergence of Optimizers for Naive vs Profile Formulation

We first compare the naive/direct procedure which explicitly estimates B and Ω to the profiled procedure which only involves B . For both procedures, we use the original EL and adjusted EL. We also consider a hybrid method, which first finds the maximum AEL point to initialize a search which then uses original EL. We randomly generate

acyclic mixed graphs with 8 nodes, 10 directed edges, and 6 bidirected edges. We randomly select directed edges $u \rightarrow v$ from all pairs such that $u < v$ and then select bidirected edges $u \leftrightarrow v$ from the remaining unselected pairs. This setup ensures that (B, Ω) are generically identifiable from $\Sigma(B, \Omega)$ by the result of [Brito & Pearl \(2002\)](#).

We generate random true parameter matrices $B = (\beta_{uv})$ and $\Omega = (\omega_{uv})$ as follows. The coefficients β_{uv} are drawn uniformly from $(-1, . - .2) \cup (.2, 1)$. For Ω , we draw off-diagonal elements $\omega_{uv} = \omega_{vu}$, $u \neq v$, uniformly from $(-.8, . - .3) \cup (.3, .8)$. We then use exponential draws to set $\omega_{vv} = \sum_{u \neq v} |\omega_{uv}| + 1 + \exp(1)$.

We consider errors from four distributions. First, we generate centered multivariate Gaussian errors with covariance matrix Ω . Second, we generate them from a multivariate T -distribution with 4 degrees of freedom, which we denote by T_4 , again with expectation zero and covariance matrix Ω . Third, we consider log-normal errors. In this case, we simulate a multivariate Gaussian vector Z , centered and with covariance matrix equal to the correlation matrix C that corresponds to Ω . We then set the error vector to $\epsilon = \exp(Z) - \sqrt{e}$, which yields covariance matrix $e(\exp(C) - 1)$. Finally, in order to draw a multivariate distribution with recentered gamma marginals and covariance Ω , we follow the steps:

1. Draw $\epsilon_v \sim \text{gamma}(\text{shape} = \omega_{vv} - \sum_{u \neq v} |\omega_{uv}|, \text{scale} = 1)$.
2. For each $\{u, v\} \in E_{\leftrightarrow}$, generate $\delta_{uv} \sim \text{gamma}(\text{shape} = |\omega_{uv}|, \text{scale} = 1)$ and a random sign $\xi_{uv} \in \{-1, 1\}$.
3. If $\omega_{uv} > 0$, add $\xi_{uv}\delta_{uv}$ to ϵ_{ui} and ϵ_{vi} . If $\omega_{uv} < 0$, add $\xi_{uv}\delta_{uv}$ to ϵ_{ui} and $-\xi_{uv}\delta_{uv}$ to ϵ_{vi} .
4. Subtract the true mean from each error term so that it has mean 0.

All optimizations are initialized with a procedure from [Drton et al. \(2017\)](#), where the free elements of B are calculated via least squares. The resulting residuals are used to initialize the non-zero values ω_{uv} . If a row is not diagonally dominant, the off-diagonal elements are scaled so that $\sum_{j \neq i} |\omega_{ij}| < .9 \times \omega_{ii}$ to ensure Ω is positive definite.

Figure 1 shows that in all cases the profiled formulation converges at least as often as the naive formulation. AEL converges more often than original EL, and the hybrid procedure converges the most often. Even at a sample size of $n = 100$, the profiled problem converges nearly every single time, except in the case of log-normal errors. Figure 2 shows that the profiled form can be up to 40 times faster on average than the naive form.

4.2. Estimation Error

We now explore the estimation errors resulting from different approaches. We compare both original EL and AEL to the Gaussian MLE computed as in [Drton et al. \(2009\)](#), GLS, and WLS. The latter two estimates are computed using the R package `lavaan` ([Rosseel, 2012](#)). We also include a hybrid procedure that finds the Gaussian MLE \hat{B} and then uses the resulting residuals and the maximum EL weights at \hat{B} to form an estimate $\check{\Omega} = (I - \hat{B})^t Y \Pi Y^t (I - \hat{B})^t$. Note that the T_4 distribution does not have finite 6th moments, so the limiting distributions from Proposition 1 may not hold; however, all estimation procedures still appear to be consistent.

Proceeding as in Section 4.1, we generate 1000 graphs for each error distribution and sample size. To measure estimation accuracy, we average the relative error $\|\text{vech}(\check{\Sigma}) - \text{vech}(\Sigma)\|^2 / \|\text{vech}(\Sigma)\|^2$ for $\Sigma(B, \Omega)$ across each of the simulation runs in which all methods converge; recall Figure 1. The results are shown in Figure 3.

In general, there is no substantial difference in accuracy between the adjusted and original empirical likelihood methods. For the Gaussian case, MLE and GLS perform better than the methods which do not assume Gaussianity, but the improvement is slight. In the T_4 and log-normal case, the EL procedures perform substantially

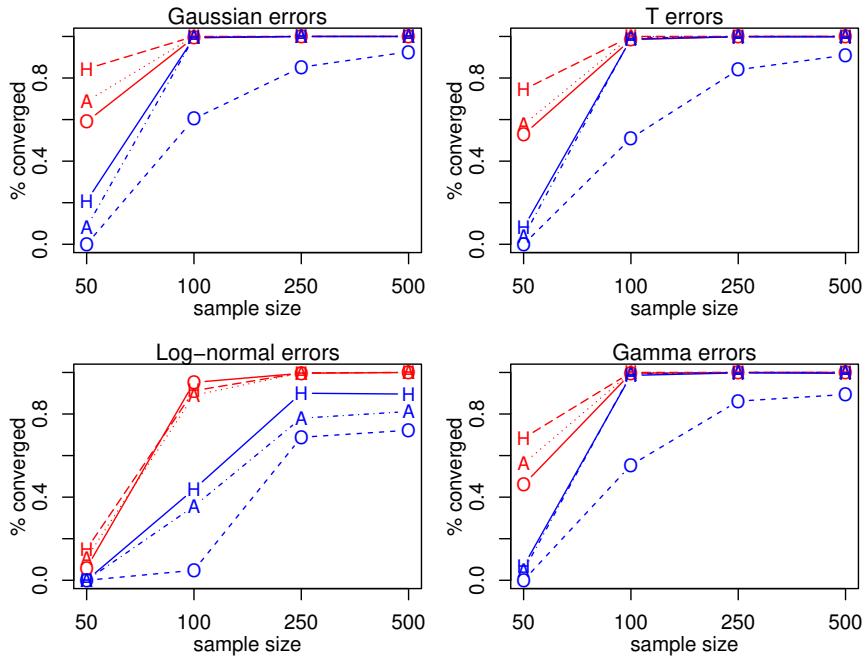


Figure 1. Proportion of 500 simulations which converge to a valid stationary point, plotted versus the sample size. O- original EL; A- adjusted EL; H- hybrid EL. Red points indicate the profile formulation; blue points indicate the naive formulation.

better than the other methods. Finally, for the gamma case, the hybrid method seems to outperform the other methods, followed closely by the EL methods; however, the differences between the methods are not substantial. In Figure 4, all methods converge more than 95% of the time in all distributions, except for the log-normal case. In this case, the WLS procedure still only converges roughly 90% at $n = 1000$.

4.3. Confidence Regions

We examine the coverage frequencies of joint confidence regions for the parameters β_{uv} and ω_{uv} . We construct Wald regions using the estimates of $\text{Var}(\hat{\theta})$ from the Gaussian MLE, GLS, and WLS. We also calculate a sandwich variance estimator using the Gaussian likelihood as the estimating equations and the asymptotic EL variance via [Qin & Lawless \(1994\)](#). Alternatively, we calculate the EL at (B_0, Ω_0) using original EL, EEL, AEL. We then compare the resulting EL ratio to its asymptotic χ^2 distribution. If a method does not converge, we count this as a case in which the confidence region does not cover the true parameters.

At each sample size and error distribution, we construct 1000 graphs with 6 nodes, 8 directed edges and 4 bidirected edges from the procedure described in Section 4.1. For the T distribution, we increase the degrees of freedom to 7 to ensure Proposition 1 applies. The coverage rates for 90% confidence intervals are shown in Figure 5. Based on the displayed results, regions obtained from the Gaussian MLE and GLS can only be recommended when the errors are (close to) Gaussian. The EEL method performs the best, staying close to the parametric methods in the Gaussian case and doing the best in most non-Gaussian scenarios. The sandwich method is another good choice. However, we also observe that in order to achieve nominal coverage levels very large sample sizes may be required.

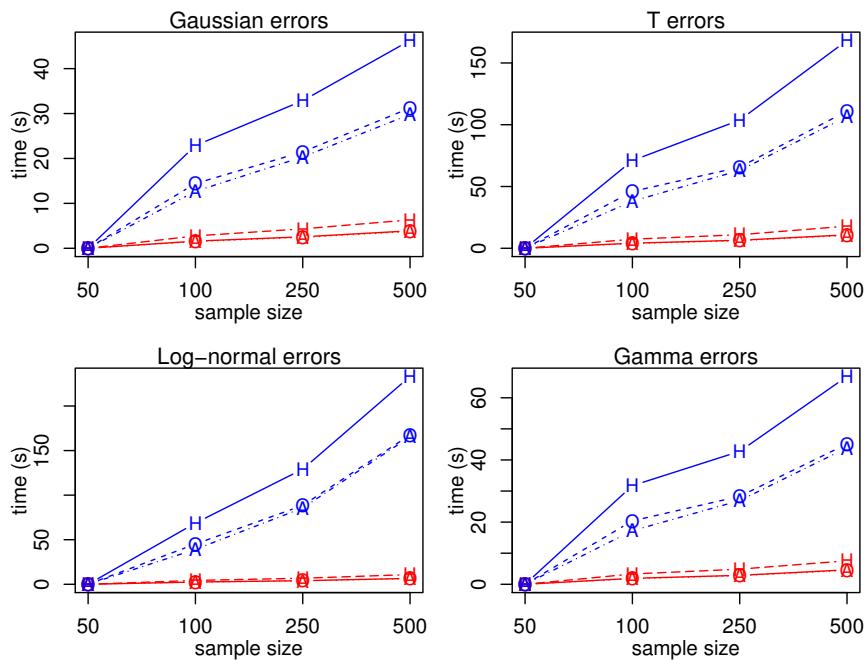


Figure 2. The average run time in seconds among the simulations in which all methods converge to a valid stationary point, plotted versus the sample size. O- original EL; A- adjusted EL; H- hybrid EL. Red points indicate the profile formulation; blue points indicate the naive formulation.

4.4. Protein Signaling Network

Sachs et al. (2005, Figure 2) present a signaling network of 11 observed molecules and 13 unobserved molecules. The black edges in Figure 6 give a plausible mixed graph representation of that network and was also considered by Drton et al. (2017). A log-transformation of the available protein expression data improves Gaussianity but leaves the distribution of some of the variables skewed and/or multimodal. We consider two separate tests; each compares the SEM sub-model corresponding to the graph of black edges against a full model which adds one of the two red edges also shown in Figure 6. Note that the added red edge from Mek \rightarrow PKA induces a directed cycle. For the log-transformed data, we perform a Gaussian as well as an empirical likelihood ratio test. For the test involving the directed edge Mek \rightarrow PKA, the Gaussian LR is .416 (p-value = .52) and the ELR is 4.379 (p-value = .04). For the test involving the bidirected edge Akt \leftrightarrow PIP2, the Gaussian LR is 15.216 (p-value $< .001$) and the ELR is .782 (p-value = .37). While we do not have a certified gold standard network, and the implicit assumption of linearity may not be appropriate for all postulated relationships, these examples present situations in which the Gaussian assumption is particularly inappropriate and may cause concerns for a practitioner.

5. Discussion

In this article, we showed that EL methods are an attractive alternative for estimation and testing of non-Gaussian linear SEMs. Our approach of profiling out the error covariance matrix Ω drastically reduces computational effort and creates a far more tractable and reliable estimation procedure. Furthermore, we showed that the use of AEL

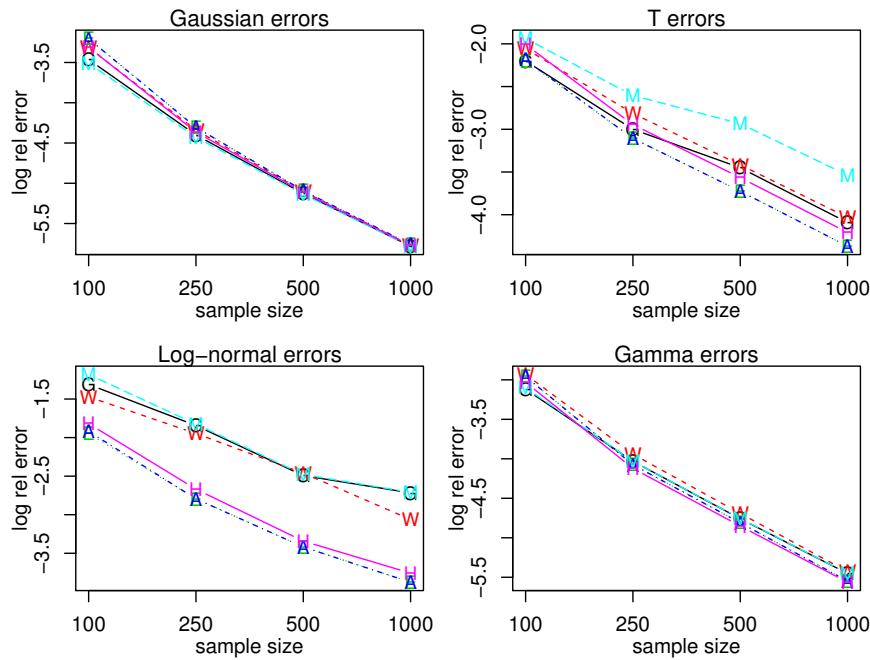


Figure 3. Log mean relative squared estimation error in Σ over 1000 simulations, plotted versus the sample size. Average is only taken on simulations in which all methods converged. A- adjusted EL; E- empirical likelihood; G- generalized least squares; H- hybrid Gauss/EL; M- Gaussian MLE; W- weighted least squares.

may further improve convergence of optimizers, particularly, when the sample size is small and the errors are skewed. EEL was seen to drastically improve the coverage rate of the joint confidence intervals.

Our EL methods are applicable under very few distributional assumptions, all the while allowing statistical inference in close to analogy to parametric modeling. When the data is non-Gaussian, the modified EL methods outperform the other methods we considered in almost all scenarios we explored. This concerns the proportion of times a valid estimate is returned, statistical efficiency, and also confidence region coverage. While there remains significant room for improvement in the design of confidence regions, we conclude that EL methods are a valuable tool for applications of linear SEMs to non-Gaussian data.

Appendix

Proof of Proposition 1

We recall our notation $\theta = (B, \Omega)$ and $\theta_0 = (B_0, \Omega_0)$. Based on the right-most expression in (9), the considered naive/direct estimating equations may be based on the function $G(y, B)$ with coordinates

$$G_v(y, B) = y_v, \quad G_{uv}(y, B) = \left(y_v - \sum_{s \in \text{pa}(v)} \beta_{vs} y_s \right) \left(y_u - \sum_{t \in \text{pa}(u)} \beta_{ut} y_t \right) - \omega_{uv},$$

for $v \in V$ and $\{u, v\} \in V \times V$, respectively.

Our claim follows from Theorem 1 of Qin & Lawless (1994) under the following conditions:

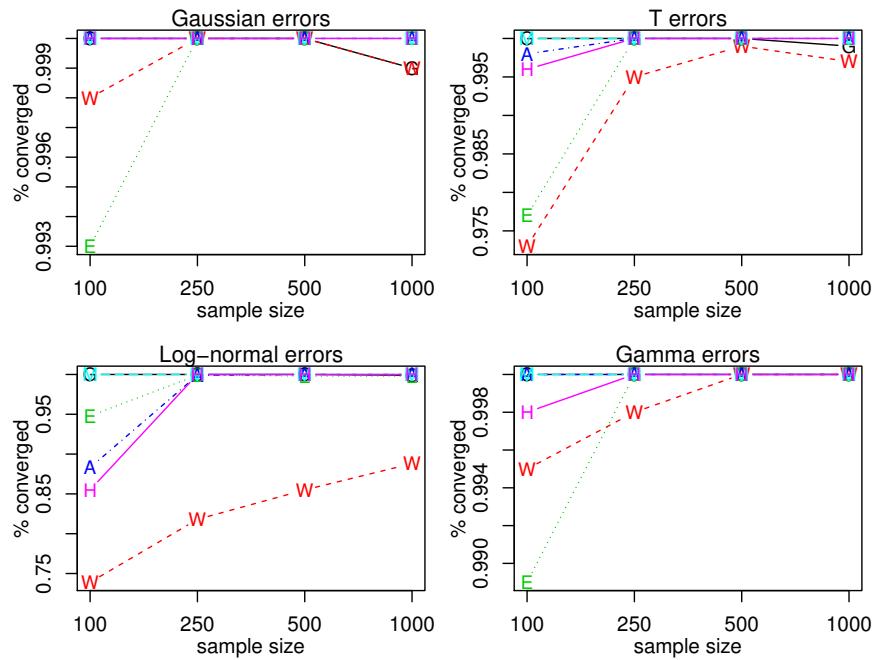


Figure 4. Proportion of times (over 1000 simulations) the method converged to a local maximum of the objective function, plotted versus the sample size. A- adjusted EL; E- empirical likelihood; G- generalized least squares; H- hybrid Gauss/EL; M- Gaussian MLE; W- weighted least squares.

- (1) $\mathbb{E}(G(Y_i, \theta_0)G(Y_i, \theta_0)^t)$ is positive definite.
- (2) In a neighborhood of the d -dimensional parameter θ_0 , the derivative $\frac{\partial G(y, \theta)}{\partial \theta}$ is continuous, and $\left\| \frac{\partial G(y, \theta)}{\partial \theta} \right\|$ and $\|G(y, \theta)\|^3$ are bounded by an integrable function $M_1(y)$.
- (3) $\mathbb{E}\left(\frac{\partial G(Y_i, \theta_0)}{\partial \theta}\right)$ has rank d .
- (4) $\partial^2 G(y, \theta)/\partial \theta \theta^T$ is continuous and $\left\| \partial^2 G(y, \theta)/\partial \theta \theta^T \right\|$ is bounded by an integrable function $M_2(y)$ in a neighborhood of the true parameter θ_0 .

Here, $\|\cdot\|$ denotes the Euclidean norm.

Noting that $G_{uv}(Y_i, B_0) = \epsilon_v \epsilon_u - \omega_{uv}$, condition (1) is an immediate consequence of assumption (b) in our proposition. Condition (3) is implied by assumption (a). With polynomial estimating equations, all derivatives in conditions (2) and (4) exist. Now, G and its first and second partial derivatives are at most quadratic functions of Y_i , which in turn is a linear function of a realization of the error vector ϵ . Local bounds on the concerned quantities are easily obtained and assumption (b) ensures their integrability. \square

6. Acknowledgments

This material is based upon work supported by the U.S. National Science Foundation under Grants No. DMS 1561814 and 1712535. We would also like to thank the Institute for Mathematical Sciences at the National

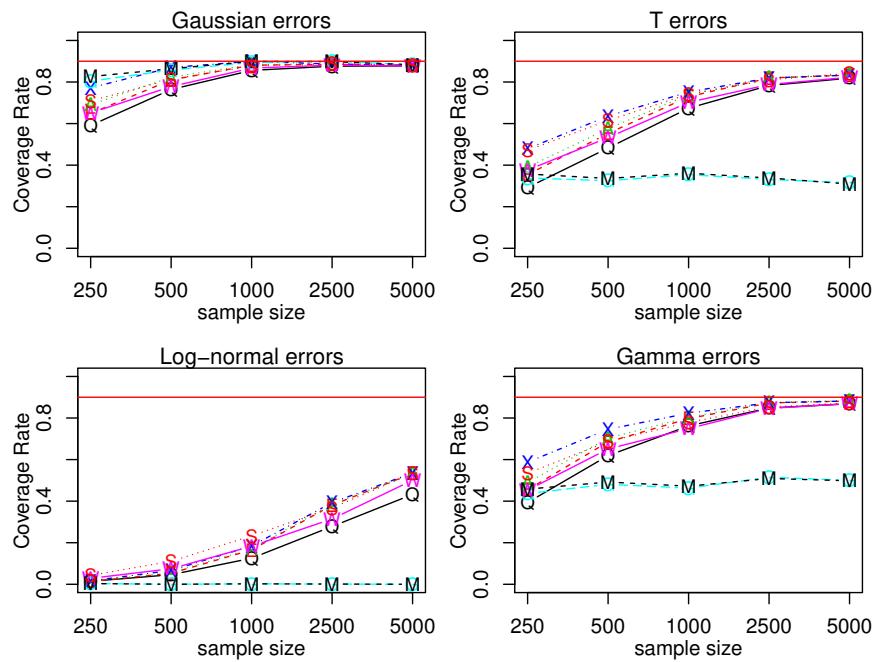


Figure 5. Coverage frequencies of joint confidence intervals. A- adjusted EL; X- extended EL; G- generalized least squares; Q- Qin and Lawless asymptotic EL variance; M- Gaussian MLE; S- sandwich estimator; W- weighted least squares; E- EL direct χ^2 calibration.

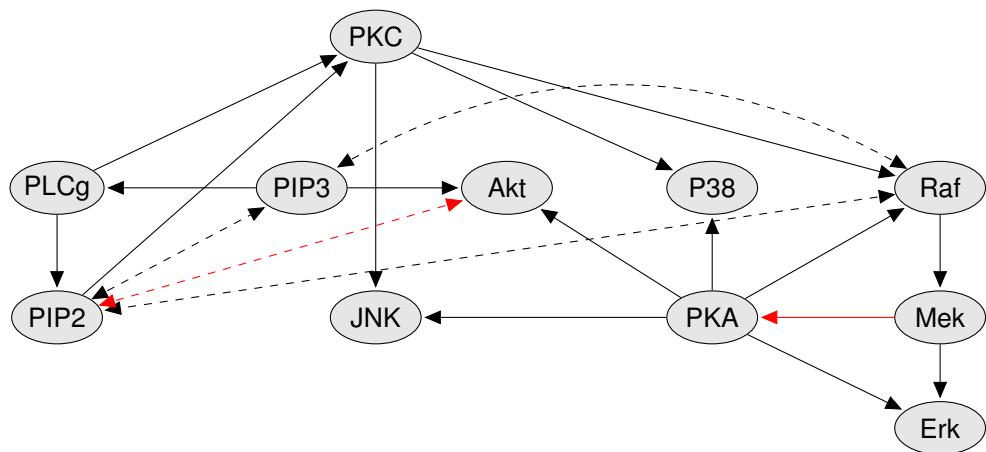


Figure 6. Plausible mixed graph for the protein-signaling network dataset. The relevant sub-model can be formed by removing the red bidirected edge between PIP2 and Akt and the red edge from MEK to PKA.

University of Singapore for travel support to the 2016 workshop: Empirical Likelihood Based Methods in Statistics.

References

Bollen, KA (1989), *Structural equations with latent variables*, Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics, John Wiley & Sons, Inc., New York, a Wiley-Interscience Publication.

Brito, C & Pearl, J (2002), 'A new identification condition for recursive models with correlated errors,' *Struct. Equ. Model.*, **9**(4), pp. 459–474.

Chaudhuri, S, Drton, M & Richardson, TS (2007), 'Estimation of a covariance matrix with zeros,' *Biometrika*, **94**(1), pp. 199–216.

Chaudhuri, S, Mondal, D & Yin, T (2017), 'Hamiltonian Monte Carlo sampling in Bayesian empirical likelihood computation,' *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, **79**(1), pp. 293–320.

Chen, B (2016), 'Identification and overidentification of linear structural equation models,' in Lee, DD, Sugiyama, M, Luxburg, UV, Guyon, I & Garnett, R (eds.), *Advances in Neural Information Processing Systems 29*, Curran Associates, Inc., pp. 1587–1595.

Chen, J, Variyath, AM & Abraham, B (2008), 'Adjusted empirical likelihood and its properties,' *J. Comput. Graph. Statist.*, **17**(2), pp. 426–443.

Colombo, D, Maathuis, MH, Kalisch, M & Richardson, TS (2012), 'Learning high-dimensional directed acyclic graphs with latent and selection variables,' *Ann. Statist.*, **40**(1), pp. 294–321.

Drton, M, Eichler, M & Richardson, TS (2009), 'Computing maximum likelihood estimates in recursive linear models with correlated errors,' *J. Mach. Learn. Res.*, **10**, pp. 2329–2348.

Drton, M, Fox, C & Wang, YS (2017), 'Computation of maximum likelihood estimates in cyclic structural equation models,' *Ann. Statist.*, to appear, arXiv:1610.03434.

Drton, M & Weihs, L (2016), 'Generic identifiability of linear structural equation models by ancestor decomposition,' *Scand. J. Stat.*, **43**(4), pp. 1035–1045.

Fox, J, Nie, Z & Byrnes, J (2017), *sem: Structural Equation Models*, r package version 3.1-9.

Foygel, R, Draisma, J & Drton, M (2012), 'Half-trek criterion for generic identifiability of linear structural equation models,' *Ann. Statist.*, **40**(3), pp. 1682–1713.

Grendár, M & Judge, G (2009), 'Empty set problem of maximum empirical likelihood methods,' *Electron. J. Stat.*, **3**, pp. 1542–1555.

Kolenikov, S & Yuan, Y (2009), 'Empirical likelihood estimation and testing in covariance structure models,' Available at <http://staskolenikov.net>.

Muthén, B & Kaplan, D (1992), 'A comparison of some methodologies for the factor analysis of non-normal Likert variables: A note on the size of the model,' *British Journal of Mathematical and Statistical Psychology*, **45**(1), pp. 19–30.

Olsson, UH, Foss, T, Troye, SV & Howell, RD (2000), 'The performance of ML, GLS, and WLS estimation in structural equation modeling under conditions of misspecification and nonnormality,' *Structural Equation Modeling*, **7**(4), pp. 557–595.

Owen, AB (1988), 'Empirical likelihood ratio confidence intervals for a single functional,' *Biometrika*, **75**(2), pp. 237–249.

Owen, AB (2001), *Empirical likelihood*, CRC press.

Pearl, J (2009), *Causality*, Cambridge University Press, Cambridge, second edn., models, reasoning, and inference.

Qin, J & Lawless, J (1994), 'Empirical likelihood and general estimating equations,' *Ann. Statist.*, **22**(1), pp. 300–325.

R Core Team (2017), *R: a language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria.

Richardson, T & Spirtes, P (2002), 'Ancestral graph Markov models,' *Ann. Statist.*, **30**(4), pp. 962–1030.

Rosseel, Y (2012), 'lavaan: An R package for structural equation modeling,' *Journal of Statistical Software*, **48**(2), pp. 1–36.

Sachs, K, Perez, O, Pe'er, D, Lauffenburger, DA & Nolan, GP (2005), 'Causal protein-signaling networks derived from multiparameter single-cell data,' *Science*, **308**(5721), pp. 523–529.

Shpitser, I & Pearl, J (2006), 'Identification of joint interventional distributions in recursive semi-Markovian causal models,' in *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pp. 1219–1226.

Spirtes, P, Glymour, C & Scheines, R (2000), *Causation, prediction, and search*, Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA, second edn., with additional material by David Heckerman, Christopher Meek, Gregory F. Cooper and Thomas Richardson, A Bradford Book.

Tsao, M & Wu, F (2014), 'Extended empirical likelihood for estimating equations,' *Biometrika*, **101**(3), pp. 703–710.

Wermuth, N (2011), 'Probability distributions with summary graph structure,' *Bernoulli*, **17**(3), pp. 845–879.

Wright, S (1921), 'Correlation and causation,' *J. Agric. Res.*, **20**, pp. 557–585.