

Confirmatory path analysis in a generalized multilevel context

BILL SHIPLEY¹

Département de Biologie, Université de Sherbrooke, Sherbrooke, Québec J1K2R1 Canada

Abstract. This paper describes how to test, and potentially falsify, a multivariate causal hypothesis involving only observed variables (i.e., a path analysis) when the data have a hierarchical or multilevel structure, when different variables are potentially defined at different levels of such a hierarchy, and when different variables have different sampling distributions. The test is a generalization of Shipley's d-sep test and can be conducted using standard statistical programs capable of fitting generalized mixed models.

Key words: causal models; d-separation; graphical models; path analysis; structural equation models.

INTRODUCTION

Hypothetical explanations of many ecological phenomena are both inherently multivariate and also implicitly or explicitly causal in nature. Even when such explanations cannot be fully tested using randomized manipulative experiments, many of the causal implications of such explanations can still be tested using confirmatory structural equations modeling (SEM; Shipley 2000a, Grace 2006). When the causal explanation does not involve unmeasured (latent) variables, this reduces to confirmatory path analysis.

However, many ecological phenomena are also inherently hierarchical or otherwise multilevel in nature; examples are when repeated measurements are taken on the same individuals, when observations are nested in different geographical areas, when individuals are nested in different species, and so on. Standard methods of testing path models based on maximum likelihood are difficult or even impossible to apply when data have such a hierarchical structure and when intercepts and path coefficients therefore potentially vary between hierarchical levels. Such problems are especially acute when different variables are measured at different hierarchical levels.

There are two alternative, but asymptotically equivalent, ways of conducting a confirmatory path analysis. Standard methods are based on a comparison between observed and predicted covariance matrices, such as those derived from maximum likelihood estimators. However, Shipley (2000a, b, 2003, 2004) derived a different, and more general, method of testing path models, based on the graph theoretic notion of “d-separation” (directional separation; Pearl and Verma 1987, Geiger et al. 1990, Geiger and Pearl 1993, Pearl 2000) and on the relationship between d-separation of a directed acyclic (i.e., causal) graph and conditional

independence claims in the probability distribution generated by such a graph. The purpose of this note is to explain how such d-sep tests of path models can be easily generalized to deal with data having a hierarchical structure; I will call these “generalized multilevel path models.” *D-sep tests of causal graphs* describes the logic and mechanics of the d-sep test. *Generalized multilevel path models* gives an ecological example of a simple generalized multilevel path model and shows how to test a model using d-sep tests combined with the mechanics of generalized mixed model regression.

D-SEP TESTS OF CAUSAL GRAPHS

A multivariate causal hypothesis consists of specifying how the variables are linked together in terms of direct and indirect causal effects. This gives rise to “box-and-arrow” diagrams familiar to ecologists, showing how causal effects should flow through the system. Fig. 1 shows a simple path model involving four measured variables and two sets of mutually independent residual causes (ϵ) of variables X_3 and X_4 . In graph theory, such box-and-arrow diagrams are called directed graphs and, when there are no feedback relationships ($A \rightarrow B \rightarrow C \rightarrow A$), they are called directed acyclic graphs (DAGs). Testing the causal structure of a multivariate causal hypothesis using observational data is therefore equivalent to testing the hypothetical cause–effect structure of a directed graph.

A classical way of testing such a causal hypothesis is through a series of experimental manipulations, in which some variables are experimentally fixed to constant values, thus preventing them from changing in response to their normal causes. One would then deduce, given the causal hypothesis, which variables must be dependent or independent in their natural state and how these patterns of dependence and independence must change following the experimental control of other sets of variables. For example, the variables X_1 and X_2 in Fig. 1 are causally independent if we hold constant none of the other variables, i.e., $\{\emptyset\}$; if we hold constant only variable X_3 , i.e., $\{3\}$; if we hold constant only variable X_4 , i.e., $\{4\}$; and if we simultaneously hold constant

Manuscript received 2 June 2008; revised 14 August 2008; accepted 20 August 2008. Corresponding Editor: A. Ellison.

¹ E-mail: Bill.Shipley@USherbrooke.ca




FIG. 1. A multivariate causal hypothesis expressed as a directed acyclic graph involving four observed variables (X_1 to X_4) and unobserved causes (ϵ_3, ϵ_4) generating the residual variances X_3 and X_4 .

both variables X_3 and X_4 , i.e., $\{3,4\}$. Symbolically, we would write down the following independence claims: $(1, 2) | \{\emptyset\}$, $(1, 2) | \{3\}$, $(1, 2) | \{4\}$, $(1, 2) | \{3, 4\}$. In Fig. 1, we also see that variable X_4 is causally dependent on the behavior of variable X_1 if no physical controls are imposed because X_1 causes X_3 which then causes X_4 ; symbolically, we would write down the following dependence claim: $(1,4) | \{\emptyset\}$. However, if we were to physically prevent X_3 from responding to changes in X_1 (i.e., hold X_3 constant) then the causal dependence of X_4 on X_1 would be removed; X_1 and X_4 would become independent conditional on X_3 remaining constant. Symbolically, we would write down the following independence claim: $(1, 4) | \{3\}$. It is possible to write down all claims of dependence and independence that are logically implied by the hypothesized causal graph given in Fig. 1 and these are listed in Table 1 under the heading “Experimental control.”

Comparing the agreement between predicted patterns of dependence and independence between all pairs of variables, following experimental control of all possible sets of other variables, and the observed patterns of dependence and independence following such experimental controls, allows one to falsify the causal hypothesis. If any of the predicted patterns are contradicted by the observed results then the causal hypothesis is wrong. This is simply the logic of a controlled experiment. The total number (N) of such predicted patterns of dependence and independence between pairs of variables (a “claim”), given a model with V variables is simply the product of the number of unique pairs of variables multiplied by the number of possible unique sets of remaining control variables:

$$N = \left[\frac{V!}{2!(V-2)!} \right] \left[\sum_{n=0}^{V-2} \frac{(V-2)!}{n!(V-2-n)!} \right]. \quad (1)$$

Table 1 lists all $6 \times 4 = 24$ possible experimental controls given the causal hypothesis shown in Fig. 1 and their predicted dependence or independence.

A d-sep (directional separation) test of a causal hypothesis follows the same logic as in the case of experimental controls except that experimental control is replaced with statistical control. Determining the dependence or independence of any two variables (X_1, X_2) after statistically holding constant some other variables $\{X_3, \dots, X_n\}$ consists of determining the dependence or independence of the residuals of X_1 from the expected value of X_1 given the values of $\{X_3, \dots, X_n\}$ and the residuals of X_2 from the expected value of X_2 given the values of $\{X_3, \dots, X_n\}$. Tests of independence based on Pearson correlations or partial correlations or on the slopes of regressions or multiple regressions are particular examples of such statistical controls. Since statistical and experimental controls do not always give the same predictions of dependence or independence (Pearl 2000, Shipley 2000a), it is necessary to know how a causal hypothesis involving V variables (a directed acyclic graph) is translated into the N claims of statistical, as opposed to experimental, dependence or independence. This translation is given by a manipulation of the graph called “d-separation.” Pearl (1988: theorem 10) proved that if two variables (X, Y) are d-separated given a conditioning set Z of other variables in a directed acyclic graph then they must also be conditionally independent in any probability distribution that is generated by such a graph; Z can also be the empty $\{\emptyset\}$ set, meaning that no conditioning variables are used. Additionally, if the two variables are not d-separated in the graph, then they cannot be condition-

TABLE 1. Patterns of dependence and independence predicted by the causal graph shown in Fig. 1 when the controls are determined by either experimentally controlling (i.e., holding constant) a variable or statistically controlling (i.e., conditioning on) a variable.

Predicted independences	Predicted dependences
Experimental control	
$(1, 2) \{\emptyset\}$, $(1, 2) \{3\}$, $(1, 2) \{4\}$, $(1, 2) \{3, 4\}$, $(1, 4) \{3\}$, $(1, 4) \{2, 3\}$, $(2, 4) \{3\}$, $(2, 4) \{1, 3\}$	$(1, 3) \{\emptyset\}$, $(1, 3) \{2\}$, $(1, 3) \{4\}$, $(1, 3) \{2, 4\}$, $(2, 3) \{\emptyset\}$, $(2, 3) \{1\}$, $(2, 3) \{4\}$, $(2, 3) \{1, 4\}$, $(1, 4) \{\emptyset\}$, $(1, 4) \{2\}$, $(3, 4) \{\emptyset\}$, $(3, 4) \{1\}$, $(3, 4) \{2\}$, $(3, 4) \{1, 2\}$, $(2, 4) \{\emptyset\}$, $(2, 4) \{1\}$
Statistical control	
$(1, 2) \{\emptyset\}$, $(1, 4) \{3\}$, $(1, 4) \{2, 3\}$, $(2, 4) \{3\}$, $(2, 4) \{1, 3\}$	$(1, 3) \{\emptyset\}$, $(1, 3) \{2\}$, $(1, 3) \{4\}$, $(1, 3) \{2, 4\}$, $(2, 3) \{\emptyset\}$, $(2, 3) \{1\}$, $(2, 3) \{4\}$, $(2, 3) \{1, 4\}$, $(1, 4) \{\emptyset\}$, $(1, 4) \{2\}$, $(3, 4) \{\emptyset\}$, $(3, 4) \{1\}$, $(3, 4) \{2\}$, $(3, 4) \{1, 2\}$, $(2, 4) \{\emptyset\}$, $(2, 4) \{1\}$, $(1, 2) \{3\}$, $(1, 2) \{4\}$, $(1, 2) \{3, 4\}$

Notes: The notation “ $(X, Y) | \{A, B\}$ ” means “the pair of variables (X, Y) when controlling the values of variables A and B ” and “ $\{\emptyset\}$ ” means holding nothing constant (a null set). Boldface type indicates differences in claims of dependence or independence between experimental and statistical control.




Fig. 2. A hypothetical causal process involving six observed variables (independent latent causes are not shown for simplicity). Latitude and year generate the number of degree-days at each site. Degree-days then cause the date of bud burst of a tree species. The date of bud burst causes the amount of diameter growth, and diameter growth determines the survival in the subsequent winter.

ally independent in the resulting probability distribution. Therefore, applying d-separation to the graph for each of the N independence claims gives us the complete description of the statistical patterns of conditional dependence and independence that must be true of any data generated by such a causal process; this applies independently of the nature of the variables, the functional form of the relationships between the variables (i.e., linear or nonlinear), or the type of multivariate probability distribution that is generated by the causal process. Table 1 also lists the N statistical independence claims associated with Fig. 1 and highlights in bold type those claims that differ between experimental and statistical control.

In order to test the full causal hypothesis that is represented by a causal graph, we must perform a simultaneous test of all N independence claims. Shipley (2000a, b, 2003) described such an inferential test for all N claims based on a particular “basis set” (\mathbf{B}_U) of the N independence claims that has the three following properties: (1) all independence claims not in this basis (and all dependence claims) set are logical consequences of some combination of those within it, (2) no independence claim within this basis set can be derived from some combination of the others within it, and (3) if data are generated according to the causal graph then the null probabilities of each independence claim are mutually independent. The first two properties mean that a test of the \mathbf{B}_U basis set is also a test of the entire set of patterns of statistical dependence and independence implied by the causal process. The third property means that the following statistic, calculated on the independence claims of \mathbf{B}_U , follows a chi-square distribution with $2k$ degrees of freedom if the observed data are generated according to the hypothesized causal structure of the causal graph, where k is the number of independence claims in the basis set \mathbf{B}_U and p_i is the null probability of the independence test associated with the i th independence claim:

$$C = -2 \sum_{i=1}^k \ln(p_i). \quad (2)$$

The basis set \mathbf{B}_U consists of each pair of variables (X , Y) in the graph that do not have an arrow between them (i.e., one is not a direct cause of the other) and the conditioning set $\{\mathbf{Z}\}$ for each such pair contains all variables that are direct causes of either X or Y (their “causal parents”). Thus, the basis set for Fig. 1 contains three independence claims: $\mathbf{B}_U = \{(X_1, X_2) | \{\emptyset\}, (X_1, X_4) | \{X_3\}, (X_2, X_4) | \{X_3\}\}$ that, together, imply all other claims of dependence or independence made by Fig. 1.

The notation “ $(X_i, X_j) | \{\mathbf{Z}\}$ ” means that variables X_i and X_j are independent after conditioning on the variables in the set $\{\mathbf{Z}\}$. In order to calculate the null probability (p_i) for each of the k independence claims in the \mathbf{B}_U basis set, one must use a test of independence that is appropriate for each claim, based on the distributional properties of the variables in question. In this sense, a d-sep test is very general. In the specific context of mixed-model path analysis, this means that one must use tests of independence that are appropriate for the assumptions of the mixed model.

The steps in performing a d-sep test are, therefore:

- 1) Express the hypothesized causal relationships between the variables in the form of a directed acyclic graph.
- 2) List each of the k pairs of variables in the graph that do not have an arrow between them.
- 3) For each of the k pairs of variables (X_i, X_j), list the set of other variables, $\{\mathbf{Z}\}$ in the graph that are direct causes of either X_i or X_j . The pair of variables (X_i, X_j) along with its conditioning set $\{\mathbf{Z}\}$ define an independence claim, $(X_i, X_j) | \{\mathbf{Z}\}$, and the full set of the k independence claims defines the basis set \mathbf{B}_U .
- 4) For each element in this basis set, obtain the probability, p_k , that the pair (X_i, X_j) is statistically independent conditional on the variables \mathbf{Z} .
- 5) Combine the k probabilities using Eq. 2 and compare the resulting C value to a chi-squared distribution with $2k$ degrees of freedom. Reject the causal model if the C value is unlikely to have occurred by chance (i.e., below the chosen significance level).

GENERALIZED MULTILEVEL PATH MODELS

Consider a hypothetical study beginning in 1970 in which 20 sites are chosen differing in latitude (X_1). Five individual trees of a particular species are chosen within each site. Each tree is followed every second year until 2006 or until it dies (thus, repeated measures). At each site, in each sample year, and for each living individual you measure the cumulative degree days until bud break (X_2), the Julian date (day of year) of bud break (X_3), the increase in stem diameter per tree (X_4), and a binary variable indicating survival (1) or death (0) during the subsequent growing season (X_5). Fig. 2 shows a hypothetical causal structure involving these five measured variables. Such a path model would be very difficult to test using standard structural unilevel equations models. First, there are three hierarchical levels (between sites, between individuals within sites, between years within individuals within sites) thus making observations nonindependent. Individuals grow-

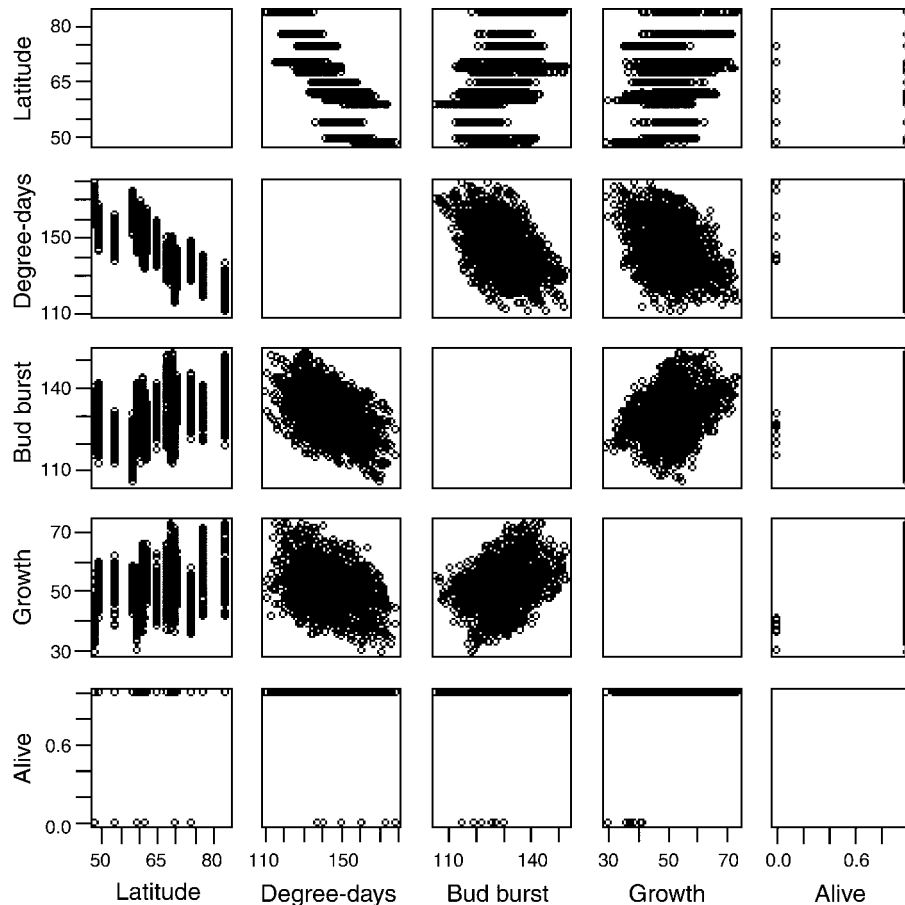


FIG. 3. A simulated data set, generated according to the causal structure shown in Fig. 2, involving 20 sites and five trees per site. Each tree is measured every second year from 1970 until 2006 or until it dies. Variables are latitude (degrees), degree days ($^{\circ}\text{C}$), bud burst (Julian days), growth (cm/yr), and alive (binary 0/1).

ing in the same site will tend to respond similarly to many site-specific characteristics besides those explicit in the model. Similarly, the repeated measures of the same individual will tend to respond similarly due to the many characteristics specific to that individual besides those explicit in the model. Second, because of this nested structure, the strength of the path coefficients could potentially vary between individuals and sites. Third, the

different nature of the variables (binary for survival, continuous for the other variables) complicates testing the conditional independencies using correlations or covariances. Finally, different variables vary at different hierarchical levels; for example, “latitude” varies only between sites, “degree days” varies both between sites and between years but not between trees in the same site and year, while “growth” and “survival” varies between

TABLE 2. The \mathbf{B}_U basis set of d-separation (d-sep; directional separation) claims implied by Fig. 3 (the correct causal structure).

D-sep claim of independence	Mixed model†	Variable whose partial regression slope should be zero	Null probability (distribution)
$(X_1, X_3) \mid \{X_2\}$	$X_3 \sim X_2 + X_1 + (1 \mid \text{site}) + (1 \mid \text{tree})$	X_1	0.9373 (normal)
$(X_1, X_4) \mid \{X_3\}$	$X_4 \sim X_3 + X_1 + (1 \mid \text{site}) + (1 \mid \text{tree})$	X_1	0.3837 (normal)
$(X_1, X_5) \mid \{X_4\}$	$X_5 \sim X_4 + X_1 + (1 \mid \text{site}) + (1 \mid \text{tree})$	X_1	0.2800 (binomial)
$(X_2, X_4) \mid \{X_1, X_3\}$	$X_4 \sim X_3 + X_1 + X_2 + (1 \mid \text{site}) + (1 \mid \text{tree})$	X_2	0.9839 (normal)
$(X_2, X_5) \mid \{X_1, X_4\}$	$X_5 \sim X_4 + X_1 + X_2 + (1 \mid \text{site}) + (1 \mid \text{tree})$	X_2	0.9839 (binomial)
$(X_3, X_5) \mid \{X_2, X_4\}$	$X_5 \sim X_4 + X_2 + X_3 + (1 \mid \text{site}) + (1 \mid \text{tree})$	X_3	0.1890 (binomial)

Notes: Key to variables: X_1 = latitude, X_2 = degree days, X_3 = Julian date of bud burst, X_4 = diameter growth following bud burst, and X_5 = survival (1) or death (0) during the following winter. The grouping variables are site, trees within sites, and repeated measures per tree.

† The associated mixed model regression for each d-sep claim using the lmer function in R to test the independence claims; the grouping variables are site, trees within sites, and repeated measures per tree.




FIG. 4. An alternative causal explanation from that shown in Fig. 3.

sites, between years and between trees. However, none of these complications affect the logic of the d-sep test since one can test the different predicted conditional independencies using different tests of independence that are appropriate for the data.

One can test each hypothesized conditional independency separately using a generalized mixed model, obtain a null probability separately for each, and then combine them using Eq. 2. This assumes that the appropriate functional form linking the variables can be expressed using a generalized mixed model. Given any d-separation claim, for example $(X, Y) | \{A, B\}$, one would regress either X or Y on the set of conditioning variables (here A and B) plus the other variable in the pair, specifying the appropriate error distribution for the dependent variable (i.e., normal, Poisson, binomial, and so on) and also specifying the correct hierarchical structure of the data. Many statistical programs can fit mixed (or multilevel) models; here I use the lme or lmer functions in the nlme and lme4 packages of R (R Development Core Team 2008). One then obtains the probability that the partial slope of the dependent variable of the pair (either X or Y) is zero in the statistical population given the conditioning variables (here, A and B) since, if this is true, then X and Y are conditionally independent.

Fig. 3 shows a scatter plot matrix of data simulated from Fig. 2. Table 2 lists the \mathbf{B}_U basis set of Fig. 2, the associated mixed model regressions of each element of \mathbf{B}_U using the lmer function, the relevant partial regression slope whose value is predicted to be zero in the statistical population, and the null probability

assuming the true partial regression slope is zero; note that the assumed sampling distribution will change depending on whether the dependent variable is normally distributed (all variables except for survival) or binomially distributed (survival). The Supplement gives the code to do this analysis. The C statistic (Eq. 2) is 9.53 with $k = 12$ degrees of freedom. The probability of observing this value by chance if the data were actually generated by the causal process (which they were) is 0.66, thus correctly preventing us from rejecting the model at a significance level of 0.05.

Fig. 4 shows an alternative (incorrect) causal hypothesis concerning these same data. The only difference is that in this incorrect hypothesis the date of budburst is not a cause of growth; rather, it is spuriously correlated with growth and survival because it shares a common cause with them (i.e., degree days). Table 3 lists the \mathbf{B}_U basis set of Fig. 4, the associated mixed-model regressions of each element of \mathbf{B}_U , the relevant partial regression slope and its null probability assuming the true partial regression slope is zero. Note that two of the d-sep claims in this incorrect model differ from those implied by the correct model. The C statistic for this incorrect model is 31.75 with $k = 12$ degrees of freedom. The probability of observing this value by chance if the data were actually generated by the causal process (they were not) is 0.002, thus correctly forcing us to reject the model at a significance level of 0.05.

Once a causal structure is obtained that is consistent both with the predicted patterns of direct and indirect statistical independence and with any other background information about the causal process, then the path

TABLE 3. The \mathbf{B}_U basis set of d-separation (d-sep) claims implied by Fig. 4 (the incorrect causal structure).

D-sep claim of independence	Mixed model†	Variable whose partial regression slope should be zero	Null probability (distribution)
$(X_1, X_3) \{X_2\}$	$X_3 \sim X_2 + X_1 + (1 \text{site}) + (1 \text{tree})$	X_1	0.9373 (normal)
$(X_1, X_4) \{X_2\}$	$X_4 \sim X_2 + X_1 + (1 \text{site}) + (1 \text{tree})$	X_1	0.4026 (normal)
$(X_1, X_5) \{X_4\}$	$X_5 \sim X_4 + X_1 + (1 \text{site}) + (1 \text{tree})$	X_1	0.2800 (binomial)
$(X_3, X_4) \{X_2\}$	$X_4 \sim X_2 + X_3 + (1 \text{site}) + (1 \text{tree})$	X_3	0.000014 (normal)
$(X_2, X_5) \{X_1, X_4\}$	$X_5 \sim X_4 + X_1 + X_2 + (1 \text{site}) + (1 \text{tree})$	X_2	0.9839 (binomial)

Notes: Key to variables: X_1 = latitude, X_2 = degree days, X_3 = Julian date of bud burst, X_4 = diameter growth following bud burst, and X_5 = survival (1) or death (0) during the following winter. The grouping variables are site, trees within sites, and repeated measures per tree. The d-sep claims in bold are those that differ from Table 2.

† The associated mixed model regression for each d-sep claim using the lmer function in R to test the independence claims; the grouping variables are site, trees within sites, and repeated measures per tree.

coefficients can be obtained by fitting the series of mixed models that follow from the structure. This series of models is not the same as the one required in testing the model because the path coefficients are obtained by regressing each variable only on each of its direct causes.

The extension of the d-sep test to a mixed-model context, including complications beyond those illustrated here, is straightforward and can be implemented using standard statistical programs. For instance, models in which path coefficients vary between groups would have randomly varying slopes in the mixed model and some models with correlated error structures can also be accommodated (Shipley 2003). An empirical example is given in Thomas et al. (2007). Finally it is also possible to apply the method described in this paper to multilevel Bayesian (MCMC) models (McCarthy 2007) simply by testing independence of the (partial) slopes by determining the 95% credible intervals.

ACKNOWLEDGMENTS

This study was financially supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada. Partick Bergeron and Patrice Bourgault commented on an earlier version of this paper.

LITERATURE CITED

- Geiger, D., and J. Pearl. 1993. Logical and algorithmic properties of conditional independence and graphical models. *Annals of Statistics* 21:2001–2021.
- Geiger, D., T. Verma, and J. Pearl. 1990. Identifying independence in Bayesian networks. *Networks* 20:507–534.
- Grace, J. B. 2006. *Structural equation modeling and natural systems*. Cambridge University Press, Cambridge, UK.
- McCarthy, M. A. 2007. *Bayesian methods for ecology*. Cambridge University Press, Cambridge, UK.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Mateo, California, USA.
- Pearl, J. 2000. *Causality*. Cambridge University Press, Cambridge, UK.
- Pearl, J., and T. Verma. 1987. The logic of representing dependencies by directed graphs. Pages 374–379 in K. Forbus and H. E. Shrobe, editors. *Proceedings AAAI-87: Sixth National Conference on Artificial Intelligence*, 13–17 July 1987. MIT Press, Cambridge, Massachusetts, USA.
- R Development Core Team. 2008. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. (www.r-project.org)
- Shipley, B. 2000a. *Cause and correlation in biology: a user's guide to path analysis, structural equations, and causal inference*. Oxford University Press, Oxford, UK.
- Shipley, B. 2000b. A new inferential test for path models based on directed acyclic graphs. *Structural Equation Modeling* 7: 206–218.
- Shipley, B. 2003. Testing recursive path models with correlated errors using d-separation. *Structural Equation Modeling* 10: 214–221.
- Shipley, B. 2004. Analysing the allometry of multiple interacting traits. *Perspectives in Plant Ecology, Evolution and Systematics* 6:235–241.
- Thomas, D. W., B. Shipley, J. Blondel, P. Perret, A. Simon, and M. M. Lambrechts. 2007. Common paths link food abundance and ectoparasite loads to physiological performance and recruitment in nestling Blue Tits. *Functional Ecology* 21:947–955.

SUPPLEMENT

A text file and detailed instructions for implementing the statistical test using mixed model regressions in R (*Ecological Archives* E090-028-S1).