TRAIT-DEPENDENT EXTINCTION LEADS TO GREATER EXPECTED BIODIVERSITY LOSS[*]

BEÁTA FALLER[†] AND MIKE STEEL[‡]

Abstract. We use a classical combinatorial inequality to establish a Markov inequality for multivariate binary Markov processes on trees. We then apply this result, alongside the Fortuin–Kasteleyn–Ginibre (FKG) inequality, to compare the expected loss of biodiversity under two models of species extinction. One of these models is the generalized version of an earlier model in which extinction is influenced by some trait that can be classified into two states and which evolves on a tree according to a Markov process. Since more than one trait can affect the rates of species extinction, it is reasonable to allow, in the generalized model, k binary states that influence extinction rates. We compare this model to one that has matching marginal extinction probabilities for each species but for which the species extinction events are stochastically independent.

1. Introduction. The concept of a "Markov process on a tree" generalizes the notion of a Markov chain and has been extensively studied in physics, information theory, and evolutionary biology. A Markov process on a tree can also be regarded as a special type of Markov random field on a graph (i.e., a "graphical model") in which the graph is a directed tree, and in evolutionary biology these models are used to model the stochastic evolution of traits on a phylogenetic tree [5, 15]. In this work, we first apply a classic lattice-theoretic inequality (a form of the "four functions theorem" of Ahlswede and Daykin) to establish a generic Markov inequality for multivariate Markov processes that consist of k independent (but not necessarily identical) two-state Markov processes on a tree. Briefly, this inequality states that the probability p of observing given binary sequences at the tips of the tree is log-supermodular (i.e., p satisfies an inequality of the form $p(\mathbf{y}) \cdot p(\mathbf{z}) \leq p(\mathbf{y} \vee \mathbf{z}) \cdot p(\mathbf{y} \wedge \mathbf{z})$, where $\mathbf{y}$ and $\mathbf{z}$ refer to assignments of binary (0,1) sequences to the tips of the tree, and $\vee$ and $\wedge$ refer to componentwise maximum and minimum operations, respectively).

The inequality has been specifically formulated for the purpose of comparing a new species-extinction model with existing models in conservation biology. This new model is the generalized version of a state-dependent species extinction model [18], in which the extinction risk of a species is associated with an underlying binary state that evolves on an evolutionary tree. In the more general setting, extinction is influenced by k independently evolved binary traits rather than only one, giving a more realistic model. We will refer to this model as a "trait-dependent field-of-bullets" model or, more briefly, a "t-FOB model." Here the phrase "field-of-bullets" is a generic term in biology for random extinction of species, the simplest model for which is one in

which all species become extinct independently and with the same probability (see, e.g., [13, 14]); a minor extension of this allows extinction probabilities to vary across species, while still treating extinction events independently—we refer to this minor extension as a "generalized field-of-bullets" model (or, more briefly, g-FOB), and its stochastic properties were recently studied in [4].

We compare the expected loss and the variance of "phylogenetic diversity" under the t-FOB model to the corresponding values of a g-FOB model in which each species has the same extinction probability as in the trait-dependent model (we refer to the latter model as an "associated g-FOB model"). Notice that, in the t-FOB model, species extinctions are not independent events, since they depend on the traits at those species, and the traits associated with different leaves of the tree are correlated because they have been generated by some Markovian process on the tree linking those leaves, while in the g-FOB model the traits are treated as independent random variables.

Using the log-supermodularity result and a classical correlation inequality on lattices (due to Fortuin, Kasteleyn, and Ginibre) we show that the expected loss of phylogenetic diversity under a t-FOB model is always greater than or equal to that of its associated g-FOB model. In other words, when extinction events reflect the evolutionary history of binary traits, the expected loss of phylogenetic diversity is always at least that predicted under a model with independent extinction events. This generalizes the result presented in section 3 of [18] and suggests that simple models that treat species extinctions independently may systematically underestimate the loss of phylogenetic diversity.

Given this inequality between the expected future phylogenetic diversity under these two models, we might expect a similar inequality to apply for the variance. However, we show that there is no similar relationship between the variances corresponding to the two models. There are examples for which the variance of future phylogenetic diversity under an independent extinction scenario can be either smaller or greater than the variance under the model in which extinction events are influenced by k characteristics, even for $k = 1$.

In the next section, we define the multivariate Markov processes under scrutiny and then state and prove the Markov inequality. To demonstrate the phylogenetic application, section 3 presents the inequality between the expected loss of phylogenetic diversity and our findings concerning the variance of future phylogenetic diversity.

2. The Markov inequality. Let T be a rooted tree with root vertex ρ and with leaf set X. Consider k independent, nonidentical two-state Markov processes on T, each of which has the state space $\{0, 1\}$ (for a formal definition of Markov processes on trees, see, for example, [2, 15, 17]). For each vertex v of T and for $j = 1, \ldots, k$, let $\xi_j(v)$ denote the random state that v is assigned in the jth Markov process. Furthermore, for $j = 1, \ldots, k$ and for $i \in \{0, 1\}$, let $\pi_i^{(j)}$ be the probability that $\xi_j(\rho) = i$. Viewing the edges of T as arcs directed away from the root, let $P^{(j)}(r, s)$ be the transition matrix assigned to arc (r, s) in the jth process. The il-entry $P^{(j)}(r, s)_{il}$ of this 2×2 matrix is, by definition, the conditional probability that $\xi_j(s) = l$ given that $\xi_j(r) = i$. For each j, having specified the probabilities $\pi_i^{(j)}$ and the transition matrices $P^{(j)}(r, s)$, $i \in \{0, 1\}, (r, s) \in A_T$ (the arc set of T), the jth Markov process on T is uniquely defined [2, 15, 17].

We now combine these k Markov processes into a vector (having jth coordinate ξ_j) to provide a multivariate Markov process on T with state space $\{0, 1\}^k$. In this process, each vertex v of T is assigned state $\boldsymbol{\xi}(v) = (\xi_1(v), \ldots, \xi_k(v))$. Let $\mathbf{i} = (i_1, \ldots, i_k) \in$

$\{0,1\}^k$, and let $\pi_{\mathbf{i}}$ be the probability that $\boldsymbol{\xi}(\rho) = \mathbf{i}$. Then, by the independence of the k processes, we get $\pi_{\mathbf{i}} = \prod_{j=1}^{k} \pi_{i_j}^{(j)}$. Similarly, for the transition matrix $P(r,s)$ corresponding to arc (r,s) in the multivariate process, the entry $P(r,s)_{\mathbf{il}}$ in row $\mathbf{i}$ and column $\mathbf{l}$ (for $\mathbf{i} = (i_1,\ldots,i_k), \mathbf{l} = (l_1,\ldots,l_k) \in \{0,1\}^k$) becomes $\prod_{j=1}^{k} P^{(j)}(r,s)_{i_j l_j}$. This is the conditional probability that $\boldsymbol{\xi}(s) = \mathbf{l}$ given that $\boldsymbol{\xi}(r) = \mathbf{i}$. With these, the multivariate Markov process is uniquely defined.

We will assume throughout that all the π values are strictly positive and that $\det P^{(j)}(r,s) \geq 0$ for each arc (r,s) and for each j. Note that this implies that $\det P(r,s) \geq 0$. Namely, it can be seen that $P(r,s)$ is the *Kronecker product* of the k matrices $P^{(j)}(r,s)$, and so $\det P(r,s) = (\det P^{(1)}(r,s) \times \cdots \times \det P^{(k)}(r,s))^{2^{k-1}}$ (see [9] for the definition and properties of the Kronecker product). However, we are neither assuming that any of the k processes are identical, nor that within any of them, the arcs are assigned the same transition matrix.

Consider now a realization $\mathbf{U} = (U_1,\ldots,U_k)$ of $\boldsymbol{\xi} = (\xi_1,\ldots,\xi_k)$. Note that $\mathbf{U}$ is a function from V into the set $\{0,1\}^k$ of character states. Let $P(\mathbf{U})$ denote the probability that $\boldsymbol{\xi} = \mathbf{U}$, that is, the probability that for each $v \in V$, v is assigned $\mathbf{U}(v)$. For $j = 1,\ldots,k$, let $\delta_j(\mathbf{U},v) = 0$ if the jth coordinate $U_j(v)$ of $\mathbf{U}(v)$ is 0, and let $\delta_j(\mathbf{U},v) = 1$ if $U_j(v) = 1$. Also, let $\delta(\mathbf{U},v)$ denote the state that v is assigned in $\mathbf{U}$. Now we are able to express $P(\mathbf{U})$ in terms of the transition matrices and the π values of the multivariate process, using the Markov property (we follow [15]). We have

$$P(\mathbf{U}) = \pi_{\delta(\mathbf{U},\rho)} \cdot \prod_{(r,s)\in A_T} P(r,s)_{\delta(\mathbf{U},r)\delta(\mathbf{U},s)},$$

which, by the independence of the k two-state processes, gives

$$(1) \qquad P(\mathbf{U}) = \prod_{j=1}^{k} \pi_{\delta_j(\mathbf{U},\rho)}^{(j)} \cdot \prod_{(r,s)\in A_T} \prod_{j=1}^{k} P^{(j)}(r,s)_{\delta_j(\mathbf{U},r)\delta_j(\mathbf{U},s)}$$

$$= \prod_{j=1}^{k} \left(\pi_{\delta_j(\mathbf{U},\rho)}^{(j)} \prod_{(r,s)\in A_T} P^{(j)}(r,s)_{\delta_j(\mathbf{U},r)\delta_j(\mathbf{U},s)} \right).$$

Recall that a *lattice* $\mathscr{L}$ is a partially ordered set in which any two elements $a, b \in \mathscr{L}$ have a unique least upper bound $a \vee b$, called their *join*, and a unique greatest lower bound $a \wedge b$, which is their *meet*. A lattice is *distributive* if $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$ for all $a, b, c \in \mathscr{L}$ or, equivalently, $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$ for all $a, b, c \in \mathscr{L}$.

Let $\mathscr{L}_V$ be the set of all possible realizations of $\boldsymbol{\xi}$. Let $\mathbf{Y}, \mathbf{Z} \in \mathscr{L}_V$, and let $\leq$ be the partial order over $\mathscr{L}_V$ in which $\mathbf{Y} \leq \mathbf{Z}$ whenever $Y_j(v) \leq Z_j(v)$ for each vertex $v \in V$ and for each $j = 1,\ldots,k$, and in which $\mathbf{Y}$ and $\mathbf{Z}$ are incomparable otherwise. Clearly, any two elements $\mathbf{Y}$ and $\mathbf{Z}$ of the partially ordered set $(\mathscr{L}_V, \leq)$ have a join $\mathbf{Y} \vee \mathbf{Z}$ and a meet $\mathbf{Y} \wedge \mathbf{Z}$. These are the realizations of $\boldsymbol{\xi}$ that, to each vertex $v \in V$, assign state $(\max\{Y_1(v), Z_1(v)\},\ldots,\max\{Y_k(v), Z_k(v)\})$ and state $(\min\{Y_1(v), Z_1(v)\},\ldots,\min\{Y_k(v), Z_k(v)\})$, respectively. It follows that $(\mathscr{L}_V, \vee, \wedge)$ is a lattice on $\mathscr{L}_V$. It is easy to see that this lattice is distributive.

Recall that X denotes the leaf set of T, and fix a nonempty subset W of X. For each function $\mathbf{U}$ in $\mathscr{L}_V$, define $\mathbf{u} = (u_1,\ldots,u_k)$ to be the restriction of $\mathbf{U}$ to W; that is, $\mathbf{u} = \mathbf{U}|_W$. With this we have $\mathbf{u}(v) = \mathbf{U}(v)$ for each leaf v in W. Since $\mathbf{u}$ is a function from the nonempty subset W of X into a set of character states, it is also

called a *character on X* [15]. Let $\mathscr{L}_W$ be the set that contains, for each $\mathbf{U} \in \mathscr{L}_V$, the restricted function $\mathbf{u} = \mathbf{U}|_W$. Let $\mathbf{y}, \mathbf{z} \in \mathscr{L}_W$, and let $\leq$ be the partial order over $\mathscr{L}_W$ such that if $y_j(v) \leq z_j(v)$ for each $v \in W$ and for each $j = 1, \dots, k$, we have $\mathbf{y} \leq \mathbf{z}$; otherwise $\mathbf{y}$ and $\mathbf{z}$ are incomparable. The join $\mathbf{y} \vee \mathbf{z}$ and the meet $\mathbf{y} \wedge \mathbf{z}$ can be obtained for any two elements $\mathbf{y}, \mathbf{z}$ of $\mathscr{L}_W$ analogously to the case of $\mathscr{L}_V$, defining the finite distributive lattice $(\mathscr{L}_W, \vee, \wedge)$. Now let $p(\mathbf{u})$ be the probability that for each leaf v in W, v is assigned $\mathbf{u}(v)$.

This marginal probability is given by

$$(2) \qquad p(\mathbf{u}) = \sum_{\mathbf{U} \in \mathscr{A}_{\mathbf{u}}} P(\mathbf{U}), \text{ where } \mathscr{A}_{\mathbf{u}} := \{\mathbf{U} \in \mathscr{L}_V : \mathbf{U}|_W = \mathbf{u}\}.$$

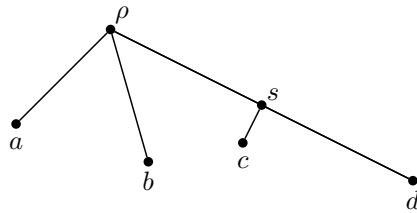An example to illustrate this concept is provided in Figure 1.

FIG. 1. *Let $k = 1$, and let $\mathbf{u}$ be denoted by u. In this example, if $W = \{a, b, c, d\}$, $u(a) = u(c) = 0$, and $u(b) = u(d) = 1$, then $p(u) = \pi_0 P(\rho, a)_{00} P(\rho, b)_{01} P(\rho, s)_{00} P(s, c)_{00} P(s, d)_{01} + \pi_0 P(\rho, a)_{00} P(\rho, b)_{01} P(\rho, s)_{01} P(s, c)_{10} P(s, d)_{11} + \pi_1 P(\rho, a)_{10} P(\rho, b)_{11} P(\rho, s)_{10} P(s, c)_{00} P(s, d)_{01} + \pi_1 P(\rho, a)_{10} P(\rho, b)_{11} P(\rho, s)_{11} P(s, c)_{10} P(s, d)_{11}.$*

The following proposition extends a result from [18], which dealt with the special case $k = 1$.

PROPOSITION 2.1. *Consider k independent two-state Markov processes on a tree with leaf set X. Assume that for each of them, all the determinants of the transition matrices are nonnegative. Then, for the corresponding multivariate process and for any two characters $\mathbf{y}, \mathbf{z} \colon W \to \{0,1\}^k$ on X from a fixed nonempty subset W of X, we have*

$$p(\mathbf{y}) \cdot p(\mathbf{z}) \leq p(\mathbf{y} \vee \mathbf{z}) \cdot p(\mathbf{y} \wedge \mathbf{z}).$$

Proof. Consider any two elements $\mathbf{Y}$ and $\mathbf{Z}$ of $\mathscr{L}_V$. We first prove the following:

$$(3) \qquad P(\mathbf{Y}) \cdot P(\mathbf{Z}) \leq P(\mathbf{Y} \vee \mathbf{Z}) \cdot P(\mathbf{Y} \wedge \mathbf{Z}).$$

Denote the term in the brackets of (1) by $P_j(\mathbf{U})$ to get $P(\mathbf{U}) = \prod_{j=1}^k P_j(\mathbf{U})$. Applying this to $\mathbf{U} \in \{\mathbf{Y}, \mathbf{Z}, \mathbf{Y} \vee \mathbf{Z}, \mathbf{Y} \wedge \mathbf{Z}\}$, inequality (3) can be written as $\prod_{j=1}^k P_j(\mathbf{Y}) \prod_{j=1}^k P_j(\mathbf{Z}) \leq \prod_{j=1}^k P_j(\mathbf{Y} \vee \mathbf{Z}) \prod_{j=1}^k P_j(\mathbf{Y} \wedge \mathbf{Z})$. It is clear that proving $P_j(\mathbf{Y}) P_j(\mathbf{Z}) \leq P_j(\mathbf{Y} \vee \mathbf{Z}) P_j(\mathbf{Y} \wedge \mathbf{Z})$ for each j establishes (3).

So let j be an arbitrary index in $\{1, \dots, k\}$, and consider the products $P_j(\mathbf{Y}) P_j(\mathbf{Z})$ and $P_j(\mathbf{Y} \vee \mathbf{Z}) P_j(\mathbf{Y} \wedge \mathbf{Z})$. Each can be written as a product of two $\pi^{(j)}$ values multiplied by a product over the arcs (r, s) of T of two entries of $P^{(j)}(r, s)$. The products of the two $\pi^{(j)}$ terms agree in $P_j(\mathbf{Y}) P_j(\mathbf{Z})$ and $P_j(\mathbf{Y} \vee \mathbf{Z}) P_j(\mathbf{Y} \wedge \mathbf{Z})$, that is, $\pi^{(j)}_{\delta_j(\mathbf{Y}, \rho)} \pi^{(j)}_{\delta_j(\mathbf{Z}, \rho)} = \pi^{(j)}_{\delta_j(\mathbf{Y} \vee \mathbf{Z}, \rho)} \pi^{(j)}_{\delta_j(\mathbf{Y} \wedge \mathbf{Z}, \rho)}$. The products of the two $P^{(j)}(r, s)$ entries agree in $P_j(\mathbf{Y}) P_j(\mathbf{Z})$ and $P_j(\mathbf{Y} \vee \mathbf{Z}) P_j(\mathbf{Y} \wedge \mathbf{Z})$, except for the cases in which either

(i) $\delta_j(\mathbf{Y}, r) = 0$, $\delta_j(\mathbf{Y}, s) = 1$, $\delta_j(\mathbf{Z}, r) = 1$, and $\delta_j(\mathbf{Z}, s) = 0$, or (ii) $\delta_j(\mathbf{Y}, r) = 1$, $\delta_j(\mathbf{Y}, s) = 0$, $\delta_j(\mathbf{Z}, r) = 0$, and $\delta_j(\mathbf{Z}, s) = 1$. However, in both cases (i) and (ii), the product $P^{(j)}(r, s)_{01} P^{(j)}(r, s)_{10}$ appears in the term for $P_j(\mathbf{Y})P_j(\mathbf{Z})$, while $P^{(j)}(r, s)_{00} P^{(j)}(r, s)_{11}$ appears in the term for $P_j(\mathbf{Y} \vee \mathbf{Z})P_j(\mathbf{Y} \wedge \mathbf{Z})$. The former term is less than or equal to the second since $P^{(j)}(r, s)_{00} P^{(j)}(r, s)_{11} - P^{(j)}(r, s)_{01} P^{(j)}(r, s)_{10} = \det P^{(j)}(r, s)$, which is nonnegative by our assumption. Consequently, all the terms in $P_j(\mathbf{Y})P_j(\mathbf{Z})$ are less than or equal to the corresponding terms in $P_j(\mathbf{Y} \vee \mathbf{Z})P_j(\mathbf{Y} \wedge \mathbf{Z})$. This establishes (3).

We now recall a form of the "four functions theorem," a classical result of Ahlswede and Daykin [1]. Let $(\mathscr{L}, \vee, \wedge)$ be a finite distributive lattice, and let α be a function that assigns a nonnegative real number to each element of $\mathscr{L}$. For a subset $\mathscr{A} \subseteq \mathscr{L}$, set $\alpha(\mathscr{A}) = \sum_{A \in \mathscr{A}} \alpha(A)$. If α satisfies the property that for any two elements A, B of $\mathscr{L}$, $\alpha(A)\alpha(B) \leq \alpha(A \vee B)\alpha(A \wedge B)$, then

$$(4) \qquad \alpha(\mathscr{A})\alpha(\mathscr{B}) \leq \alpha(\mathscr{A} \vee \mathscr{B})\alpha(\mathscr{A} \wedge \mathscr{B}),$$

where $\mathscr{A} \vee \mathscr{B} = \{A \vee B : A \in \mathscr{A}, B \in \mathscr{B}\}$ and $\mathscr{A} \wedge \mathscr{B} = \{A \wedge B : A \in \mathscr{A}, B \in \mathscr{B}\}$.

We apply this theorem by taking $\mathscr{L} = \mathscr{L}_V, \alpha = P$ and noting that α satisfies the required hypothesis by (3). Consider any fixed nonempty subset W of X, and recall the definition (for $\mathbf{u} \in \mathscr{L}_W$) of $\mathscr{A}_\mathbf{u}$ in (2). Note that

$$\mathscr{A}_\mathbf{y} \vee \mathscr{A}_\mathbf{z} = \mathscr{A}_{\mathbf{y} \vee \mathbf{z}} \text{ and } \mathscr{A}_\mathbf{y} \wedge \mathscr{A}_\mathbf{z} = \mathscr{A}_{\mathbf{y} \wedge \mathbf{z}}.$$

Thus, taking $\mathscr{A} = \mathscr{A}_\mathbf{y}$ and $\mathscr{B} = \mathscr{A}_\mathbf{z}$ in (4), we deduce that

$$\alpha(\mathscr{A}_\mathbf{y})\alpha(\mathscr{A}_\mathbf{z}) \leq \alpha(\mathscr{A}_{\mathbf{y} \vee \mathbf{z}})\alpha(\mathscr{A}_{\mathbf{y} \wedge \mathbf{z}}),$$

which is, by $\alpha = P$ and (2), equivalent to $p(\mathbf{y})p(\mathbf{z}) \leq p(\mathbf{y} \vee \mathbf{z})p(\mathbf{y} \wedge \mathbf{z})$. □

3. Application: Predicting future phylogenetic diversity.

3.1. Expected future phylogenetic diversity. In this section, we use Proposition 2.1 to obtain an inequality concerning the expected loss of biodiversity under species extinction models. Consider a rooted directed tree $T = (V_T, A_T)$ in which all the arcs are directed away from the root and with leaf set X. Let each arc a in A_T be assigned a nonnegative length λ_a. Here, T represents the evolutionary history of the species in X, while λ_a refers either to the amount of the genetic change on arc a, or to its temporal duration, or to some other feature such as morphological diversity.

Given a subset Y of X, the *phylogenetic diversity* (PD) of Y, denoted φ_Y, is the sum of the lengths of the arcs of the minimal subtree of T that connects the root and the leaves in Y. PD has been widely used to measure the biodiversity of a group of species [3, 4, 13, 18]; informally, the PD-score of a subset Y measures how much total "genetic" or "evolutionary" diversity in the tree is spanned just by the species in Y (depending on whether the lengths assigned to the edges reflect the amount of genetic change or evolutionary time, respectively).

As a function from 2^X to $\mathbb{R}^{\geq 0}$, φ has some attractive properties for the discrete mathematician: as well as being a submodular, increasing function, it also has the property that the subsets of X of given cardinality that have maximal φ value form a (strong) greedoid, and so can be quickly constructed by the greedy algorithm (for details, see [12]).

Assume that species in X undergo random extinction, and let E_x denote the event that a species $x \in X$ is extinct at some fixed future time t. Consider the phylogenetic
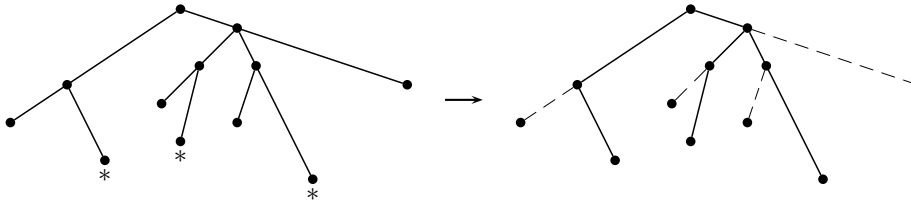
FIG. 2. *If only the species marked * in the tree on the left survive, then the future PD is the sum of the lengths of the solid edges in the tree on the right.*

diversity φ of the group of species that are still extant at time t. This random variable is referred to as *future PD*. An example to illustrate this notion is given in Figure 2.

The expected value of φ is

$$(5) \quad \mathbb{E}[\varphi] = \sum_{a=(u,v)\in A_T} \lambda_a \cdot \left(1 - \mathbb{P}\left(\bigcap_{x\in C_v} E_x\right)\right) = \varphi_X - \sum_{a=(u,v)\in A_T} \lambda_a \cdot \mathbb{P}\left(\bigcap_{x\in C_v} E_x\right),$$

where C_v denotes the subset of X which is separated from the root by v and which equals $\{v\}$ if v is a leaf vertex. $\mathbb{E}[\varphi]$ is referred to as *expected future PD*.

In the *generalized field-of-bullets* (g-FOB) model [4], the events $E_x^{(g)} := E_x$ are independent, and so the probability $\mathbb{P}(\bigcap_{x\in C_v} E_x^{(g)})$ that all the species descended from v become extinct can be written as

$$(6) \quad \mathbb{P}\left(\bigcap_{x\in C_v} E_x^{(g)}\right) = \prod_{x\in C_v} p_x,$$

where p_x denotes the probability $\mathbb{P}(E_x^{(g)})$.

However, the assumption that the events E_x are independent is likely to be unrealistic in most settings (see, for example, [8, 16]). In particular, rates at which lineages become extinct may be influenced by some species traits [11, 6]. The model referred to as the *state-based field-of-bullets* (s-FOB) model in [18] is based on the idea that closely related species in T are more likely to share attributes that may put them at risk in a hostile future environment. It assumes that the extinction risk of each species is influenced by some associated binary state with values 0 and 1, where state 0 confers an elevated risk of extinction, for example, under climate change.

Here, we generalize this model and suppose that the extinction risk of each species x is influenced by k binary states, each of which takes values in $\{0,1\}$, where state 1 is always advantageous over state 0 for x. We suppose that it is not known what features will help species survive, and so the states are not known for the species in X. However, we assume that the k states have evolved under k independent Markovian models on T assigning a state in $\{0,1\}^k$ to each species.

We assume further that if the states were determined at the leaves, then extinction would proceed according to the g-FOB model in which species x is extinct at time t with probability $p_x^{\mathbf{i}}$ if it is in state $\mathbf{i} \in \{0,1\}^k$. Finally, we suppose that for each species $x \in X$ and any two states $\mathbf{i} = (i_1,\ldots,i_k)$ and $\mathbf{l} = (l_1,\ldots,l_k)$,

$$(7) \quad p_x^{\mathbf{i}} \le p_x^{\mathbf{l}} \text{ whenever } l_j \le i_j \text{ for each } j = 1,\ldots,k.$$

This condition says that state $\mathbf{l}$ confers at least as high an extinction risk on a species x as state $\mathbf{i}$ if all the binary states in $\mathbf{i}$ are at least as "advantageous" for x as the

binary states in $\mathbf{l}$. Note, however, that if condition $l_j \leq i_j$ is not satisfied for every j, there is no prescribed relationship between $p_x^{\mathbf{i}}$ and $p_x^{\mathbf{l}}$. We have the freedom to either specify these relationships according to the needs of the model being studied, or leave them unspecified. For example, we may assume that the k binary states are ordered in a decreasing manner by their importance for survival and that $p_x^{\mathbf{i}} \leq p_x^{\mathbf{l}}$, whenever $l_j \leq i_j$ for the smallest coordinate $j \in \{1, \ldots, k\}$ for which $i_j \neq l_j$. Alternatively, we may assume that all the states are equally important for survival and that $p_x^{\mathbf{i}} \leq p_x^{\mathbf{l}}$, whenever $\sum_{j=1}^{k} l_j \leq \sum_{j=1}^{k} i_j$; that is, the more coordinates of the state assigned to x are 1 the smaller is the extinction probability of x. In the following, we assume only the relationships described in (7).

We call the model described above the *trait-dependent field-of-bullets* (t-FOB) model. In the case when $k = 1$, this model is the s-FOB model, whereas the case where for each x, $p_x^{\mathbf{i}} = p_x^{\mathbf{l}}$ for any two states $\mathbf{i}, \mathbf{l} \in \{0, 1\}^k$ gives the g-FOB model.

Given a t-FOB model, consider the g-FOB model in which the extinction probability of each species x is the same as in the t-FOB model. That is, if $\boldsymbol{\xi}$ describes the multivariate Markov process and the values $p_x^{\mathbf{i}}$ are the conditional extinction probabilities in the t-FOB model, then, in the associated g-FOB model, each species $x \in X$ goes extinct with probability

$$(8) \qquad p_x = \mathbb{P}[E_x^{(g)}] = \mathbb{P}[E_x^{(t)}] = \sum_{\mathbf{i} \in \{0,1\}^k} p_x^{\mathbf{i}} \mathbb{P}(\boldsymbol{\xi}(\mathbf{x}) = \mathbf{i}),$$

where $E_x^{(t)}$ denotes the event E_x under t-FOB. Theorem 3.1 compares the loss of PD under a t-FOB model with the PD loss under the associated g-FOB model.

THEOREM 3.1. *Consider a t-FOB model on a fixed tree T with nonnegative arc lengths and with leaf set X. The expected future PD of this model is less than or equal to the expected future PD of the associated g-FOB model.*

Proof. Let $\boldsymbol{\xi}$ and $p_x^{\mathbf{i}}$ denote the Markov process and the extinction probabilities of the t-FOB model, respectively. In view of (5) and (6), it suffices to show that

$$(9) \qquad \prod_{x \in C_v} p_x \leq \mathbb{P}\left(\bigcap_{x \in C_v} E_x^{(t)} \right),$$

where p_x is as given in (8). Recall how we defined the lattice $(\mathscr{L}_W, \vee, \wedge)$ for a Markov process on a tree and for a nonempty subset W of the leaf set of the tree in the previous section, and consider $(\mathscr{L}_{C_v}, \vee, \wedge)$. Since, for $\mathbf{u} \in \mathscr{L}_{C_v}$, $p(\mathbf{u})$ denotes the probability that for each $x \in C_v$, x is assigned $\mathbf{u}(x) \in \{0, 1\}^k$, we get

$$\mathbb{P}\left(\bigcap_{x \in C_v} E_x^{(t)} \right) = \sum_{\mathbf{u} \in \mathscr{L}_{C_v}} p(\mathbf{u}) \prod_{x \in C_v} f_x(\mathbf{u}),$$

where $f_x(\mathbf{u})$ is the probability that x becomes extinct given that it is in state $\mathbf{u}(x)$; that is, $f_x(\mathbf{u}) = p_x^{\mathbf{u}(x)}$. Moreover, for each $x \in C_v$, we have

$$p_x = \sum_{\mathbf{i} \in \{0,1\}^k} p_x^{\mathbf{i}} \mathbb{P}(\boldsymbol{\xi}(x) = \mathbf{i}) = \sum_{\mathbf{i} \in \{0,1\}^k} p_x^{\mathbf{i}} \left(\sum_{\mathbf{u} \in \mathscr{L}_{C_v} : \mathbf{u}(x) = \mathbf{i}} p(\mathbf{u}) \right) = \sum_{\mathbf{u} \in \mathscr{L}_{C_v}} p(\mathbf{u}) f_x(\mathbf{u}).$$

Now we can rewrite (9) as

$$(10) \qquad \prod_{x \in C_v} \left(\sum_{\mathbf{u} \in \mathscr{L}_{C_v}} p(\mathbf{u}) f_x(\mathbf{u}) \right) \leq \sum_{\mathbf{u} \in \mathscr{L}_{C_v}} p(\mathbf{u}) \prod_{x \in C_v} f_x(\mathbf{u}).$$

The proof of (10) makes use of Proposition 2.1 as well as the following multivariate form of the FKG inequality [7]. Given a finite distributive lattice $(\mathscr{L}, \vee, \wedge)$, suppose that $f_1, f_2, \ldots, f_n$ are functions from $\mathscr{L}$ into the nonnegative real numbers that satisfy, for any two elements A and B of $\mathscr{L}$, the condition that

$$(11) \qquad A \leq B \Rightarrow f_i(A) \geq f_i(B).$$

Furthermore, suppose that μ is a probability measure on the elements of $\mathscr{L}$ which satisfies the condition that

$$(12) \qquad \mu(A)\mu(B) \leq \mu(A \vee B)\mu(A \wedge B) \text{ for any pair } A, B \in \mathscr{L}.$$

Then

$$(13) \qquad \prod_{i=1}^{n}\left(\sum_{A \in \mathscr{L}} \mu(A)f_i(A)\right) \leq \sum_{A \in \mathscr{L}} \mu(A)\prod_{i=1}^{n} f_i(A).$$

We apply this inequality by setting $\mathscr{L} = \mathscr{L}_{C_v}$, $\mu = p$, and $f_x(\mathbf{u}) = p_x^{\mathbf{u}(x)}$ for $\mathbf{u} \in \mathscr{L}_{C_v}$, $x \in C_v$. Note that f_x satisfies (11). Namely, $\mathbf{u} \leq \mathbf{y}$ (for $\mathbf{u}, \mathbf{y} \in \mathscr{L}_{C_v}$) means that $u_j(x) \leq y_j(x)$ for each coordinate j, which, by (7), implies $p_x^{\mathbf{u}(x)} \geq p_x^{\mathbf{y}(x)}$. Note also that μ satisfies (12) by Proposition 2.1. In view of these, (13) provides inequality (10), and the proof is complete. $\square$

3.2. Variance of future PD. Consider now the variance of φ,

$$(14) \qquad \mathrm{Var}[\varphi] = \mathrm{Cov}[\varphi, \varphi] = \sum_{a,b \in A_T} \lambda_a \lambda_b \, \mathrm{Cov}[Y_a, Y_b],$$

where Y_a is the random variable that takes value 1 if arc a is part of the subtree connecting the survival species and the root and takes value 0 otherwise. Our goal is to compare the variance under a t-FOB model to the variance under the associated g-FOB model.

Theorem 3.1 demonstrated a global inequality that holds between the first moment (mean) of the t-FOB and the associated g-FOB, and so it is natural to ask whether a similar inequality holds for the second moment (variance). We will show that it does not, by exhibiting a simple example. The variance of future PD is a quantity that has recently been studied by [10] as a measure of the extent to which one can predict the loss of future biodiversity (in that paper, simulations revealed that the variance in future PD is generally considerably larger in models that include ecological dependencies, such as food webs, than for an associated g-FOB model).

It is easy to find examples in which the variance under a t-FOB model is greater than that of the associated g-FOB model, and so we will only show that the former variance can be less than the latter. To this end, let T be the tree with leaf set $\{x, y\}$ in which the arcs b and c pointing to x and y, respectively, are incident with the single interior vertex of the tree, which is adjacent to the root by arc a. Consider $\mathrm{Cov}[Y_a, Y_a] = (1 - \mathbb{P}[E_x \cap E_y])\mathbb{P}[E_x \cap E_y]$, which is written as $(1 - \mathbb{P}[E_x^{(t)} \cap E_y^{(t)}])\mathbb{P}[E_x^{(t)} \cap E_y^{(t)}]$ in t-FOB and which becomes $(1 - p_x p_y)p_x p_y$ under g-FOB. Note that $\mathrm{Cov}[Y_a, Y_a]$ is less under a t-FOB than under the associated g-FOB if and only if $\mathbb{P}[E_x^{(t)} \cap E_y^{(t)}] > p_x p_y$ and $\mathbb{P}[E_x^{(t)} \cap E_y^{(t)}] + p_x p_y > 1$ hold. It is easy to see that these conditions can be satisfied by some t-FOB model (together with its g-FOB model) on T. Additionally,

for any such t-FOB model, a value of λ_a can be chosen that is large enough in relation to λ_b and λ_c so that $\lambda_a^2 \operatorname{Cov}[Y_a, Y_a]$ is the dominant term in (14), resulting in a greater total variance for the corresponding g-FOB.

The following example describes an s-FOB (that is, a t-FOB with $k = 1$) under which the variance is less than the variance under the associated g-FOB.
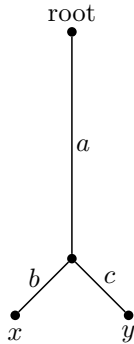

FIG. 3.

Example. Let T be the tree shown in Figure 3 with arc lengths $\lambda_a = 4$ and $\lambda_b = \lambda_c = 1$, and consider the following s-FOB model on T. Let ξ be a two-state Markov process on T with the state space $\{0, 1\}$ so that $\pi_0 = \pi_1 = \frac{1}{2}$ and each arc is assigned the transition matrix $\left(\begin{smallmatrix} \frac{3}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{3}{4} \end{smallmatrix}\right)$. Let $p_x^0 = p_y^0 = \frac{7}{8}$ and $p_x^1 = p_y^1 = \frac{6}{8}$. A careful check shows that the variance under this model is less than the variance under the associated g-FOB model (in which $p_x = p_y = \frac{13}{16}$).

REFERENCES

[1] R. Ahlswede and V. Blinovsky, *Lectures on Advances in Combinatorics*, Springer-Verlag, Berlin, 2008.

[2] J. T. Chang, *Full reconstruction of Markov models on evolutionary trees: Identifiability and consistency*, Math. Biosci., 137 (1996), pp. 51–73.

[3] D. P. Faith, *Conservation evaluation and phylogenetic diversity*, Biol. Conserv., 61 (1992), pp. 1–10.

[4] B. Faller, F. Pardi, and M. Steel, *Distribution of phylogenetic diversity under random extinction*, J. Theoret. Biol., 251 (2008), pp. 286–296.

[5] J. Felsenstein, *Inferring Phylogenies*, Sinauer Press, Victoria, Australia, 2003.

[6] R. G. FitzJohn, W. P. Maddison, and S. P. Otto, *Estimating trait-dependent speciation and extinction rates from incompletely resolved phylogenies*, Syst. Biol., 58 (2009), pp. 595–611.

[7] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre, *Correlation inequalities on some partially ordered sets*, Commun. Math. Phys., 22 (1971), pp. 89–103.

[8] S. B. Heard and A. O. Mooers, *Phylogenetically patterned speciation rates and extinction risks change the loss of evolutionary history during extinctions*, Proc. R. Soc. Lond. B, 267 (2000), pp. 613–620.

[9] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1994. Corrected reprint of the 1991 original.

[10] T. Ingram and M. Steel, *Modelling the unpredictability of biodiversity in ecological networks*, J. Theoret. Biol., 264 (2010), pp. 1047–1056.

[11] W. P. MADDISON, P. E. MIDFORD, AND S. P. OTTO, *Estimating a binary character's effect on speciation and extinction*, Syst. Biol., 56 (2007), pp. 701–710.

[12] V. MOULTON, C. SEMPLE, AND M. STEEL, *Optimizing phylogenetic diversity under constraints*, J. Theoret. Biol., 246 (2007), pp. 186–194.

[13] S. NEE AND R. M. MAY, *Extinction and the loss of evolutionary history*, Science, 278 (1997), pp. 692–694.

[14] D. M. RAUP, *Extinction: Bad Genes or Bad Luck?* W. W. Norton & Company, New York, 1992.

[15] C. SEMPLE AND M. STEEL, *Phylogenetics*, Oxford University Press, Oxford, UK, 2003.

[16] H. SIMIANER, *Accounting for non-independence of extinction probabilities in the derivation of conservation priorities based on Weitzman's diversity concept*, Conserv. Genet., 9 (2008), pp. 171–179.

[17] M. STEEL, *Recovering a tree from the leaf colourations it generates under a Markov model*, Appl. Math. Lett., 7 (1994), pp. 19–23.

[18] M. STEEL AND B. FALLER, *Markovian log-supermodularity, and its applications in phylogenetics*, Appl. Math. Lett., 22 (2009), pp. 1140–1144.