

A Theory of Diversity¹

KLAUS NEHRING

Department of Economics, University of California
Davis, CA 95616, U.S.A.
kdnehring@ucdavis.edu

and

CLEMENS PUPPE

Department of Economics, University of Bonn
Adenauerallee 24-42, D – 53113 Bonn, Germany
clemens.puppe@wiwi.uni-bonn.de

Final Revision, June 2001

¹For helpful comments we are indebted to Giacomo Bonanno, Adam Brandenburger, Alain Chateauneuf, Josef Falkinger, Karl Habermeier, Dorothea Herreiner, Jean-Yves Jaffray, Jörg Oechssler, John Roemer, Karl Schlag, Avner Shaked, Joaquim Silvestre, Martin van Hees, Martin Weitzman, as well as to the anonymous referees. We are also grateful for the open-mindedness of the Co-Editor in handling this submission. Preliminary versions of this paper have been presented at ESEM 1998 in Berlin, the Conference of the Society for Social Choice and Welfare, Vancouver 1998, the Conference on Ordinal and Symbolic Data Analysis, Amherst 1998, the Symposium on Operations Research, Dresden 2000, and in seminars at Harvard, Paris I, UC Davis, UCLA, UC Riverside, and the Universities of Berlin (Humboldt), Birmingham, Bonn, Darmstadt (Mathematics), Heidelberg, Karlsruhe, Mannheim and Vienna. We appreciate the stimulating responses of the participants. Financial support from the SFB 303 at Bonn University is gratefully acknowledged.

Title: A Theory of Diversity

Authors:

Klaus Nehring, Department of Economics, University of California,
One Shields Avenue, Davis, CA 95616, U.S.A., e-mail: kdnehring@ucdavis.edu

Clemens Puppe, Department of Economics, University of Bonn, Adenauerallee 24-42,
D – 53113 Bonn, Germany, e-mail: clemens.puppe@econ3.uni-bonn.de

Abstract *How can diversity be measured? What does it mean to value biodiversity? Can we assist Noah in constructing his preferences? To address these questions, we propose a multi-attribute approach under which the diversity of a set of species is the sum of the values of all attributes possessed by some species in the set. We develop the basic intuitions and requirements for a theory of diversity and show that the multi-attribute approach satisfies them in a flexible yet tractable manner.*

A natural starting point is to think of the diversity of a set as an aggregate of the pairwise dissimilarities between its elements. The multi-attribute framework allows one to make this program formally precise. It is shown that the program can be realized if and only if the family of relevant attributes is well-ordered (“acyclic”). Moreover, there is a unique functional form aggregating dissimilarity into diversity, the length of a minimum spanning tree. Examples are taxonomic hierarchies and lines representing uni-dimensional qualities. In multi-dimensional settings, pairwise dissimilarity information among elements is insufficient to determine their diversity. By consequence, the qualitative and quantitative behavior of diversity differs fundamentally.

KEYWORDS: Diversity, endangered species, similarity, conjugate Moebius inversion.

1 Introduction

“An angel is more valuable than a stone. It does not follow, however, that two angels are more valuable than one angel and one stone.” Thus Thomas Aquinas in *Summa contra Gentiles, III*.² While contemporary readers may not subscribe to the premise of St. Thomas’ argument, they have good reason to follow its logic, as in the following updated version: “A human being is more valuable than a chimpanzee. It does not follow, however, that 6,000,130,000 human beings and no chimpanzees are more valuable than 6,000,000,000 human beings and 130,000 chimpanzees.”³ Thomas’ argument suggests that “ontological diversity” has intrinsic value; the contemporary, narrower version makes an analogous case for biological diversity. By showing why straight counting is not good enough, the argument presents in a nutshell two major requirements on any satisfactory mathematical theory of valuing diversity: valuing as opposed to counting cannot treat single objects equally nor can it be additive. The argument points to *similarity* as the source of non-additivity: presumably, while a second angel is very similar to the first one, the stone is not. Hence, the *marginal* diversity value of the second angel is less than that of the stone. While equal valuation of single objects may sometimes be defensible, similarity and hence non-additivity are essential to the very concept of diversity. Relevance of similarity suggests also a third, cardinal requirement: since it is presumably the harder for objects to add to the diversity of an existing set the larger that set already is, one would expect diversity to be a *submodular* function.

In this paper, we explore a *multi-attribute* approach to valuing diversity whose starting point is the idea that the value of diversity consists in the realization of certain attributes/potentialities of life by some existing species; examples of such attributes are “being a primate,” “a carnivore,” “being able to live for more than 1000 years,” “possessing a certain gene,” etc. The diversity of a set is simply taken to be the sum of the numeric values (“weights”) of the attributes realized by some object in the set; set functions of this kind are called *diversity functions*.

Diversity functions satisfy the requirements outlined above in a natural and flexible manner. First, individual objects will be valued more highly if they realize higher-valued attributes; presumably, Aquinas valued the spirituality of an angel more highly than the raw physicality of a stone. Second, similarity and hence non-additivity are accounted for naturally: adding a stone to an angel enriches the universe by a significant attribute (physicality), whereas the second angel adds less (playing the harp perhaps?). The submodularity of diversity functions is also clear since the marginal diversity of an object is the total value of all of its attributes not already realized by the existing set, hence non-increasing in the size of that set. A somewhat stronger condition called “total submodularity” in fact *characterizes* diversity functions.

The basic purpose of the paper is to show that the multi-attribute approach can serve as a useful and in some contexts perhaps even canonical conceptual framework for thinking about diversity. In particular, we try to demonstrate that the framework provides a sound basis for quantifying the benefit of biodiversity preservation.

²The quotation is taken from Hurka (1983).

³In 1995, the number of chimpanzees on planet earth was less than the number of inhabitants of a small city (according to Dobson (1996)).

A Decision-Theoretic Foundation for Existence Value

The formalism developed in this paper is designed to capture the intrinsic (“existence”) rather than the instrumental (“use”) value of biodiversity; prime instances of the latter are agricultural and pharmaceutical uses of genetic resources.⁴ The instrumental value of biodiversity derives from the entailed provision of natural resources and thus depends critically on the amount of resource. For instance, the instrumental value of rhinoceroses for use of their horn, say as aphrodisiac or shaft of a dagger, clearly depends on their number. By contrast, the existence value of rhinoceroses derives from their possessing a unique combination of specific features. Additional individual rhinos have no marginal existence value, unless they have valuable features of their own.⁵ While the instrumental value is much easier to integrate into economic analyses, it is only a poor approximation of the intrinsic value. For instance, in terms of use value the rapid development of genetic engineering technology promises to more than make up for what is lost in natural genetic resources. However, such technological progress cannot substitute for the lost intrinsic value.

Since “existence value” may seem rather elusive a concept, it is highly desirable to endow it with a decision-theoretic foundation. For this purpose, we provide in Section 2 a representation theorem⁶ for a decision-maker’s preferences whom, borrowing a metaphor from Weitzman (1998), we shall call Noah. Noah’s preferences are defined over probability distributions of sets of existing species (“species lotteries”). The representation justifies the functional form assumed by the multi-attribute approach; specifically, it is shown that Noah’s von-Neumann-Morgenstern utility function must be a diversity function which, moreover, *uniquely* reveals his attribute weights.⁷ Noah’s (i.e. *our*) very real problem of *constructing* a preference relation thus boils down to one of assigning appropriate attribute weights.

Diversity as Aggregate Dissimilarity

The basic contribution of the multi-attribute approach is to explain diversity in terms of something simpler, the weight of different attributes. It is attractive partly because of its generality and flexibility; this flexibility comes, however, at an (apparent) price, namely the lack of specific structure and the great number of possible ways to fill it. One way to obtain more specific structure is to restrict the set of admissible attributes to a patterned family by setting the weight of all others to zero; this strategy will be extensively explored in this paper.

A natural alternative is to conceptualize diversity as “aggregate dissimilarity,” that is: as determined by the pairwise dissimilarities between objects. This indeed was the program of Weitzman (1992), who realized it for contexts in which attributes have an especially simple organization as “taxonomic hierarchies.” However, the hierarchical model entails strong restrictions on the associated dissimilarity metric.

⁴By contrast, the value supplied by the largest industrial user of biodiversity, the tourist and recreational industry, is arguably derived to a significant part from its intrinsic value.

⁵Indirectly, the number of currently existing individuals of a species matters also from an existence value perspective as it influences the survival probability of the species in the long run.

⁶The result is a re-interpretation of Nehring (1999a, Th.2).

⁷Strictly speaking, the latter requires an identification of co-extensive features, i.e. features possessed by exactly the same set of objects/species in the domain; thus, technically, attributes have to be understood as equivalence classes of co-extensive features (see Section 2.1).

Here, we will show that aggregation of dissimilarities is possible for more general dissimilarity metrics using more flexible organizations of attributes (“models”). Indeed, a main concern of this paper is to determine exactly when the view of diversity as aggregate dissimilarity is compatible with the multi-attribute approach. We find that the domain of compatibility extends significantly beyond that of taxonomic hierarchies, but still remains restricted to “one-dimensional” models. We conclude from this that the view of diversity as aggregate dissimilarity has inherent limitations, and that multi-dimensional situations need to be understood within a multi-attribute framework.

As a first step beyond the hierarchical model, we study of the line model, in which objects (species) are described in terms of a uni-dimensional quality characteristic. The line model assumes that all attributes are quality *intervals* such as “weighing more than 1 ton” or “weighing between 1 and 10 pounds.” The line model in effect generalizes the hierarchy model, and diversity is again a simple function of dissimilarity.

When, in general, can the diversity of an arbitrary set be determined from the pairwise dissimilarities between its elements? This question is made formally precise in terms of a property of *binary determinacy* in Section 4. The first main novel result of the paper, Theorem 4.1, shows binary determinacy to be equivalent to the condition that there be no cycles of three attributes, i.e. that there be no triple of objects such that for each pair of them there is an attribute possessed by both but not by the third object.⁸

Somewhat paradoxically, for certain binary determinate models the diversity of a set may decrease as the dissimilarities between its elements increase. Since it seems unsatisfactory to *explain* a decrease in diversity by an increase in the underlying dissimilarities, an intuitively compelling notion of diversity as aggregate dissimilarity requires the diversity of a set to be a *monotone* function of the dissimilarities between its elements. The second main result, Theorem 6.2, shows this requirement to correspond to a stronger condition of acyclicity of the underlying attribute family which in addition excludes cycles of length greater than three. As biologically interesting instances of acyclic models generalizing lines and hierarchies, we introduce “phylogenetic trees” that simultaneously incorporate evolutionary and taxonomic information. Theorem 6.2 has the corollary that within the multi-attribute approach, there is a *unique* way to aggregate pairwise dissimilarities into diversity, namely via the minimal length of a spanning tree. We also identify the major restriction imposed by acyclicity on the underlying dissimilarity metrics (Proposition 6.4).

While the view of diversity as aggregate dissimilarity thus turns out to have a limited domain of applicability, it offers powerful advantages. Since dissimilarity judgements translate directly into preference judgements over sets with at most two elements, it simplifies Noah’s task of constructing a preference ordering substantially.

The Distinction Between One- and Multi-Dimensional Models

An important example of a model that is not acyclic is the “hypercube,” the simplest and most basic multi-dimensional model, in which objects are described as strings of zeros and ones (embodying, for example, a genome), and attributes correspond to particular shared substrings (see Section 3.3). Starting from this example, we argue that acyclicity can be viewed as a natural definition of the *one-dimensionality* of an

⁸The exact characterizing condition (“no-free-triples”) is slightly weaker.

attribute family that is particularly well-suited for a theory of diversity. This perspective motivates an alternative reading of Theorem 6.2 as a result about one-dimensional (i.e. acyclic) attribute structures. Together with closely related results of the paper, it has two main implications. First, diversity behaves in a unified way in one-dimensional models since it is characterized as the length of a minimum spanning tree; this behavior exhibits the common-sense comparative statics property of monotonicity in dissimilarity. Secondly, the global behavior of diversity in one-dimensional models differs fundamentally from that in multi-dimensional models. In one-dimensional models the diversity of a set must grow roughly in proportion to its cardinality, while it may grow much more slowly in multi-dimensional models. As a result, ignoring the multi-dimensional nature of a situation may lead to misestimations of the true diversity by orders of magnitude, as illustrated in Section 7.

Relation to the Literature

Diversity theory can be viewed as a branch of utility theory with a focus on how marginal utility depends on the set of goods consumed. Valuing biodiversity is clearly a special case, with society consuming the existence of species as heterogeneous goods. In terms of utility theory, similarity between goods translates into their closeness as substitutes. While closeness as substitutes has been understood traditionally as elasticity of substitution between pairs of goods, the novel contribution of the diversity viewpoint is an analysis of the “combinatorics of substitution,” i.e. of the potentially complex dependence of marginal utilities on the *entire set* of goods consumed. A natural application is to models of differentiated products. Indeed, the multi-attribute framework can be viewed as a discrete version of the Lancasterian approach in which consumption goods are described as bundles of ultimately valued characteristics (see Lancaster (1966), (1979)). Our analysis contributes to this approach the beginnings of a rich representation theory, both for the general model as well as for the restrictions entailed by specific “consumption technologies.” In the present paper, we study a few especially natural and basic examples. A more systematic treatment is provided in a companion paper, “Diversity and the Geometry of Similarity” (Nehring (1999b)).

In decision theory, preferences over a domain of sets have been studied in the context of “preference for flexibility” following Kreps (1979). It follows from the analysis in Nehring (1999a) that “preference for flexibility” can be thought of as preference for diversity of opportunity sets. The role of diversity has also been emphasized, following Pattanaik and Xu (1990), in the related literature on the measurement of opportunity and “freedom of choice” (see, e.g., Sugden (1998) for a recent overview). In the context of valuing biodiversity, the flexibility value of existing species for instrumental purposes is often referred to as their *option value*. In Section 2.5.1, we indicate how this paper’s framework can be interpreted along these lines.

The literature most directly relevant to our paper consists of the two pioneering contributions by M. Weitzman, “On Diversity” (1992, henceforth W92), and “The Noah’s Ark Problem” (1998, henceforth W98). The general program of W92 was to derive the diversity of a set from pairwise dissimilarities between its elements. The multi-attribute approach is conceptually more general as it derives both diversity *and* dissimilarity from more primitive data, namely attribute weights: the dissimilarity of an object x from another object y is defined as the total weight of all attributes possessed by x but not by y . With similar intent, Weitzman proposed in W98 the “species/library” metaphor which identifies species with collections of books (“libraries”); the multi-attribute approach

developed here can be viewed as a systematic and mathematically fully articulated development of Weitzman’s metaphor, confirming (we hope) his belief in the metaphor’s “staying power.” Both W92 and W98 focus on taxonomic hierarchies. While analytically especially transparent and well-behaved, taxonomic hierarchies have very special properties and are thus of limited applicability. The vast space of non-hierarchical attribute structures is virgin territory.

Since the early 90s, biologists have independently begun to conceptualize diversity in terms of the counting of attributes in a spirit akin to Weitzman’s species/library metaphor; see Gaston (1996) for an overview. Most of the specific mathematical models proposed in this literature attempt to model phylogenetic diversity based on evolutionary trees, and can be viewed as variants or extensions of the hierarchy model. The most elaborate of these (as far as we know), proposed by Faith and Walker (1994), can be shown to correspond to an acyclic attribute structure with specific assumptions about attribute weights; it is thus a special case of the theory developed in Section 6.

The remainder of the paper is organized as follows. Section 2 introduces the multi-attribute framework. Section 3 describes and analyzes the three paradigmatic models of taxonomic hierarchies (Section 3.1), lines (Section 3.2), and the hypercube (Section 3.3). A further natural model, the “tree model” is analyzed in Section 5. Sections 4 and 6 form the mathematical core of the paper, providing the characterizations of binary determinacy and monotonicity in dissimilarity in terms of triple acyclicity and acyclicity, respectively. Central to the latter is a characterization of the geometry of acyclic attribute structures in terms of trees (Theorem 6.1). Section 7 illustrates the fundamentally different quantitative behavior of diversity in one- versus multi-dimensional models. In the concluding Section 8, we briefly indicate the existence of promising applications of diversity theory to two other contexts, economies of scope and case-based decision theory. An application to the measurement of social diversity is sketched in the body of the paper; see 2.5.2 for the basic idea. All proofs are collected in an appendix.

2 The Multi-Attribute Model of Diversity

2.1 Diversity Functions

Let X be a finite universe of objects, with 2^X denoting its power set. In the context of biodiversity, an object $x \in X$ may correspond to a biological species, for instance.⁹ The basic idea underlying the proposed multi-attribute approach is to view the diversity of a set $S \subseteq X$ as determined by the value of the different features possessed by the objects (species) in S . An elementary but fundamental step in the development of the theory is to describe features in terms of their *extension*, i.e. to identify each feature with the subset A of all objects in X that possess the feature in question. A subset $A \subseteq X$ thus defines an equivalence class of co-extensive features, henceforth referred to as the *attribute* associated with A . For all $A \subseteq X$, denote by $\lambda_A \geq 0$ the *weight* of the attribute A , i.e. the weight of all features with extension A . Upon normalization, λ_A can be thought of as the relative importance of the corresponding attribute. If there is no feature possessed exactly by the objects in A , then $\lambda_A = 0$. An object x possesses the attribute A , i.e. one and therefore all features with extension A , if $x \in A$. A set

⁹Other entities may play the role of the fundamental “unit of conservation,” such as genes, habitats, or individual organisms.

S realizes the attribute A if $A \cap S \neq \emptyset$, i.e. if there is at least one element in S that possesses the attribute A .¹⁰

Definition A function $v : 2^X \rightarrow \mathbf{R}$ is called a **diversity function** if there exists a positive measure (additive set-function) λ on 2^{2^X} such that for all $S \subseteq X$,

$$v(S) = \lambda(\{A \subseteq X : A \cap S \neq \emptyset\}) = \sum_{A \subseteq X : A \cap S \neq \emptyset} \lambda_A, \quad (2.1)$$

where $\lambda_A := \lambda(\{A\})$ and, by convention, $v(\emptyset) = \lambda_\emptyset = 0$.¹¹ The function $\lambda : A \mapsto \lambda_A$ is referred to as the **attribute weighting function** associated with v . The support $\Lambda := \{A \subseteq X : \lambda_A \neq 0\}$, i.e. the set of attributes with non-zero weight, will be called the family of **relevant attributes**.

By (2.1), the diversity value of a set S is given by the sum of the weights of all attributes realized by the set S . Note that each attribute occurs at most once in the sum. In particular, each single object contributes to diversity the weight of all those attributes that are not possessed by any already existing object. The functional form (2.1) may look rather special, both in assuming additive separability and in disregarding how often attributes are realized. A rigorous justification will be given via a representation theorem in Section 2.4 below.

A central aspect of the multi-attribute approach is to characterize the qualitative and quantitative behavior of diversity in terms of properties of the underlying family Λ of relevant attributes. In principle, every conceivable attribute might be deemed relevant, i.e. one may take $\Lambda = 2^X \setminus \{\emptyset\}$. This, however, easily gets too complex since with n objects there are as many as $2^n - 1$ conceivable attributes. As a drastic countermeasure to this combinatorial explosion, one might consider restricting the family of relevant attributes to the family of all singleton attributes, i.e. to take $\Lambda = \{\{x\} : x \in X\}$. This, however, leads to an additive diversity function, i.e. $v(S) = \sum_{x \in S} v(\{x\})$, and thus ignores any similarities between elements. In most applications, the family of relevant attributes will exhibit particular structure that one will wish to incorporate.

Example 1 (Whales, Rhinos, and Sharks) Consider a universe consisting of the three species whales (wh), rhinoceroses (rh) and sharks (sh).¹² Considering the biological classification of species, two natural attributes in this example are “being a mammal,” corresponding to the set $\{wh, rh\}$, and “being a fish,” corresponding to $\{sh\}$.¹³ A finer distinction would also include the two remaining idiosyncratic attributes $\{wh\}$ (“being a whale”) and $\{rh\}$ (“being a rhino”) as relevant attributes. Suppose for now that this list exhausts the attributes deemed relevant. Applying formula (2.1), one obtains $v(\{sh, rh\}) = v(\{sh\}) + v(\{rh\})$, i.e. the diversity value of shark and rhino species together equals the sum of the value of each species taken

¹⁰Subsets of X thus take on a double role as sets to be evaluated in terms of diversity on the one hand, and as weighted attributes on the other. In order to notationally distinguish these roles we will always denote generic subsets by the symbol “ A ” whenever they are viewed as attributes, and by the symbol “ S ” otherwise.

¹¹Technically, diversity functions that are further normalized so that $v(X) = 1$ are known as *plausibility functions* (“conjugate belief functions”) in the sense of Dempster-Shafer (Dempster (1967), Shafer (1976)). Observe that in this case λ is a probability measure.

¹²Frequently, our use of the term “species” is that of a layman rather than that of a biologist.

¹³Note that, due to their extensional identification, the interpretation of attributes depends on a prespecified universe. For instance, the idiosyncratic attribute $\{sh\}$ corresponds to “being a fish” only in a universe that contains no other fish.

separately. Intuitively, this reflects the heterogeneity of the two species. On the other hand, the joint attribute “mammal” renders rhinos and whales similar with respect to their biological classification. Applying (2.1) to the set of these two species yields $v(\{wh, rh\}) = v(\{wh\}) + v(\{rh\}) - \lambda_{\{wh, rh\}}$, i.e. the diversity value of whale and rhino species together is less than the sum of the corresponding individual values by the weight $\lambda_{\{wh, rh\}}$ of their common attribute “mammal.” In this way, the multi-attribute approach captures the central intuition that the diversity of a set is reduced by similarities between its elements.

The family of relevant attributes in Example 1 exhibits a *hierarchical* structure in the sense that any two relevant attributes are either disjoint, or one is contained in the other (see Figure 1, left, with the small circles representing the idiosyncratic attributes). This structure underlies most of the work of Weitzman W92 and W98; it is analyzed in more detail in Section 3.1 below. While particularly transparent and well-behaved, the hierarchical model turns out to be too restrictive in many cases. To illustrate, consider again the species in Example 1.

Example 1' Suppose that $v(\{wh\}) = v(\{rh\})$, i.e. suppose that whales and rhinos have equal value in isolation. Then, given the relevant attributes as in Example 1 above, formula (2.1) entails the restriction $v(\{rh, sh\}) = v(\{wh, sh\})$. However, this conclusion fails to reflect an important commonality between sharks and whales, namely their living in the ocean. Adding the relevant attribute $\{wh, sh\}$ corresponding to the common habitat sets things right: in accordance with intuition, one now obtains $v(\{rh, sh\}) > v(\{wh, sh\})$. For verification, note that, given the new attribute $\{wh, sh\}$ (“living in the ocean”), the assumption $v(\{wh\}) = v(\{rh\})$ implies that $\lambda_{\{rh\}} > \lambda_{\{wh\}}$.

Clearly, including the new attribute $\{wh, sh\}$ destroys the hierarchical structure since it overlaps with the attribute $\{wh, rh\}$ (“mammal”). Nevertheless, the resulting family of relevant attributes still admits a tractable structure. Indeed, the three species can be ordered from left to right in such a way that all relevant attributes are *intervals* with respect to the ordering (see Figure 1, right). The resulting “line model” is analyzed in Section 3.2 below.




Figure 1: Hierarchical versus linear organization of attributes

2.2 Distinctiveness and Dissimilarity

For practical purposes, marginal diversity will often play a decisive role. Consider the following problem, henceforth referred to as *Noah’s local problem*: Suppose the set of existing species is S , but one of the species in $\{y_1, y_2\} \subseteq S$ must be sacrificed. Using a metaphor from W98, Noah’s ark cannot carry both y_1 and y_2 . Which species should Noah let go? Formally, the obvious solution is to sacrifice y_1 if and only if $v(S \setminus \{y_1\}) \geq v(S \setminus \{y_2\})$, i.e. to sacrifice the species y with the lower marginal diversity

$v(S) - v(S \setminus \{y\})$, intuitively: to sacrifice the least distinct species. Indeed, in the multi-attribute model, the *distinctiveness* of species x from the set S is defined naturally as

$$d(x, S) := v(S \cup \{x\}) - v(S). \quad (2.2)$$

This interpretation of marginal diversity as distinctiveness is justified by the observation that $d(x, S)$ is the total weight of all attributes possessed by x that are not already realized by S , i.e.

$$d(x, S) = \lambda(\{A : x \in A, A \cap S = \emptyset\}).$$

Also, let

$$d(x, y) := d(x, \{y\}) = v(\{x, y\}) - v(\{y\}) = \lambda(\{A : x \in A, y \notin A\}) \quad (2.3)$$

denote the **dissimilarity** of x from y , i.e. the total weight of all attributes possessed by x but not by y . Due to non-negativity of λ , the function $d : X \times X \rightarrow \mathbf{R}$ satisfies the triangle-inequality, i.e. $d(x, z) \leq d(x, y) + d(y, z)$ for all x, y, z . Hence, $d(\cdot, \cdot)$ is a pseudo-metric, and we will refer to it as the *dissimilarity (pseudo-)metric* associated with a diversity function. In general, the function d may be non-symmetric, and thus fail to be a proper metric. It is symmetric if and only if $v(\{x\}) = v(\{y\})$ for all x, y . This case of equal valuation of singletons will be referred to as the *uniform* case.

2.3 Recovering Attribute Weights via Conjugate Moebius Inversion

On the multi-attribute approach, the diversity value of a set is defined as the total weight of all attributes realized. In this subsection we show how the underlying attribute weighting function can be recovered from a given diversity function using *conjugate Moebius inversion*.

Recovering attribute weights from a given diversity function is particularly simple in the case of idiosyncratic attributes: the weight $\lambda_{\{x\}}$ of the attribute $\{x\}$ is given by the total weight of all attributes ($= v(X)$) minus the total weight of all attributes possessed by any object different from x ($= v(X \setminus \{x\})$). In general, one has the following result.¹⁴

Fact 2.1 (Conjugate Moebius Inversion) *For any function $v : 2^X \rightarrow \mathbf{R}$ with $v(\emptyset) = 0$ there exists a unique function $\lambda : 2^X \rightarrow \mathbf{R}$, the conjugate Moebius inverse, such that $\lambda_\emptyset = 0$ and, for all S ,*

$$v(S) = \sum_{A: A \cap S \neq \emptyset} \lambda_A.$$

Furthermore, the conjugate Moebius inverse λ is given by the following formula. For all $A \neq \emptyset$,

$$\lambda_A = \sum_{S \subseteq A} (-1)^{\#(A \setminus S)+1} \cdot v(S^c),$$

where S^c denotes the complement of S in X .

¹⁴The result is obtained by applying (ordinary) Moebius inversion (cf. Shapley (1953), Rota (1964)) to the associated *loss function* \hat{v} , defined by $\hat{v}(S) := v(X) - v(X \setminus S)$ (see Nehring (1999a)).

By Fact 2.1, the only restriction on a set function $v : 2^X \rightarrow \mathbf{R}$ entailed by the multi-attribute model of diversity is $\lambda \geq 0$, i.e. positivity of the corresponding conjugate Moebius inverse. In order to characterize this restriction directly in terms of properties of the function v we need the following definitions. A function $v : 2^X \rightarrow \mathbf{R}$ is *monotone* if $W \subseteq S$ implies $v(W) \leq v(S)$. Furthermore, v is called *totally submodular* if, for any collection $\{S_i\}_{i \in I}$,

$$v\left(\bigcap_{i \in I} S_i\right) \leq \sum_{J: \emptyset \neq J \subseteq I} (-1)^{\#J+1} \cdot v\left(\bigcup_{i \in J} S_i\right). \quad (2.4)$$

The following fact can be derived from Chateauneuf and Jaffray (1989, Cor. 1).

Fact 2.2 *The function $v : 2^X \rightarrow \mathbf{R}$ is a diversity function, i.e. has a positive conjugate Moebius inverse, if and only if v is monotone and totally submodular.*

The basic instance of total submodularity is given by the case of $\#I = 2$ in which (2.4) specializes to the following condition known as *submodularity*. For all S_1, S_2 ,

$$v(S_1 \cap S_2) + v(S_1 \cup S_2) \leq v(S_1) + v(S_2).$$

Equivalently, for all S_1, S_2 with $S_2 \subseteq S_1$, and all x ,

$$v(S_1 \cup \{x\}) - v(S_1) \leq v(S_2 \cup \{x\}) - v(S_2), \quad (2.5)$$

i.e. the marginal value of an additional object decreases with the set of objects already available. In the multi-attribute model, the property of decreasing marginal diversity is immediate since $v(S \cup \{x\}) - v(S) = d(x, S) = \lambda(\{A : x \in A, A \cap S = \emptyset\})$, which is decreasing in S due to the positivity of λ .

2.4 A Decision-Theoretic Representation Theorem

In this subsection, we provide a decision-theoretic foundation of the proposed concept of diversity. Specifically, we derive the cardinal scale of diversity functions and the positivity of their conjugate Moebius inverse from properties of an underlying preference relation. The key to this is to view diversity functions as von Neumann-Morgenstern utility functions representing binary relations \succeq defined on the simplex Δ^{2^X} of *set-lotteries* (“species lotteries”) with appropriate properties. The intended interpretation of \succeq is that for all lotteries $p, q \in \Delta^{2^X}$, $p \succeq q$ if and only if p is (weakly) preferred to q in terms of entailed *expected diversity*. As usual, we denote by \sim and \succ the symmetric and asymmetric part of \succeq , respectively. In the biodiversity context, set-lotteries can be identified with *acts* that result in the probabilistic survival of certain species. The desired representation of \succeq is that for all acts $p, q \in \Delta^{2^X}$,

$$p \succeq q \Leftrightarrow \sum_{S \subseteq X} p_S \cdot v(S) \geq \sum_{S \subseteq X} q_S \cdot v(S), \quad (2.6)$$

for some diversity function v , where p_S and q_S denote the probability of S under p and q , respectively. The representation (2.6) in terms of expected diversity is guaranteed by the following axioms imposed on \succeq . The first is the set of von Neumann-Morgenstern axioms.

Axiom vNM (von Neumann-Morgenstern)

- i) **(Completeness)** For all $p, q \in \Delta^{2^X}$, $p \succeq q$ or $q \succeq p$.
- ii) **(Transitivity)** For all $p, q, r \in \Delta^{2^X}$, $p \succeq q$ and $q \succeq r$ imply $p \succeq r$.
- iii) **(Independence)** For all $p, q, r \in \Delta^{2^X}$, and all $a \in (0, 1)$,
 $p \succeq q \Leftrightarrow ap + (1 - a)r \succeq aq + (1 - a)r$.
- iv) **(Continuity)** For all $p, q, r \in \Delta^{2^X}$, $p \succeq q \succeq r$ implies the existence of $a \in [0, 1]$ such that $q \sim ap + (1 - a)r$.

It is well-known that vNM implies the existence of a function $v : 2^X \rightarrow \mathbf{R}$, unique up to a positive affine transformation, such that (2.6) holds. Furthermore, by Fact 2.1, for any such function v with $v(\emptyset) = 0$ there exists a unique measure λ satisfying (2.1) for all S . Consequently, the only additional restriction entailed by the representation of \succeq in terms of expected *diversity* is that $\lambda_A \geq 0$ for all A . It is non-negativity of all weights attached to attributes that allows one to interpret v as a diversity function, and thus demarcates the proposed concept of diversity from overall subjective utility which may include instrumental components. For instance, an arachnophobic individual attaches negative instrumental utility to individual spiders, but may well value the existence of spiders as a species.

Positivity of the conjugate Moebius inverse λ is guaranteed by the following condition. For any family $\mathcal{S} \subseteq 2^X$, let $p(\mathcal{S}) = \sum_{S \in \mathcal{S}} p_S$ denote the probability of \mathcal{S} .

Axiom POS (Positivity) For all $p, q \in \Delta^{2^X}$, $p \succeq q$ whenever

$$p(\{S : S \cap A \neq \emptyset\}) \geq q(\{S : S \cap A \neq \emptyset\}) \text{ for all } A.$$

Axiom POS can be paraphrased as follows: If, for any conceivable attribute A , the probability that A is realized under p is at least as great as the corresponding probability under q , then p is (weakly) preferred to q . To illustrate, consider subsets S_1, S_2, W such that $W \supseteq (S_1 \cup S_2)$ and $S_1 \cap S_2 = \emptyset$. Then, condition POS implies

$$\left[\frac{1}{2} \cdot \mathbf{1}_{W \setminus S_1} + \frac{1}{2} \cdot \mathbf{1}_{W \setminus S_2} \right] \succeq \left[\frac{1}{2} \cdot \mathbf{1}_W + \frac{1}{2} \cdot \mathbf{1}_{W \setminus (S_1 \cup S_2)} \right], \quad (2.7)$$

where, for any S , $\mathbf{1}_S$ denotes the lottery that gives S with certainty. Hence, losing one of the subsets S_1 or S_2 for sure (with equal odds) is always preferred to facing a fifty-percent chance of losing both S_1 and S_2 .¹⁵ This can be verified as follows. Any attribute A that intersects $W \setminus (S_1 \cup S_2)$ will be realized with certainty under both lotteries in (2.7). On the other hand, all attributes A that intersect W but not $W \setminus (S_1 \cup S_2)$ will be realized with probability $\frac{1}{2}$ under the right-hand lottery, and with probability $\geq \frac{1}{2}$ under the left-hand lottery. By POS, the left-hand lottery must be preferred. Observe that, given a representation in terms of expected diversity, the preference in (2.7) is strict whenever there exists a relevant attribute $A \subseteq (S_1 \cup S_2)$ such that $A \cap S_1 \neq \emptyset$ and $A \cap S_2 \neq \emptyset$. For instance, suppose that W consists of whales, rhinos and sharks. Then, with $S_1 = \{wh\}$ and $S_2 = \{rh\}$, the lottery on the right-hand side of (2.7) entails a survival probability of $\frac{1}{2}$ for mammals, while the probability of survival of at least one mammal is 1 under the lottery on the left-hand side. The strict

¹⁵A risk-aversion implication of this kind has already been highlighted in W98.

preference for the left-hand lottery thus follows from the attribute “mammal” having a strictly positive weight.

The following result shows that vNM and POS jointly imply the desired representation.

Theorem 2.1 *Let \succeq be a binary relation on Δ^{2^X} . The following statements are equivalent.*

(i) \succeq satisfies vNM and POS.

(ii) There exists a diversity function $v : 2^X \rightarrow \mathbf{R}$, unique up to multiplication by a positive scalar, such that for all $p, q \in \Delta^{2^X}$,

$$p \succeq q \Leftrightarrow \sum_{S \subseteq X} p_S \cdot v(S) \geq \sum_{S \subseteq X} q_S \cdot v(S).$$

(iii) There exists a non-negative measure λ on 2^{2^X} , unique up to multiplication by a positive scalar, such that for all $p, q \in \Delta^{2^X}$,

$$p \succeq q \Leftrightarrow \sum_{A \subseteq X} p(\{S : S \cap A \neq \emptyset\}) \cdot \lambda_A \geq \sum_{A \subseteq X} q(\{S : S \cap A \neq \emptyset\}) \cdot \lambda_A.$$

The proof of the equivalence of (i) and (ii) parallels the proof of Theorem 2 in Nehring (1999a) and is therefore omitted here. The representation in part (iii) can be paraphrased as follows: p is weakly preferred to q if and only if p realizes more valuable attributes than q in expectation. The equivalence of (ii) and (iii) is straightforward in view of Fact 2.1.

W98 developed a cost-benefit calculus for the case of independent survival probabilities π_x of species $x \in X$, i.e. with $p_S = \prod_{x \in S} \pi_x$. Consider an increase in the survival probability of species y by ϵ , while those of other species remain unchanged. By (iii) in the above theorem, the resulting gain in expected diversity is given by

$$\epsilon \cdot \sum_{A: y \in A} \left(\prod_{x \in A \setminus \{y\}} (1 - \pi_x) \right) \cdot \lambda_A.$$

Note that the weight of any attribute possessed by y is multiplied by the probability that no other species possessing the attribute survives. In many cases, including almost all models studied below, the summation ranges over no more than $(\#X)^2$ non-zero summands, thus ensuring practical computability.

2.5 Other Interpretations

The multi-attribute framework allows for a variety of interpretations. Besides the interpretation in terms of biodiversity, the following two seem to be natural and important from an economic perspective.

2.5.1 Option Value

A frequently cited rationale for valuing biodiversity is derived from uncertainty concerning which species possess specific instrumentally desirable features. This is often referred to as the “option value” of biodiversity. To see how this can be modelled

using the multi-attribute approach, assume that society cares for the existence of some specified natural drug. For simplicity, assume that the utility of a certain set S of existing species is 1 if there is at least one species $x \in S$ that carries the drug, and 0 otherwise. Moreover, while the experts are uncertain about which species (of plants) carry the drug they agree on the relevant probabilities. Any subset A of species can be interpreted as the “event” that exactly the species in A carry the drug. Denoting by λ_A the probability of event A , the probability that some species in S will carry the drug is given by

$$v(S) = \sum_{A: A \cap S \neq \emptyset} \lambda_A.$$

Clearly, for any probability distribution λ , the function v is a diversity function. By Theorem 2.1, a preference relation over set-lotteries satisfying POS uniquely reveals, for each event A , the subjective probability λ_A of A (conditional on the event that some species carries the drug).¹⁶

2.5.2 Sociodiversity

The steady growth of cultural diversity in contemporary societies has provoked a lively debate on the entailed benefits and costs both to individual organizations and society as a whole. Our methodology promises to provide useful tools for modelling both costs and benefits of social diversity. A benefit of diversity derives from the fact that different people tend to have different ideas, and thus tend to be more productive as teams. This can be modelled along the lines of the option value interpretation with λ_A as the (subjective) probability that exactly the group A of agents is able to come up with a solution to some problem; the probability that someone in S is able to solve the problem is given by $v(S) = \lambda(\{A : A \cap S \neq \emptyset\})$ as above.

On the other hand, a significant part of the costs of diversity consists in the greater need for and greater difficulty of communication.¹⁷ In the following, we take these costs as being determined by the heterogeneity of experience of groups of agents. An individual is described by a set of formative experiences; the heterogeneity of a group of agents will be measured by the overall range of experiences of members of the group.¹⁸ In the language developed above, a formative experience thus corresponds to a feature of a person; an attribute, therefore, is a set of individuals sharing the same formative experience. Attribute weights reflect the relative importance of specific experiences in defining someone's identity.

Using a metaphor from W98, one may refer to a formative experience as a book read (or, alternatively, a web page hit). The diversity of a set of individuals is then given by the number of different books that have shaped its members; a group of individuals, and even a single individual can thus be viewed as the library of these books. With even greater metaphorical license, the diversity of a group can be viewed as the number of books that have to be read by an outsider in order to fully understand every member of the group. Each book can be viewed as a *meme* in the felicitous expression of

¹⁶Due to the assumption $\lambda_\emptyset = 0$ in Fact 2.1, only this conditional probability is revealed.

¹⁷An equilibrium model for cross-cultural teams that emphasizes the tradeoff between productivity benefits and communication costs of teams can be found in Lazear (1999).

¹⁸For a complementary approach to the measurement of group heterogeneity, see Esteban and Ray (1994). While these authors conceptualize social polarization in terms of an average of socio-psychological distances among pairs of individuals, our approach focuses “holistically” on the coherence of a group as a whole.

R. Dawkins. Note that *counting* the number of books is only justified for “atomic” (genuinely different) books.

2.5.3 Valuing and Measuring Diversity

In the sequel, we repeatedly invoke and contrast the bio- and sociodiversity interpretations. We distinguish between the two not primarily because of a putative fundamental difference of the bio- and the socio-spheres in terms of diversity. Rather, the distinction reflects the difference between *valuing* diversity (in the sense of weighting attributes) and *measuring* diversity (in the sense of counting books). On the former interpretation, diversity theory is a branch of utility theory (with sets as consequences), on the latter a branch of measure/probability theory (where points have the structure of sets). The measurement interpretation has the advantage of greater objective reality, albeit at a highly metaphorical level. On the other hand, since valuing can be operationalized in terms of preferences (i.e. hypothetical choices), the value interpretation has the advantage of greater rigor.

While we have sited “valuing” in the biological and “measuring” in the sociological context in the present exposition, the opposite pairings are also meaningful. Indeed, there is an inherent interplay between the two dimensions: valuing presumably reflects in part measured “objective” dissimilarity and distinctiveness; on the other hand, measuring will typically involve subjective judgements of the relevance of factual differences. We thus view the value and the measurement interpretations as complementary.

3 Paradigmatic Models

While conceptually compelling, the multi-attribute framework is not readily applicable in its unconstrained form since with n objects there are as many as $2^n - 1$ potentially relevant attributes. Depending on the specific context, particular applications will often exhibit additional structure that one may wish to incorporate. A natural way to represent such additional structure within the multi-attribute framework is through appropriate restrictions on the family of relevant attributes. In this section, we study three basic structural restrictions which promise to be particularly useful in applications. The resulting models are referred to as hierarchies, lines and hypercubes, respectively. In hierarchies and lines, diversity can be viewed as aggregate dissimilarity, but not in hypercubes.

3.1 Hierarchies

A concrete starting point for thinking about diversity is to try to derive the diversity of arbitrary sets from the pairwise dissimilarities between its elements. This, in a nutshell, is the program of W92 who demonstrated its success in the context of taxonomic hierarchies in the uniform case. Specifically, he proposed and studied the recursion formula (3.2) below, and extended it to the non-uniform case in W98. In this section, we will see that the recursion formula applies *only* in the hierarchical case; the larger program, however, can be realized under weaker, but still restrictive assumptions, as we will show later.

A family $\mathcal{A} \subseteq 2^X$ of sets is called a (*taxonomic*) *hierarchy* if the elements $A \in \mathcal{A}$ are nested in the following sense.

Nesting Property For all $A, B \in \mathcal{A}$,

$$A \cap B \neq \emptyset \Rightarrow [A \subseteq B \text{ or } B \subseteq A].$$

Accordingly, we will refer to a diversity function v , as well as to the associated attribute weighting function λ , as *hierarchical* if the support Λ of relevant attributes has the Nesting Property. The hierarchical case represents the simplest interesting class of diversity functions and will therefore serve as an important benchmark model. It is of evident interest in the context of biodiversity, both under an evolutionary and under a properly taxonomic interpretation.

The hierarchical model is characterized by the following condition on the underlying preference relation over set-lotteries. For all $\pi \in [0, 1]$, $x \in X$, $S \in 2^X \setminus \{\emptyset\}$, $q, r \in \Delta^{2^X}$,

$$\begin{aligned} [\pi \cdot \mathbf{1}_{\{x, y\}} + (1 - \pi)q \succeq \pi \cdot \mathbf{1}_{\{y\}} + (1 - \pi)r \text{ for all } y \in S] \\ \Rightarrow \pi \cdot \mathbf{1}_{\{x\} \cup S} + (1 - \pi)q \succeq \pi \cdot \mathbf{1}_S + (1 - \pi)r. \end{aligned} \quad (3.1)$$

Note that for this condition to have a bite it must be that $r \succ q$, since otherwise it is implied by vNM and POS. Condition (3.1) can be paraphrased as follows: If the disadvantage of getting q instead of r (with probability $1 - \pi$) is outweighed by a chance of having x in addition to $\{y\}$, for all $y \in S$, then it must also be outweighed by the chance of having x in addition to S itself.

Theorem 3.1 *Let \succeq satisfy vNM and POS, and let $v : 2^X \rightarrow \mathbf{R}$ be the corresponding diversity function with attribute weighting function λ . The following statements are equivalent.*

- (i) Λ has the Nesting Property.
- (ii) \succeq satisfies (3.1).
- (iii) For all x and S ,

$$v(S \cup \{x\}) - v(S) = \min_{y \in S} [v(\{x, y\}) - v(\{y\})]. \quad (3.2)$$

Of particular interest is the formula (3.2) which allows one to compute the diversity function v recursively from its values on the set $\mathcal{B}(X) := \{W : 1 \leq \#W \leq 2\}$ of all non-empty subsets of X containing at most two elements. We refer to $\mathcal{B}(X)$ as the family of *binary sets* with the understanding that these include all singletons as well. The recursion formula can be rewritten suggestively in terms of distinctiveness as follows. For all x, S ,

$$d(x, S) = \min_{y \in S} d(x, y). \quad (3.2')$$

From a purely mathematical point of view, this is the natural way to define a point-set metric from a metric between objects, namely the familiar Hausdorff metric. By Theorem 3.1, however, the validity of (3.2') is restricted to the special case of hierarchies.¹⁹

In view of the recursion formula, Noah's local problem simplifies dramatically in the hierarchical case: sacrifice the species that is least dissimilar from its closest neighbour among the remaining species. Formally, for a given set S and $y_1, y_2 \in S$, let z_1 and z_2 be the closest neighbours of y_1 and y_2 , respectively, i.e. $z_i := \arg \min_{z \in S \setminus \{y_i\}} d(y_i, z)$. By

¹⁹W98 uses (3.2') to *define* the distance of a point to a set in general, and notes the equality of his definition with marginal diversity in the hierarchical case.

(3.2'), the loss in diversity by giving up y_i is then simply $d(y_i, z_i)$. On the other hand, Theorem 3.1 also shows that the “closest neighbor” criterion typically fails outside the hierarchy model.

Although particularly transparent and well-behaved, the hierarchy model is clearly highly restrictive, as has already been observed in Section 2.1. This is also reflected in the properties of the underlying dissimilarity metric. In particular, it is well-known that in the uniform case (equal valuation of singletons) the recursion (3.2') holds if and only if d is an *ultrametric*, i.e. if and only if the two greatest distances between any three points are equal (see, e.g. Johnson (1967)).

3.2 Lines

In many applications, one will want to exploit pre-existing geometric structure of the object space. As a simple but fundamental example of geometric structure, consider a line, i.e. the case in which the universe X is ordered by a given linear (i.e. complete, transitive and antisymmetric) ordering \geq according to some one-dimensional quality. Plausible biological examples are mass, maximal speed, annual number of offspring, life expectancy, etc.²⁰

In order to account for the prespecified geometric structure, one needs a concept of compatibility with that structure. For the line, a natural compatibility requirement is that all relevant attributes be *intervals* in the ordering \geq . Specifically, denote by \mathcal{L} the family of all intervals with respect to \geq . We will impose the following condition on a diversity function $v : 2^X \rightarrow \mathbf{R}$ with attribute weighting function λ .

Line Compatibility $\Lambda \subseteq \mathcal{L}$.

Examples of attributes in \mathcal{L} are “being able to run at least 30 mph,” “weighing between 1 and 10 grams.” By contrast, the conceivable attribute “having an odd number of offspring” is excluded as a relevant attribute by Line Compatibility. The family \mathcal{L} of all intervals with respect to the ordering \geq is referred to as the *line model*. Observe that Line Compatibility does not require *all* intervals to be relevant attributes. For instance, to any hierarchical attribute family \mathcal{A} one can associate a (non-unique) linear ordering such that all elements of \mathcal{A} are intervals with respect to that ordering. The line model is more general in that it allows attributes to (non-trivially) overlap. For instance, the augmented attribute family $\{\{rh, wh\}, \{wh, sh\}, \{wh\}, \{rh\}, \{sh\}\}$ in Example 1' above is compatible with the linear ordering $rh < wh < sh$ (cf. Figure 1 above).

3.2.1 Characterizing the Line Model

The key to characterizing diversity in the line model is the following observation. Consider any set $\{y_1, \dots, y_m\}$ with the elements ordered such that $y_1 < y_2 < \dots < y_m$. The distinctiveness $d(y_m, \{y_1, \dots, y_{m-1}\})$ of the right endpoint y_m from the remaining elements is simply given by its dissimilarity $d(y_m, y_{m-1})$ from its closest neighbor. Indeed, any interval containing y_m intersects $\{y_1, \dots, y_{m-1}\}$ if and only if it contains y_{m-1} . By induction, one thus obtains the following formula which we henceforth refer to as the

²⁰In many applications more than one quality dimension will be of interest. This can be incorporated with the help of an appropriate notion of the product of lines, as shown in Nehring (1999b).

Line Equation. For all $y_1 < y_2 < \dots < y_m$,

$$v(\{y_1, \dots, y_m\}) = v(\{y_1\}) + \sum_{i=2}^m d(y_i, y_{i-1}). \quad (3.3)$$

In terms of the underlying preference relation \succeq over set-lotteries, the Line Equation (3.3) translates into the following “risk-neutrality” condition. For all x, y, z and all S with $y \in S$,

$$x \leq y \leq z \Rightarrow \left[\frac{1}{2} \cdot \mathbf{1}_{S \cup \{x\}} + \frac{1}{2} \cdot \mathbf{1}_{S \cup \{z\}} \right] \sim \left[\frac{1}{2} \cdot \mathbf{1}_S + \frac{1}{2} \cdot \mathbf{1}_{S \cup \{x, z\}} \right]. \quad (3.4)$$

The following result shows that (3.3) and (3.4) are indeed equivalent, and that either condition characterizes the line model.

Theorem 3.2 *Let \succeq satisfy vNM and POS, and let $v : 2^X \rightarrow \mathbf{R}$ be the corresponding diversity function with attribute weighting function λ . The following statements are equivalent.*

- (i) Λ satisfies Line Compatibility.
- (ii) \succeq satisfies (3.4).
- (iii) v satisfies the Line Equation (3.3).

Note that by the Line Equation, a diversity function on a line is determined by its values on the family $\mathcal{B}(X)$ of all binary sets, as in the hierarchical case.

3.2.2 Noah’s Local Problem on a Line

Consider a set $S = \{w, x, y, z\}$ with $w < x < y < z$ (see Figure 2). Assume that species x and z are safe, while species w and y are endangered, and that in fact one of them has to be sacrificed. Which one of w and y should Noah choose?

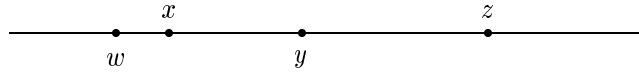


Figure 2: Noah’s local problem in a line

Using the Line Equation (3.3), one easily verifies that the loss in diversity is given by $v(S) - v(S \setminus \{w\}) = d(w, x)$ and

$$v(S) - v(S \setminus \{y\}) = d(z, y) + d(y, x) - d(z, x), \quad (3.5)$$

respectively. Observe that the right-hand side of (3.5) is zero whenever the dissimilarity metric d is *additive* in the following sense. For all x, y, z ,

$$x \leq y \leq z \Rightarrow d(z, x) = d(z, y) + d(y, x). \quad (3.6)$$

In particular, if d were additive, Noah would sacrifice y no matter how “close” w is to x . As a general rule, this seems implausible. Indeed, additivity implies that interior elements have zero marginal diversity value as follows.

Proposition 3.1 *Let $X = \{x_1, \dots, x_n\}$ with $x_1 < x_2 < \dots < x_n$, and let $v : 2^X \rightarrow \mathbf{R}$ be a diversity function such that the support Λ of the corresponding attribute weighting function satisfies Line Compatibility. The following statements are equivalent.*

- (i) *The associated dissimilarity function d is additive, i.e. satisfies (3.6).*
- (ii) *Every $A \in \Lambda$ contains at least one of the endpoints, x_1 or x_n , of $X = \{x_1, \dots, x_n\}$.*
- (iii) *The diversity of any set equals the diversity of its two endpoints, i.e. for all $y_1 < \dots < y_m$, $v(\{y_1, \dots, y_m\}) = v(\{y_1, y_m\})$.*

Two general lessons can be gleaned from this discussion. First, rather than being additive, dissimilarity can in general only be expected to be *subadditive*. Indeed, in view of (3.5), Noah’s choice between giving up the outer species w or the middle species y precisely depends on the “extent of subadditivity” of the dissimilarity metric d (at (x, y, z)). By contrast, “objective” measures of distance such as difference in log mass, for example, are additive. Hence, dissimilarity cannot be equated with objective distance. This suggests that judgements about diversity and similarity contain an irreducibly subjective element.

Secondly, the above analysis reveals that the distinctiveness $d(y, S)$ of an object y to a set S is in general *not* determined by the dissimilarities $d(y, z)$ of y to any $z \in S$. For instance, in the above example $d(y, \{w, x, z\})$ also depends on $d(z, x)$. The dependence of $d(y, S)$ on dissimilarities within the set S shows that the relationship between (marginal) diversity and dissimilarity is subtle and complex even in simple models such as the line model. The general analysis of this relationship is the central theme of Sections 4 and 6 below.

3.3 The Hypercube

In this subsection, we consider the simplest multi-dimensional model in which objects are described as vectors of binary characteristics. Let $K = \{1, \dots, k\}$, and consider the set $X = \{0, 1\}^K$ of binary sequences of length k , which we refer to as the (k -dimensional) *hypercube*. The intended interpretation is that an element $x = (x^1, \dots, x^k) \in \{0, 1\}^K$ corresponds to a particular combination of k basic properties with $x^j = 1$ ($x^j = 0$) denoting possession (non-possession) of property $j \in K$. In the context of biodiversity, a natural interpretation of x is as a genome. Alternatively, the coordinates may correspond to physical or behavioral properties. It will often be useful to describe e.g. the size or the maximal speed of motion of species in first approximation by the opposites large vs. small, slow vs. fast, etc.

In the context of sociodiversity, an example of the hypercube structure is the Myers-Briggs typology of psychological orientations that is used, among other things, to facilitate understanding and communication in the workplace (see, e.g., Keirse and Bates (1984)). It is based on a classification according to the following four characteristics: introversion ($x^1 = 1$) vs. extroversion ($x^1 = 0$), intuition ($x^2 = 1$) vs. sensation ($x^2 = 0$), thinking ($x^3 = 1$) vs. feeling ($x^3 = 0$), and perception ($x^4 = 1$) vs. judgement ($x^4 = 0$). Hence, e.g. the vector $(1, 0, 1, 0)$ corresponds to an introvert, fact-oriented (“sensation”) person guided by thought (rather than feeling) and normative judgement.

A natural notion of compatibility of a diversity function with the hypercube structure is the requirement that all relevant attributes be *sub(-hyper)cubes*. Specifically,

denoting by

$$\mathcal{C} := \{A = \{0\}^{K_1} \times \{1\}^{K_2} \times \{0, 1\}^{K_3} : K_i \cap K_j = \emptyset \text{ and } K_1 \cup K_2 \cup K_3 = K\}$$

the family of all subcubes of the hypercube, we will impose the following condition on a diversity function $v : 2^{\{0,1\}^K} \rightarrow \mathbf{R}$ with attribute weighting function λ .

Hypercube Compatibility $\Lambda \subseteq \mathcal{C}$.

Clearly, the grand cube $\{0, 1\}^K$ corresponds to the universal attribute. The family of all $(k - 1)$ -dimensional subcubes corresponds to the k basic properties and their negations; for instance, the subcube $A = \{1\}^{\{j\}} \times \{0, 1\}^{K \setminus \{j\}}$ describes the set of all objects that possess property $j \in K$. Subcubes of lower dimension describe *conjunctions* of the basic properties (and their negation); for instance, the $(k - 2)$ -dimensional subcube $A = \{1\}^{\{j\}} \times \{0\}^{\{l\}} \times \{0, 1\}^{K \setminus \{j, l\}}$ identifies the set of objects that possess property j but not property l . In general, a $(k - m)$ -dimensional subcube describes conjunctions of m basic properties. If the elements of X are pictured as genomes, subcubes intuitively correspond to genetically determined properties. In the four-dimensional Myers-Briggs typology, a prominent role is played by the basic characteristics themselves, i.e. three-dimensional subcubes that fix only one coordinate, by maximal conjunctions of characteristics, i.e. zero-dimensional subcubes that fix every coordinate, and by certain conjunctions of two characteristics, i.e. two-dimensional subcubes fixing exactly two coordinates (cf. Keirse and Bates (1984)).

The hypercube model differs from the line and hierarchical models in a fundamental respect: unlike in the hierarchy and line models, pairwise dissimilarities between objects (plus their values as singletons) fail to determine the diversity of their set in the hypercube, as shown by the following example.

Example 2 (Insufficiency of Binary Information) Consider the three-dimensional hypercube $X = \{0, 1\}^3$ and the subset $S = \{(1, 1, 0), (1, 0, 1), (0, 1, 1)\}$ (see Figure 3).

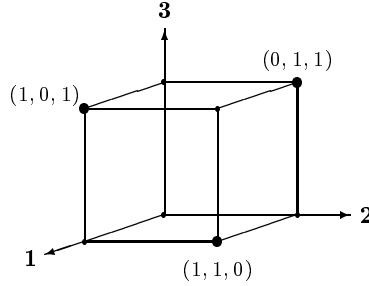


Figure 3: The 3-hypercube

Construct two diversity functions v_1 and v_2 satisfying Hypercube Compatibility from their respective attribute weighting functions λ_1 and λ_2 as follows. Set λ_1 equal to zero except at the grand cube X , for which $\lambda_1(X) = 1/3$, and at all singletons $\{x\}$ (zero-dimensional subcubes), for which $\lambda_1(\{x\}) = 2/3$. Furthermore, set λ_2 equal to zero except for all “faces” (two-dimensional subcubes) A , for which $\lambda_2(A) = 1/3$. One has, for all $x, y \in S$ with $x \neq y$,

$$\begin{aligned} v_1(\{x\}) &= \lambda_1(\{x\}) + \lambda_1(X) &= \frac{2}{3} + \frac{1}{3} &= 1, \\ v_1(\{x, y\}) &= \lambda_1(\{x\}) + \lambda_1(\{y\}) + \lambda_1(X) &= \frac{2}{3} + \frac{2}{3} + \frac{1}{3} &= \frac{5}{3}. \end{aligned}$$

To compute the corresponding values for v_2 , observe that any $x \in S$ is contained in exactly three faces, and that any pair $\{x, y\} \subseteq S$ intersects with five different faces; hence, for all $x, y \in S$ with $x \neq y$,

$$\begin{aligned} v_2(\{x\}) &= \sum_{A \ni x} \lambda_2(A) = 3 \cdot \frac{1}{3} = 1, \\ v_2(\{x, y\}) &= \sum_{A \cap \{x, y\} \neq \emptyset} \lambda_2(A) = 5 \cdot \frac{1}{3} = \frac{5}{3}. \end{aligned}$$

Thus, v_1 and v_2 agree on all one- and two-element subsets of S , so that the value of singletons and the induced dissimilarities within S coincide. Yet, v_1 and v_2 give different diversity value to the set S itself. Indeed, $v_1(S) = 3 \cdot \frac{2}{3} + \frac{1}{3} = \frac{7}{3}$, but $v_2(S) = 6 \cdot \frac{1}{3} = 2$.

4 Diversity Determined by Binary Information

In the previous section, we have observed a structural difference between the hypercube on the one hand and the line and hierarchy models on the other. From the Line Equation (3.3) it follows that the diversity value $v(S)$ of any subset S of the line is determined by the dissimilarities within S (together with the values of all singletons in S). As we have just seen, this is no longer true for the 3-hypercube. This suggests the following question: under what condition on a model is binary information (value on singletons plus dissimilarity) on an arbitrary subset S sufficient to determine the diversity of S ? To provide an answer, think of a “model” as a family $\mathcal{A} \subseteq 2^X$ of *potentially* relevant attributes. A diversity function v is *compatible* with a model \mathcal{A} if $\Lambda \subseteq \mathcal{A}$, i.e. if the induced family of relevant attributes is contained in \mathcal{A} . Note that the inclusion can be strict, that is, a model only gives an upper bound on the family of actually relevant attributes. The line model corresponds to $\mathcal{A} = \mathcal{L}$, the hypercube model to $\mathcal{A} = \mathcal{C}$.

Definition Say that a model \mathcal{A} is *binarily determinate*, henceforth: **bideterminate**, if it satisfies the following condition. For all diversity functions $v_1, v_2 : 2^X \rightarrow \mathbf{R}$ with $\Lambda_1, \Lambda_2 \subseteq \mathcal{A}$, and all S ,

$$v_1^{\mathcal{B}(S)} = v_2^{\mathcal{B}(S)} \Rightarrow v_1(S) = v_2(S),$$

where $v^{\mathcal{B}(S)}$ denotes the restriction of v to the family $\mathcal{B}(S)$ of all binary subsets of S with one or two elements.²¹

Bideterminacy is introduced as a property of a model; this reflects the fact that, evidently, the binary information cannot determine the diversity of arbitrary sets *all by itself*, as many extensions to a diversity function are mathematically possible. A particular model restricts the set of possible extensions by specifying which attributes are admissible; it therefore represents additional *qualitative* information.

A simple necessary condition for bideterminacy can be derived from Example 2 above. Consider again the triple $S = \{(1, 1, 0), (1, 0, 1), (0, 1, 1)\}$ in the 3-hypercube. The key feature of the example is that, for any non-empty subset $W \subseteq S$, there is an attribute (subcube) that, among the elements of S , is possessed exactly by the objects in W ; formally, with \mathcal{A} as the family of all subcubes,

$$\text{for all non-empty } W \subseteq S \text{ there exists } A \in \mathcal{A} \text{ such that } A \cap S = W. \quad (4.1)$$

²¹Observe that $v^{\mathcal{B}(S)}$ specifies both the dissimilarities within S and the value of all singletons in S ; conversely, the restriction $d|_S$ of d to S and the values $v(\{x\})$ for all $x \in S$ jointly determine $v^{\mathcal{B}(S)}$.

Indeed, each singleton in S is distinguished by a zero-dimensional subcube, and the set S itself is obtained as the intersection of S with the grand cube; furthermore, for each pair there is a two-dimensional subcube (a “face”) that distinguishes the given pair from the third element of S . For instance, $(1, 1, 0)$ and $(1, 0, 1)$ but not $(0, 1, 1)$ are contained in the front face of the cube in Fig. 3 above. In general, say that a **triple** S is **free** in a model \mathcal{A} , if S and \mathcal{A} satisfy (4.1). Thus, a triple S is free in \mathcal{A} if for any subset $W \subseteq S$, there is an attribute in \mathcal{A} possessed by all elements of W but by no other object in S . Bideterminacy in a free triple would require solving six linear equations – as in Example 2 – for seven unknowns, which obviously is not possible. More generally, any free triple S in an arbitrary model \mathcal{A} allows one to construct, in analogy to Example 2, two diversity functions v_1 and v_2 such that $v_1^{B(S)} = v_2^{B(S)}$ but $v_1(S) \neq v_2(S)$.²² Bideterminacy thus presupposes that there be no free triples. This condition turns out to be sufficient as well.

Theorem 4.1 *A model \mathcal{A} is bideterminate if and only if no triple is free in \mathcal{A} .*

As we have just seen, hypercubes of dimension ≥ 3 admit free triples. On the other hand, all hierarchical models and the line model clearly satisfy the no-free-triple condition. As to the former, the Nesting Property implies that, for any triple S , at most one pair $\{x, y\} \subseteq S$ can be obtained as the intersection of an attribute A with S , since otherwise attributes would non-trivially overlap. As to the latter, Line Compatibility implies that for any linearly ordered triple $x < y < z$, the set $\{x, z\}$ of the two outer points cannot jointly be distinguished from the middle point y , since any interval containing x and z must also contain y . A less obvious example of a bideterminate model is the 2-hypercube.

Example 3 (The 2-Hypercube) Consider the two-dimensional hypercube $X = \{0, 1\}^2 = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$. Any triple $S \subseteq \{0, 1\}^2$ necessarily contains two diagonally opposite elements; by Hypercube Compatibility, such elements cannot share an attribute that is not also possessed by the third element of S . Hence, no triple is free. By Theorem 4.1, the diversity of any subset of the 2-hypercube can thus be determined from binary information alone. An explicit formula, e.g. for total diversity $v(X)$, can be derived as follows. Let x, y, z, w be a consecutive numbering of the four vertices of the 2-hypercube, say in clockwise orientation as depicted in Figure 4. Then, for any diversity function v on $X = \{0, 1\}^2$ satisfying Hypercube Compatibility,

$$v(X) = v(\{x\}) - d(x, z) + d(x, y) + d(y, z) + d(z, w) + d(w, x). \quad (4.2)$$

For verification, note that the first term $v(\{x\}) - d(x, z)$ gives the weight λ_X of the universal attribute; each of the following four dissimilarities quantifies the weight of one singleton plus the weight of one adjacent “edge” (one-dimensional subcube).

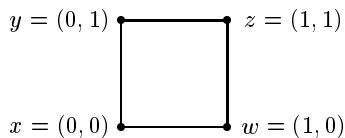


Figure 4: The 2-hypercube

²²Specifically, for any $W \subseteq S$, choose an attribute $A^W \in \mathcal{A}$ such that $A^W \cap S = W$; for all distinct $x, y \in S$ set: $\lambda_1(\{A^{\{x\}}\}) = 2/3$, $\lambda_1(\{A^S\}) = 1/3$ and $\lambda_1 = 0$ otherwise; similarly, $\lambda_2(\{A^{\{x\}}\}) = 1/3$, $\lambda_2(\{A^{\{x, y\}}\}) = 1/3$ and $\lambda_2 = 0$ otherwise.

A striking feature of formula (4.2) is the negative sign with which one of the dissimilarities enters; it implies that the diversity $v(X)$ can *decrease* as the dissimilarities between elements *increase*.²³ To exclude such somewhat paradoxical behavior, we will study in Section 6 below a condition of “monotonicity in dissimilarity” which requires diversity to be a monotone function of dissimilarity. It will turn out that in this case the diversity of any set S can always be represented as the value of a singleton plus the sum of an appropriate subset of pairwise dissimilarities, generalizing the Line Equation (3.3). Specifically, one can write,

$$v(S) = v(\{w\}) + \sum_{x \in S \setminus \{w\}} d(x, y(x)), \quad (4.3)$$

for some $w \in S$ and appropriately chosen elements $y(x) \neq x$. Note that each $x \in S \setminus \{w\}$ occurs exactly *once* in the sum on the r.h.s. of (4.3). In particular, identifying each pair $\{x, y(x)\}$ with an edge in a graph, formula (4.3) will be seen to define a connected and acyclic graph, i.e. a *tree* on S . Thus, trees, to which we turn now, will play a central role in the analysis of monotonicity in dissimilarity.

5 Trees

In this section, we study a natural generalization of the line model, the tree model. An application in the context of biodiversity is to evolutionary trees. The results of this section also serve as an important step towards the analysis of monotonicity in dissimilarity provided in Section 6 below.

A *graph* γ on a set $S \subseteq X$ is a symmetric binary relation on S ; the elements of γ are referred to as the edges of the graph. A *tree* is a connected and acyclic graph. Given a tree τ on X , denote by \mathcal{T} the family of all τ -connected subsets of X , where $A \subseteq X$ is τ -connected if for all $x, y \in A$ there exists a path in τ that connects x and y and that lies entirely in A . The *tree model* associated with a tree τ on X requires all potentially relevant attributes to be τ -connected, i.e. the tree model imposes the following condition on a diversity function $v : 2^X \rightarrow \mathbf{R}$.

Tree Compatibility $\Lambda \subseteq \mathcal{T}$.

Important examples in the context of biodiversity are evolutionary trees. Here, Tree Compatibility can be motivated as follows. Let X be the set of all species that ever existed, and consider a rooted tree τ_{ev} on X with the interpretation that y is an ancestor of x if y is on the path from the root x_0 to x (see Figure 5 in which w is an ancestor of y which, in turn, is an ancestor of x). Denote by \mathcal{T}_{ev} the family of all τ_{ev} -connected subsets. Compatibility with the evolutionary tree ($\Lambda \subseteq \mathcal{T}_{ev}$) amounts to the following two conditions on relevant attributes. First, any two species sharing an attribute must descend from a common ancestor with the same attribute. Secondly, an attribute once left behind by an evolutionary lineage is never recovered. Both are natural requirements under a genealogical interpretation of attributes; for specific examples of connected attributes, see Subsection 5.2 below.

²³One needs to verify that $d(x, z)$ can be increased without altering the other terms in formula (4.2); the proof of Proposition 6.3 below shows how this is done.

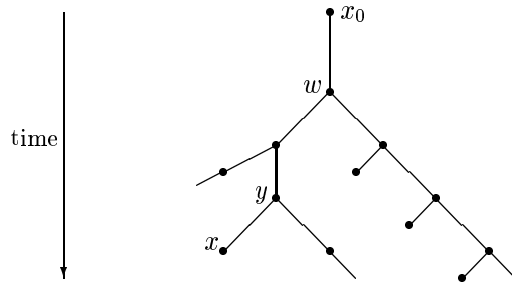


Figure 5: An evolutionary tree

5.1 Characterizing the Tree Model

In the tree model, total diversity $v(X)$ is given by the length of the underlying tree τ as follows. Fix any $w \in X$, and denote by η_w the set of all “edges directed away” from w , i.e. $\eta_w := \{(x, y) \in \tau : y \text{ lies on any path from } w \text{ to } x\}$ (see Fig. 5 for the required location of x and y relative to w). Given a diversity function v on X , the *length* of τ is defined as

$$l_d(\tau) := v(\{w\}) + \sum_{(x,y) \in \eta_w} d(x,y). \quad (5.1)$$

That the r.h.s. of (5.1) does indeed not depend on the “starting point” w is easily verified using the equality $d(x, y) + v(\{y\}) = v(\{x, y\}) = d(y, x) + v(\{x\})$.

Proposition 5.1 *A diversity function $v : 2^X \rightarrow \mathbf{R}$ is compatible with a tree τ on X (i.e. $\Lambda \subseteq \mathcal{T}$) if and only if $v(X) = l_d(\tau)$.*

Note that in case of a line, this result yields the Line Equation (3.3) applied to $v(X)$. However, in contrast to the line model, the characterization of diversity does not carry over to arbitrary subsets. Indeed, the tree model is typically not even bideterminate. The canonical example is the following.

Example 4 (The 3-Star Tree) Consider the “3-star tree” τ on $X = \{x, y_1, y_2, y_3\}$ that connects the “center” x with each y_i , $i = 1, 2, 3$, by an edge (see Figure 6).

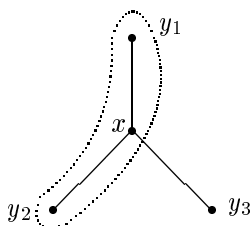


Figure 6: The 3-star tree

Clearly, the set $S = \{y_1, y_2, y_3\}$ is a free triple of \mathcal{T} since any subset $W \subseteq S$ can be obtained as the intersection with S of a τ -connected set $A \subseteq X$ (for instance, the subset $\{y_1, y_2\}$ is the intersection with S of the dotted attribute in Fig. 6); hence, by Theorem

4.1, \mathcal{T} is not bideterminate. In particular, the diversity $v(S)$ cannot be expressed in a way similar to formula (5.1) above.

Notice that the line is distinguished among all trees by the property that each element is connected to at most two other elements by an edge. In particular, *any* tree that is not a line contains a 3-star tree as a subtree; thus, in view of Example 4, the line model is the only bideterminate tree model.

5.2 From Evolutionary to Phylogenetic Trees

From a modelling perspective, one may simply accept the underdeterminacy of the tree model. However, under an evolutionary interpretation this is unsatisfactory since the tree model \mathcal{T}_{ev} yields no similarity restrictions on the set $Y \subseteq X$ of *terminal nodes*, in particular no such restrictions on the subset of currently existing species. Indeed, as exemplified by the 3-star tree above, any subset $W \subseteq Y$ of terminal nodes can be obtained as the intersection of some attribute $A \in \mathcal{T}_{ev}$ with Y . In other words, for *any* subset W of currently existing species there is a potential attribute that is possessed by all species in W but by no other currently existing species. But one would expect the evolutionary tree to deliver substantial similarity implications also for the subset of currently existing species. Such implications can be obtained by selecting appropriate *subfamilies* of \mathcal{T}_{ev} as the set of relevant attributes. The proposal made in W92 and W98 can be viewed as an example of this strategy. Given an evolutionary tree τ_{ev} , any species x defines an attribute $A_x := \{y : x \text{ is an ancestor of } y\}$; in biological terminology, the set A_x of species is the “clade” founded by x . Weitzman selects the family $\mathcal{H}_{ev} := \{A_x : x \in X\}$ of all such clades as the family of relevant attributes. As is easily verified, the subfamily $\mathcal{H}_{ev} \subseteq \mathcal{T}_{ev}$ forms a hierarchy, in particular, \mathcal{H}_{ev} is bideterminate. Moreover, by Theorem 3.1, the requirement $\Lambda \subseteq \mathcal{H}_{ev}$ leads exactly to the recursion formula (3.2) proposed by Weitzman.

Consider, however, the following segment of an evolutionary tree:

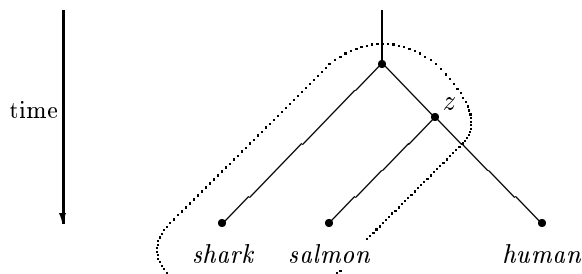


Figure 7: A segment of an evolutionary tree

In Figure 7, z denotes a common ancestor of *salmon* and *human* that is not an ancestor of *shark*. Observe that for any uniform diversity function with $\Lambda \subseteq \mathcal{H}_{ev}$ the set $\{\textit{salmon}, \textit{human}\}$ is (weakly) less diverse than the set $\{\textit{salmon}, \textit{shark}\}$, in other words, $d(\textit{salmon}, \textit{human}) \leq d(\textit{salmon}, \textit{shark})$. It seems counterintuitive that the information about evolutionary descent should *force* this ranking. In the biological literature, Weitzman’s hierarchical model has similarly been criticized for implicitly assuming a constant speed of evolution (see, e.g. Faith (1994)). On this account the evolutionary clock has run faster on the branch from z to *salmon* than on the branch from z to *human*, making *salmon* closer than *human* to *shark*. In response, Faith and Walker (1994) have

proposed a model which selects a (non-hierarchical) subfamily of τ_{ev} -connected subsets as the relevant attributes.

In the following, we propose an alternative resolution of the problem raised by the evolutionary tree in Figure 7. It is based on the observation that the hierarchy \mathcal{H}_{ev} neglects an important similarity between *salmon* and *shark* that derives from their belonging to the same taxon “fish” (corresponding to the dotted subset in Fig. 7). One would thus want to combine the underlying evolutionary tree τ_{ev} with the similarity information contained in the hierarchy \mathcal{H}_{tax} representing taxonomic classification. Since the taxonomic classification of species should be compatible with the genealogy of their descent, it is natural to require $\mathcal{H}_{tax} \subseteq \mathcal{T}_{ev}$, i.e. to assume that any taxon is connected in the evolutionary tree τ_{ev} . Without loss of generality, suppose also that $X \in \mathcal{H}_{tax}$, i.e. that the universal attribute is included in the taxonomic classification. Define the *phylogenetic tree* $\mathcal{T}_{phyl} \subseteq \mathcal{T}_{ev}$ corresponding to τ_{ev} and \mathcal{H}_{tax} as the family of all intersections of attributes in \mathcal{H}_{ev} and \mathcal{H}_{tax} , i.e.,

$$\mathcal{T}_{phyl} := \{A_x \cap A_{tax} : A_x \in \mathcal{H}_{ev}, A_{tax} \in \mathcal{H}_{tax}\}.$$

Observe that since $X \in \mathcal{H}_{tax}$ and $X \in \mathcal{H}_{ev}$, \mathcal{T}_{phyl} contains both hierarchies \mathcal{H}_{ev} and \mathcal{H}_{tax} ; including intersections of attributes is natural as it allows one to assign positive weight to *combinations* of the underlying features. Note that by including all clades $A_x \in \mathcal{H}_{ev}$, the phylogenetic tree model still allows ranking the set $\{\textit{salmon}, \textit{shark}\}$ as more diverse than the set $\{\textit{salmon}, \textit{human}\}$; indeed, salmon and humans share an important common phylogenetic feature, cartilaginous bones, which sharks do not have. This point also illustrates why it would be mistaken to simply replace \mathcal{H}_{ev} by \mathcal{H}_{tax} . One can show that the class of phylogenetic tree models is strictly more general than the line model.²⁴ Nevertheless, one has the following result.

Proposition 5.2 *Any phylogenetic tree is bideterminate.*

The proof in the appendix verifies this by showing that no triple is free in \mathcal{T}_{phyl} ; the assertion thus follows from Theorem 4.1.

6 Diversity as Aggregate Dissimilarity

The concept of bideterminacy does not fully exhaust our intuitive notion of diversity as aggregate dissimilarity. As already noted in Section 4 above, one would also expect some kind of “monotonicity in dissimilarity.”

Definition Say that a model \mathcal{A} is **monotone in dissimilarity** if it satisfies the following condition. For all diversity functions $v_1, v_2 : 2^X \rightarrow \mathbf{R}$ with $\Lambda_1, \Lambda_2 \subseteq \mathcal{A}$, and all S ,

$$\left[v_1^{S(S)} = v_2^{S(S)} \text{ and } d_1|_S \geq d_2|_S \right] \Rightarrow v_1(S) \geq v_2(S),$$

where $v^{S(S)}$ denotes the restriction of v to the family $\mathcal{S}(S) := \{\{x\} : x \in S\}$ of all singletons in S , and $d|_S$ is the restriction of the induced dissimilarity metric to $S \times S$.

²⁴For a proof of this assertion and a deeper analysis of the phylogenetic tree model, see Nehring and Puppe (2001). In particular, it is shown there that phylogenetic trees satisfy the property of “monotonicity in dissimilarity” which is defined in the next section as a natural strengthening of bideterminacy.

Holding the value on singletons fixed, the defined notion of monotonicity thus describes the dependence of diversity on a pure change of dissimilarity.²⁵ Note that, by the Line Equation (3.3), the line model is clearly monotone in dissimilarity. Also observe that, by definition, monotonicity in dissimilarity implies bideterminacy. The converse does not hold as the simple counterexample of the 2-hypercube shows (cf. Example 3 in Section 4 above).

To illustrate the above definition, assume that Noah has quasi-linear preferences over sets of species and money; specifically, interpret $v(S)$ as his willingness to pay for the preservation of all species in S (if otherwise no species survives). Furthermore, assume uniformity (equal valuation of singletons) along with the normalization $v(\{x\}) = 1$, so that single species serve as the numeraire. In particular, $d(x, y)$ can be interpreted as Noah's willingness to pay for the preservation of species x if otherwise only y survives. Consider a situation in which the set of existing species is $Y \subseteq X$ with the subset $S \subseteq Y$ as the set of endangered species. Noah had made up his mind initially about all dissimilarities $d(x, y)$ within X ; he then reconsiders some of the dissimilarities within the set S , finding that in all pairwise comparisons he is willing to pay for each species in S at least as much as before. In other words, Noah's revised assessment $d'(x, y)$ satisfies: $d'(x, y) \geq d(x, y)$ for all $x, y \in X$, and $d'(x, y) = d(x, y)$ whenever $x \notin S$. Then, monotonicity in dissimilarity ensures that his willingness to pay for the preservation of all of S cannot go down, i.e. $v'(Y) - v'(Y \setminus S) \geq v(Y) - v(Y \setminus S)$, for it implies both $v'(Y \setminus S) = v(Y \setminus S)$ and $v'(Y) \geq v(Y)$. Thus, monotonicity in dissimilarity ensures that Noah can derive his valuations of complex sets of species from those about pairs of species. This may be a much easier and more transparent task; for instance, Noah can derive his dissimilarity judgements from evolutionary or genetic distances.

6.1 Characterization of Monotonicity in Dissimilarity

The characterization of the class of models that are monotone in dissimilarity proceeds in three steps. Based on the analysis of the tree model, we first derive a sufficient condition, "tree-consistency" (Subsection 6.1.1). Based on the counterexample of the 2-hypercube, we then formulate a natural necessary condition for monotonicity in dissimilarity, "acyclicity." The third and crucial step is to establish the equivalence of tree-consistency and acyclicity (Subsection 6.1.2). The implications of this result, including the desired characterization of monotonicity in dissimilarity, are described in Subsection 6.1.3.

6.1.1 A Sufficient Condition: Tree-Consistency

By Proposition 5.1, compatibility with a tree model entails that *total* diversity $v(X)$ is given by the length of the underlying tree, and thus a monotone function of the dissimilarities in X . However, it follows from Example 4 that compatibility of v with a tree model on X does *not* imply compatibility of the restriction $v|_S$ to $S \subseteq X$ with a tree model on S . A brute force method to ensure monotonicity in dissimilarity is thus to require compatibility of the restriction $v|_S$ with *some* tree model on S for all subsets $S \subseteq X$, as follows.

²⁵The assumption $v_1^{S(S)} = v_2^{S(S)}$ fixes a common scale for v_1 and v_2 and thus ensures meaningfulness of the condition in terms of the underlying preferences over set-lotteries.

For a model \mathcal{A} on X , denote by $\mathcal{A}|_S := \{A \cap S : A \in \mathcal{A}\}$ the *relativization* of \mathcal{A} to S ; thus, $\mathcal{A}|_S$ describes the potential attributes *among* the elements of S . If v is compatible with \mathcal{A} , then $v|_S$ is compatible with $\mathcal{A}|_S$.²⁶ Say that a model \mathcal{A} is **tree-consistent** if, for all subsets $S \subseteq X$, there exists a tree τ_S on S such that any $A \in \mathcal{A}|_S$ is τ_S -connected. To see what is involved, consider again the subset $S = \{y_1, y_2, y_3\}$ of the 3-star tree in Example 4 above. We have already observed that S is a free triple in the corresponding tree model \mathcal{T} on $X = \{x, y_1, y_2, y_3\}$; this means that every non-empty subset of S is an element of the relativization to S , i.e. $\mathcal{T}|_S = 2^S \setminus \{\emptyset\}$. But this implies that there cannot exist a tree τ_S on S such that any $A \in \mathcal{T}|_S$ is τ_S -connected. Hence, the model \mathcal{T} is not tree-consistent.

A model is tree-consistent exactly when the diversity of *any* set is given by the minimal length of a tree on that set, as shown by the following result.

Proposition 6.1 *Let $v : 2^X \rightarrow \mathbf{R}$ be a diversity function with attribute weighting function λ . Then, Λ is tree-consistent if and only if, for all S ,*

$$v(S) = \min_{\tau} l_d(\tau), \quad (6.1)$$

where the minimum is taken over all trees τ on S .

In the following, we refer to a tree on S that achieves minimal length with respect to d as a *minimum spanning tree* on S .²⁷ Note that the minimum in (6.1) is taken over all trees on S , no matter whether $\Lambda|_S$ is compatible with them.

Clearly, Proposition 6.1 establishes tree-consistency as a sufficient condition for monotonicity in dissimilarity since the length of a tree on S is a monotone function of the dissimilarities within S . Moreover, formula (6.1) expresses diversity without reference to a specific underlying tree. In particular, a tree τ_S with which $\Lambda|_S$ is compatible can be found by length minimization.

Diversity in a tree-consistent model can also be understood in terms of the following condition introduced by Weitzman as the “Link Property” (see W92, p.378).²⁸ For all S with $\#S \geq 2$ there exist distinct $x, y \in S$ such that

$$v(S) - v(S \setminus \{x\}) = d(x, y). \quad (6.2)$$

The Link Property asserts that in each set S there exists an “outer point” x and its “gate” y to $S \setminus \{x\}$ such that the distinctiveness of x from $S \setminus \{x\}$ is simply the dissimilarity of x from y .

Proposition 6.2 *Let $v : 2^X \rightarrow \mathbf{R}$ be a diversity function with attribute weighting function λ . Then, Λ is tree-consistent if and only if v satisfies the Link Property (6.2).*

²⁶Indeed, the attribute weighting function $\lambda|_S$ of $v|_S$ is given by $\lambda|_S(B) = \sum_{A \in \mathcal{A}: A \cap S = B} \lambda_A$, for all non-empty $B \subseteq S$; hence, the family of relevant attributes induced by $v|_S$ on S is $\Lambda|_S$.

²⁷According to the acknowledgement of W92, the proposal to measure diversity in terms of the length of a minimum spanning tree is due to S. Hart. Note that, in contrast to standard usage in combinatorics, the metric underlying a minimum spanning tree in our sense is typically non-additive; also, the trees over which the minimum is taken are restricted to lie in S .

²⁸The Link Property does very different work in Weitzman’s approach. In particular, the framework of W92 does not allow to derive a result similar to Proposition 6.2. The reason is that W92 (in contrast to W98) does not assume diversity to be a submodular function. Indeed, the hierarchical model is the only case in which the set functions considered in W92 can be rationalized in terms of a multi-attribute model.

6.1.2 A Necessary Condition: Acyclicity

A closer examination of the no-free-triple condition of Section 4 points to a necessary condition for monotonicity in dissimilarity. A model \mathcal{A} will be called **regular** if it contains the universal attribute X and the family $\mathcal{S}(X) = \{\{x\} : x \in X\}$ of all singletons; note that both the line and hypercube models are regular. For regular models, the no-free-triple condition is equivalent to the following, otherwise slightly stronger condition. Say that a model \mathcal{A} is *triple acyclic* if in any triple S there exists a pair $\{x, y\} \subseteq S$ such that for no $A \in \mathcal{A}$, $A \cap S = \{x, y\}$. Thus, a model is triple acyclic if there is no triple of objects such that any pair of them possesses an attribute that is not shared by the third object. In Section 4, we have already seen that hierarchical models, the line model and the 2-hypercube are all triple acyclic, while hypercubes of dimension ≥ 3 are not.

Consider again the 2-hypercube as a counterexample to monotonicity in dissimilarity. While triple acyclic, hence bideterminate, the 2-hypercube admits a “cycle” of *four* attributes, namely the four edges (one-dimensional subcubes) (cf. Fig. 4 above). This motivates the following definition.

Let $m \geq 3$; say that a model \mathcal{A} is *m-cyclic* if there exist m distinct elements x_1, \dots, x_m and m attributes $A_1, \dots, A_m \in \mathcal{A}$ such that, for all $i = 1, \dots, m$,

$$A_i \cap \{x_1, \dots, x_m\} = \{x_i, x_{i+1}\}, \quad (6.3)$$

where indices are understood modulo m so that $x_{m+1} = x_1$. A model is called *m-acyclic* if it is not *m-cyclic*; finally, a model is **acyclic** if it is *m-acyclic* for all $m \geq 3$. The 2-hypercube is an example of a 4-cyclic model, the 3-star tree in Example 4 is a 3-cyclic model. Also, observe that the notion of 3-acyclicity just defined coincides with triple acyclicity as defined above. The following result shows that, for regular models, acyclicity is necessary for monotonicity in dissimilarity; for a discussion of the regularity assumption, see Section 6.1.3 below.

Proposition 6.3 *Any regular model that is monotone in dissimilarity must be acyclic.*

Being already violated in the 2-hypercube, acyclicity can be viewed as a condition of *one-dimensionality* of a model. To confirm the appropriateness of this interpretation, consider the following class of models. With any connected graph γ on X associate its *graphical model* \mathcal{M}_γ consisting of all sets A with the property that, if $\{x, y\} \subseteq A$ and z is on a shortest path²⁹ between x and y , then $z \in A$. The class of graphical models is rather flexible; in particular, it contains the line, tree and hypercube models. It is easily seen that a graphical model \mathcal{M}_γ is acyclic if and only if γ is a line, the paradigmatic one-dimensional graph. Indeed, the line model is clearly acyclic; conversely, since \mathcal{M}_γ contains all edges of γ , acyclicity of \mathcal{M}_γ forces γ to be acyclic, i.e. to be a tree. But any tree that is not a line contains a 3-star subtree, and is hence not acyclic.

Note that any *selection* $\mathcal{A}' \subseteq \mathcal{A}$ of an acyclic model \mathcal{A} is acyclic as well. For instance, hierarchical models are acyclic as selections from the line model. On the other hand, not every acyclic model is a selection from the line model. For instance, any phylogenetic tree model can be shown to be acyclic (see Nehring and Puppe (2001)). The following central result entails that any acyclic model is a selection of a tree model.

Theorem 6.1 *A model is acyclic if and only if it is tree-consistent.*

²⁹That is, a path with the minimum number of edges.

Thus, the necessary condition for monotonicity in dissimilarity, acyclicity, is in fact equivalent to the sufficient condition, tree-consistency. While it is easy to show that tree-consistency implies acyclicity, the converse implication is substantially more difficult (see the appendix for the proof and its underlying intuition).³⁰

6.1.3 The Equivalence of Acyclicity and Monotonicity in Dissimilarity

Combining Theorem 6.1 with Propositions 6.1 and 6.3, we obtain our main result.

Theorem 6.2 *An acyclic model is monotone in dissimilarity. Conversely, a regular model that is monotone in dissimilarity must be acyclic.*

In view of formula (6.1) above, Theorem 6.2 has the surprisingly strong consequence that there is essentially only *one* way for deriving diversity from dissimilarity in a monotone fashion, namely as the length of a minimum spanning tree.

Corollary 6.1 *For any regular model that is monotone in dissimilarity, the diversity of any set is given as the length of a minimum spanning tree on that set.*

Both Corollary 6.1 and the necessity of acyclicity for monotonicity in dissimilarity in Theorem 6.2 hinge on the assumption of regularity, i.e. on the requirement that a model \mathcal{A} contains the family $\mathcal{S}(X)$ of all singletons and the universal set X as potentially relevant attributes; in Appendix A, we give an example of an irregular model that is monotone in dissimilarity without being acyclic. However, regularity is arguably a weak requirement since, as an assumption on a model (and not on the support of a diversity function), it only restricts *potential* attributes: a regular model \mathcal{A} *allows* all singletons and the universal attribute to have positive weight.³¹ As already noted, both the line model \mathcal{L} and the hypercube model \mathcal{C} are regular; regularity of hierarchical models can be assumed without loss of generality, since the inclusion of the universal attribute and all singletons preserves the hierarchical structure. The condition $\mathcal{S}(X) \subseteq \mathcal{A}$ states that it must be possible to uniquely distinguish each object by a potential attribute that is not shared with any other object. In particular, the condition is equivalent to the existence of a diversity function v that is compatible with \mathcal{A} and strictly monotone in the sense that $v(W) < v(S)$ whenever W is a proper subset of S .³² The requirement

³⁰Acyclicity has been suggested as *one* way to formally define a notion of “one-dimensionality” of models. Does it make sense to view acyclic, i.e. tree-*consistent*, models as one-dimensional but to deny that status to tree models themselves? The following observation suggests a positive answer. Recall from Section 5.2 that a tree model on X entails no restrictions on the set $Y \subseteq X$ of terminal nodes. That is, *any* model on Y can be obtained as the relativization to Y of a tree model on X . In particular, the model $2^Y \setminus \{\emptyset\}$ of all conceivable attributes on Y can be obtained in this way. But clearly this model cannot be viewed as one-dimensional on Y ; hence, it would seem inappropriate to view the corresponding tree model on X as one-dimensional. While relativizations of tree models need not have any structure, those of tree-consistent models inherit the tree-like structure by definition. Tree-consistency/acyclicity thus satisfies one fundamental requirement on a concept of one-dimensionality, namely that one-dimensionality be preserved under restrictions of the domain.

³¹Note that in Proposition 6.3 and Theorem 6.2, regularity is of interest only as an assumption on a model since monotonicity in dissimilarity is; indeed, the latter can only be formulated for models and not for particular diversity functions.

³²In some contexts, one may wish to model *dominance* of objects in the sense that $v(\{x\}) \geq v(\{y\})$ for *all* diversity functions compatible with some model \mathcal{A} . Clearly, in that case \mathcal{A} cannot contain the singleton attribute $\{y\}$. In the working paper version (Nehring and Puppe (1999a)), we show that in the definition of regularity, the requirement $\mathcal{S}(X) \subseteq \mathcal{A}$ may be replaced by the requirement that \mathcal{A} be closed under intersections. Regularity of a model \mathcal{A} with $X \in \mathcal{A}$ thus in effect reduces

$X \in \mathcal{A}$, on the other hand, ensures that diversity is monotone in the values of singletons (see the example in Appendix A for an illustration).

Theorem 6.2 shows that monotonicity in dissimilarity requires acyclicity. How restrictive is acyclicity? Two natural types of restrictions come to mind: those on the underlying dissimilarity metrics and those on the resulting global behavior of diversity. The first will be described now, the second in Section 7 below.

Just as the Nesting Property of the more special hierarchical model entails strong restrictions on the associated dissimilarity metrics (ultrametricity in the uniform case), so does acyclicity. Specifically, any dissimilarity metric associated with a diversity function with acyclic support must satisfy the following “Metric Link Property.”

Metric Link Property For all S with $\#S \geq 2$, there exist distinct $x, y \in S$ such that, for all $z \in S \setminus \{x\}$,

$$d(x, y) \leq d(x, z) \text{ and } d(z, x) \geq d(z, y). \quad (6.4)$$

That is, any set S with at least two elements contains distinct x and y such that x is closer to y than to any other element of S , and conversely, any other element of S is farther away from x than from y . For verification, observe that by Theorem 6.1, acyclicity of Λ implies the existence of a tree τ_S with which $\Lambda|_S$ is compatible; then simply take in (6.4) any pair x, y such that y is the only other element of S to which x is connected by an edge of τ_S (for instance, the elements x and y in Fig. 5 above have the required properties).

The Metric Link Property (6.4) provides a fairly tight criterion to distinguish diversity functions based on one-dimensional (i.e. acyclic) models from those based on multi-dimensional models. Say that a model \mathcal{A} has a *minimal cycle of length m* if there exists $S \subseteq X$ such that $\mathcal{A}|_S$ is m -cyclic while $\mathcal{A}|_W$ is acyclic for every proper subset $W \subset S$. One has the following result.

Proposition 6.4 *Let v be a diversity function with attribute weighting function λ .*

- (i) *v satisfies the Metric Link Property whenever Λ is acyclic.*
- (ii) *v violates the Metric Link Property whenever Λ has a minimal cycle of length ≥ 4 .*
- (iii) *If all diversity functions compatible with a model \mathcal{A} satisfy the Metric Link Property, then \mathcal{A} is m -acyclic for all $m \geq 4$.*

Minimal cycles of length 4 arise already in the simplest case of structured multi-dimensionality, the hypercube of dimension at least two.³³ On the other hand, in models with “diffuse multi-dimensionality” such as the unconstrained model $\mathcal{A} = 2^X \setminus \{\emptyset\}$, all minimal cycles have length 3; in such models, the Metric Link Property may be satisfied for some attribute weights. By part (iii), it will however be violated for other attribute weights compatible with the same model, unless *all* attribute cycles have length 3; this is an extremely special case.

to the requirement that it *either* contain all singletons *or* be closed under intersections. Intersection-closedness is a natural property in that intersections of attributes correspond to conjunctions of the attribute-defining features.

³³More generally, using the notion of a product of models introduced in Nehring (1999b), one can show that *any* product of models gives rise to minimal cycles of length 4.

6.2 Discussion

Our main result, Theorem 6.2, shows that while diversity can be understood as aggregate dissimilarity in situations significantly more general than those of a taxonomic hierarchy, a “well-orderedness” of the family of relevant attributes remains essential to this possibility. This general condition of well-orderedness, or one-dimensionality, is formally expressed as acyclicity rather than the more demanding Nesting Property. Acyclicity has been shown to yield strong restrictions on the underlying dissimilarity metrics in form of the Metric Link Property; it also imposes significant restrictions on the global quantitative behavior of diversity, as will be shown in Section 7 below. Theorem 6.2 might thus be viewed as an impossibility result with an “Arrovian flavor,” with monotonicity in dissimilarity as an IIA-type condition. As Arrow’s celebrated impossibility theorem, the result shows that coherent aggregation of binary information works only on highly restricted domains. In both cases, aggregation is possible on one-dimensional domains (line viz single-peakedness and tree-like generalizations; see Demange (1982)), while on multi-dimensional domains it fails generically.

That “monotonicity in dissimilarity” should be as restrictive as it turns out to be seems rather remarkable; indeed, it is easy to formulate simple rules that incorporate this property by definition such as, for instance, the minimum spanning tree or W92’s “Link Recursion” formula. However, for metrics without the Metric Link Property, these cannot yield a diversity function or even a submodular function.³⁴ But submodularity is equivalent to distinctiveness $d(x, S)$ being weakly decreasing in the reference set S , and thus represents itself a fundamental requirement of diversity aggregation.

Nonetheless, Theorem 6.2 is far from purely negative in its implications. First of all, it precisely characterizes the domains on which dissimilarity aggregation *is* possible, namely the class of all acyclic models. This class considerably generalizes the hierarchical model. In particular, it contains lines and allows to model diversity based on evolutionary information in a rich way, for instance through phylogenetic trees. Secondly, the result singles out a *unique* way of aggregation, the length of a minimum spanning tree. Finally, the multi-attribute approach already represents a framework in which the informational constraints of relying exclusively on binary dissimilarity information can be overcome.

7 Global Behaviour of Diversity in One- versus Multi-Dimensional Models

In the previous section, we have shown that diversity can be understood as aggregate dissimilarity only when the underlying model is sufficiently “well-ordered,” essentially “one-dimensional.” This, we shall argue now, is not merely a matter of nicety of functional form and analytical convenience, but reveals a fundamental difference in the quantitative behaviour of diversity.

Let $v : 2^X \rightarrow \mathbf{R}$ be a diversity function; throughout this section, we assume uniformity, i.e. constant valuation of singletons, along with the normalization $v(\{x\}) = 1$, for all $x \in X$. Note that, in view of the maintained normalization $v(\emptyset) = 0$, this fixes a

³⁴The former follows from Propositions 6.1 and 6.2, the latter requires a more elaborate argument involving Proposition 4.3 in Nehring and Puppe (1999b), which specifies conditions under which submodularity implies total submodularity.

unique cardinal scale; in particular, the normalizations allow one to attach significance to absolute diversity values.³⁵ Define the *granularity* Γ_S of S as the minimal dissimilarity between two distinct objects in S , i.e. $\Gamma_S := \min_{x,y \in S, x \neq y} d(x,y)$. In view of the minimum spanning tree formula (6.1), it is easily verified that in any acyclic model one has

$$v(S) \geq 1 + (\#S - 1) \cdot \Gamma_S. \quad (7.1)$$

This establishes a remarkable connection between three magnitudes: the granularity, the cardinality, and the diversity of a given set. From the perspective of diversity as aggregate dissimilarity, such a connection may not seem troublesome. However, from the more general perspective of the multi-attribute framework there is no reason why significant granularity must necessarily be accompanied by large diversity of large sets as implied by (7.1). Indeed, the coexistence of significant granularity and comparatively low diversity of large sets is naturally explained by multi-dimensionality: Any two distinct objects x and y differ in some dimensions, i.e. for any given y , x possesses attributes not shared with y ; in particular, granularity can be significant. On the other hand, almost all attributes of an object x in a sufficiently large set S will be shared by *some* member of S . By consequence, the marginal value of x can be very small, and hence the diversity of S relatively low. The following example confirms this intuition.

Example 5 (The k -hypercube) Consider the k -dimensional hypercube. For expository convenience, suppose that only the $(k - 1)$ -dimensional subcubes have positive weight,³⁶ assuming equal weight $1/k$ for all of these. This implies $\Gamma_S = 1/k$ for all S containing two elements that differ in only one coordinate. Diversity in this simple example is given by the following formula. For all S ,

$$v(S) = 1 + \frac{\dim S}{k}, \quad (7.2)$$

where (in self-explanatory notation) $\dim S := \#\{j : \text{proj}_j S = \{0, 1\}\}$. In this example of a non-acyclic model, (7.1) is violated whenever $k \geq 2$. To see this, consider the set $S = X$. By (7.2), $v(X) = 2$; on the other hand, the r.h.s. of (7.1) becomes $1 + (2^k - 1)/k$ which exceeds 2 by orders of magnitudes if k is large. Thus, mistakenly assuming one-dimensionality may lead to gross overestimations of the true diversity of large sets. Note that this conclusion does not hinge on the specific weights chosen in this example; indeed, (7.1) will be violated whenever high-dimensional subcubes have sufficient weight.³⁷ Also observe that as k becomes large, the granularity becomes arbitrarily small, while the cardinality of the set X becomes large very rapidly (exponentially). Our claim that (7.1) is a substantial and potentially problematic restriction on the global behavior of diversity is thus consistent with an intuition that large sets will have small granularity.

³⁵For instance, in a larger context with quasi-linear preferences over sets and money, $v(S)$ can be interpreted as the willingness to pay for preventing the extinction of all of S , as in Section 6 above. Alternatively, diversity numbers can be interpreted in terms of probabilities. For instance, the statement “ $v(S) = \alpha$ ” means that Noah is indifferent between the certainty of saving one single species and the prospect of saving all species in S in the presence of a $(1 - \frac{1}{\alpha})$ -chance of total extinction.

³⁶Recall from Section 3.3 that this means that only the coordinate-defining characteristics themselves but no conjunctions are relevant.

³⁷Giving some weight to all subcubes would eliminate the degenerate feature that any set S containing two opposite elements achieves the same (maximal) diversity $v(S) = 2$, and would also make the model regular.

An important context in which one would like to model the coexistence of significant granularity and relatively low diversity for large sets is sociodiversity. The intuitions motivating multi-dimensionality apply very naturally here: Any two members of a society possess memes not shared with each other; on the other hand, each individual corresponds to a recombination of the same basic characteristics, so that most memes of an individual in a large society will be shared by *some* other individual. For the sake of concreteness, consider the diversity of the population of the largest nation, China ($=S$), in the context of the entire world population ($=X$). We have the following intuitions:

- (i) Any two people (in particular, any two Chinese) differ in a non-negligible way.
- (ii) The world is more diverse than China.
- (iii) In a large population, most of the total diversity can be realized by appropriately choosing a subset of relatively few extremely heterogeneous individuals that are representative for the entire population.

By (ii) and (iii), one can find a comparatively small set of individuals in the world (presumably including some Chinese) that is more diverse than all of China. Under (7.1), however, this easily conflicts with hypothesis (i). To see this, consider the following implication of (7.1) for the ordinal ranking of sets in terms of diversity. For all W, S ,

$$\left[\#S \geq \frac{\#W}{\Gamma_S} \right] \Rightarrow v(S) \geq v(W).^{38} \quad (7.3)$$

Assume that $\Gamma_S = 0.001$, i.e. that any two Chinese differ at least in 0.1% of their memes. Furthermore, it seems plausible that one can find 100,000 individuals in the world constituting a set $W \subseteq X$ that is more diverse than all of China. Considering that China has a population of roughly one billion, these hypotheses contradict (7.3), since $10^9 > \frac{10^5}{10^{-3}}$ by an order of magnitude but at the same time $v(S) < v(W)$. While the choice of numbers is clearly open to debate, our point is that these numbers represent reasonable empirical hypotheses which should not be ruled out *a priori* by adopting the restrictive framework implied by acyclicity.

Just as in the case of the underlying inequality (7.1), there is no reason to expect (7.3) to hold in multi-dimensional models. Indeed, the example of the k -hypercube above also violates (7.3). To verify this, note that for $W = \{(0, \dots, 0), (1, \dots, 1)\}$ and any $(k-1)$ -dimensional subcube S one has $v(W) = 2$ and $v(S) = 2 - (1/k)$, respectively. Thus, the comparatively large set S is dominated in diversity by the appropriately chosen set W of only two extremely heterogeneous objects. In particular, since $\#S = 2^{(k-1)}$, this involves a violation of (7.3) whenever $k \geq 4$.

8 Conclusion: Outlook and Further Applications

This paper has provided a fairly in-depth account of diversity in the one-dimensional case. It has shown the distinction between one- and multi-dimensional models to be a fundamental one for understanding and modelling diversity, and has thereby set a

³⁸For verification, note first that by uniformity, $v(W) \leq \#W$ for all W ; from this and the antecedent in (7.3), $v(S) \geq 1 + (\#S - 1)\Gamma_S \geq 1 + ((\#W/\Gamma_S) - 1)\Gamma_S = 1 + \#W - \Gamma_S \geq v(W)$.

clear task for further research: to develop tractable and empirically relevant multi-dimensional models. This task is addressed in Nehring (1999b), where various notions of a “product” of models are developed that allow to build multi-dimensional models from one-dimensional ones. Given an appropriate notion of product, the analysis of this paper can also be applied in multi-dimensional situations. Specifically, if the universe has a product structure such that each factor is acyclic, one can define the diversity function on the product as the product of the one-dimensional diversity functions in the sense of Hendon *et al.* (1996), and thus indirectly derive diversity from component-wise dissimilarity information.

On a methodological level, the paper has tried to demonstrate the great analytical usefulness of multi-attribute approach. This approach has made it possible, for example, to translate what is essentially an algebraic problem, the characterization of monotonicity in dissimilarity, into one concerning the “geometry” of families of sets, namely the characterization of acyclic models as tree-consistent.

In view of its inherent mathematical generality, the multi-attribute framework promises to be applicable in a number of contexts beyond the ones for which it was originally conceived. Indeed, diversity theory can be viewed as a general method of *non-additive similarity-based counting*. Two further applications seem particularly promising: economies of scope and case-based decision making. In another paper, Nehring and Puppe (1999b), we apply the framework to models of joint production and conceptualize economies of scope as resulting from the shared use of joint public inputs which can be thought of as “fixed cost attributes.” Secondly, the notion of similarity-based counting should also contribute to the understanding of learning from heterogeneous experience. To analyze such learning under premises of bounded rationality, Gilboa and Schmeidler (1995) have proposed a theory of “case-based decision making.” Their approach thematizes similarity of past experience to present decisions as central; this has the formal structure of a point (“case at hand”) – set (“past cases”) – metric. Since such metrics are at the center of diversity theory, it seems reasonable to hope that it should be able to contribute to the further development of case-based decision theory.

Appendix A: The Free Graph Model

Consider the family $\mathcal{B}(X)$ of all one- and two-element subsets of X . As a model, $\mathcal{B}(X)$ can be viewed as the *free graph* on X in which all elements are connected by an edge.³⁹ Note that $\mathcal{B}(X)$ does not contain the universal attribute X and is hence not regular. For any set function $v : 2^X \rightarrow \mathbf{R}$ with $\Lambda \subseteq \mathcal{B}(X)$, and all x, y , one has $\lambda_{\{x,y\}} = v(\{x\}) - d(x, y)$, and $\lambda_{\{x\}} = v(\{x\}) - \sum_{z \in X, z \neq x} \lambda_{\{x,z\}}$. This yields the following formula by straightforward computation. For all S ,

$$v(S) = \left(\frac{3}{2} - \frac{\#S}{2}\right) \sum_{x \in S} v(\{x\}) + \frac{1}{2} \sum_{x,z \in S} d(x, z). \quad (\text{A.1})$$

Clearly, $\mathcal{B}(X)$ is not acyclic whenever $\#X \geq 3$; on the other hand, it is immediate from (A.1) that $\mathcal{B}(X)$ is monotone in dissimilarity. Observe, however, that this is

³⁹With slight abuse of notation, a set $\{x, y\} \subseteq X$ may be identified with an edge connecting x and y in a graph. Any subset of $\mathcal{B}(X)$ can thus be viewed as a graph on X ; the graph associated with $\mathcal{B}(X)$ itself is “free” in that no edge is excluded.

accompanied by a degenerate behaviour of diversity with respect to changes in the value of singletons. Indeed, by (A.1) the diversity of a set S with $\#S \geq 4$ *decreases* in the value of singletons provided that all dissimilarities are held constant.

Appendix B: Proofs

Proof of Theorem 3.1 First, we prove the equivalence of (ii) and (iii). Using the representation (2.6) implied by vNM it is straightforward to show that SB implies $d(x, S) \geq \min_{y \in S} d(x, y)$. The converse inequality follows from submodularity of v (cf. (2.5)). The implication “(iii) \Rightarrow (ii)” is immediate.

Next, we show “(i) \Rightarrow (iii)”. For given x and S , consider the family $\mathcal{A}^x := \{A \in \Lambda : x \in A\}$. If $\mathcal{A}^x = \emptyset$, then $d(x, S) = d(x, y) = 0$ for all $y \in S$, hence (3.2) holds trivially. Thus, assume that $\mathcal{A}^x \neq \emptyset$. The Nesting Property implies that \mathcal{A}^x is a chain, i.e. totally ordered by set-inclusion; we may thus write $\mathcal{A}^x = \{A_1, \dots, A_l\}$ with $A_1 \subseteq A_2 \subseteq \dots \subseteq A_l$. Furthermore, set $A_0 = \emptyset$ and $A_{l+1} = X$. Let j be the maximal index such that $A_j \cap S = \emptyset$ and choose $z \in A_{j+1} \cap S$. By construction, $d(x, S) = d(x, z) = \lambda(\{A_1, \dots, A_j\})$. By submodularity, $d(x, S) \leq d(x, y)$ for all $y \in S$; hence $d(x, S) = d(x, z) = \min_{y \in S} d(x, y)$.

Finally, “(iii) \Rightarrow (i)” is shown by contraposition. Thus, suppose the Nesting Property is violated, i.e. there exist $B, C \in \Lambda$ such that $B \cap C$, $B \setminus C$, and $C \setminus B$ are all non-empty. Choose x, w, z according to $x \in B \cap C$, $w \in B \setminus C$, $z \in C \setminus B$, and consider the set $S = \{w, z\}$. Since $\{x, z\} \subseteq C$ and $w \notin C$ one has

$$d(x, \{w, z\}) = \lambda(\{A : x \in A, w \notin A, z \notin A\}) < \lambda(\{A : x \in A, w \notin A\}) = d(x, w).$$

By a completely symmetric argument, one also obtains $d(x, S) < d(x, z)$. Hence, $d(x, S) < \min_{y \in S} d(x, y)$.

Proof of Theorem 3.2 The implication “(i) \Rightarrow (iii)” has already been shown in the main text. In order to verify the converse implication, assume, by way of contraposition, that there exists $A^0 \in \Lambda$ that is not an interval, i.e. for some $y_1 < y_2 < y_3$, $\{y_1, y_3\} \subseteq A^0$ and $y_2 \notin A^0$. This implies $\lambda(\{A : y_3 \in A, A \cap \{y_1, y_2\} = \emptyset\}) < \lambda(\{A : y_3 \in A, y_2 \notin A\})$, and hence $v(\{y_1, y_2, y_3\}) - v(\{y_1, y_2\}) < d(y_3, y_2)$, in violation of (3.3).

Using the representation (2.6) implied by vNM, condition (3.4) is easily seen to be equivalent to the condition that $d(z, S) = d(z, S \cup \{x\})$ whenever $x \leq y \leq z$ and $S \ni y$. From this, the equivalence of (i) and (ii) is straightforward.

Proof of Proposition 3.1 The equivalence of (ii) and (iii) is immediate. Hence, it suffices to prove equivalence of (i) and (ii). By definition, the function d is additive in the sense of (3.6) if and only if, for all $x \leq y \leq z$,

$$\sum_{A: \{x, z\} \cap A \neq \emptyset} \lambda_A = \sum_{A: \{y, z\} \cap A \neq \emptyset} \lambda_A - \sum_{A: y \in A} \lambda_A + \sum_{A: \{x, y\} \cap A \neq \emptyset} \lambda_A. \quad (\text{B.1})$$

It is easily checked that for any interval A containing at least one endpoint, λ_A occurs as a summand on the l.h.s. of (B.1) if and only if it also occurs on the r.h.s. (possibly twice with a positive sign and once with a negative).

Conversely, let $A^0 \ni y$ be an interval that contains neither endpoint. Then, e.g. with $x = x_1$ and $z = x_n$, λ_{A^0} does not occur on the l.h.s. of (B.1) but twice with a positive

sign and once with a negative sign on the r.h.s. of (B.1). Hence, $\lambda_{A^c} > 0$ implies $d(x_n, x_1) < d(x_n, y) + d(y, x_1)$.

Proof of Theorem 4.1 Since any set function $v : 2^X \rightarrow \mathbf{R}$ with $\Lambda \subseteq \mathcal{A}$ can be represented as the difference of two diversity functions with the same support restriction (as a straightforward consequence of conjugate Moebius inversion), bideterminacy of \mathcal{A} can equivalently be formulated as follows. For all *set* functions v with $\Lambda \subseteq \mathcal{A}$ and all S ,

$$v^{\mathcal{B}(S)} = 0 \Rightarrow [v(W) = 0 \text{ for all } W \subseteq S]. \quad (\text{B.2})$$

Necessity of the no-free-triple condition has already been shown in the main text. Sufficiency is proved by induction on the cardinality of $X = \{x_1, \dots, x_n\}$. For $\#X = 3$, we distinguish the following three cases.

Case 1. $\mathcal{A} \not\supseteq X$, i.e. $\mathcal{A} \subseteq \mathcal{B}(X)$. Then bideterminacy follows at once from formula (A.1).

Case 2. \mathcal{A} does not contain some pair, say $\{x_1, x_3\}$. Then \mathcal{A} is compatible with the line model given the ordering $x_1 < x_2 < x_3$, whence the assertion follows from the Line Equation (3.3).

Case 3. $\mathcal{A} \not\supseteq \{x\}$, say $x = x_3$. Then, for any set function v with support $\Lambda \subseteq \mathcal{A}$, $v(X) - v(\{x_1, x_2\}) = \lambda_{\{x_3\}} = 0$, from which the claim follows immediately.

Suppose now that the assertion of the theorem holds for all X and \mathcal{A} such that $\#X \leq m - 1$; we will show it to hold for all X with $\#X = m$. Thus, take any set function v on X such that $\Lambda \subseteq \mathcal{A}$ and such that $v^{\mathcal{B}(X)} = 0$. Since \mathcal{A} has no free triples, evidently Λ and indeed all $\Lambda|_S$, $S \subseteq X$, have no free triples. Hence, by the induction hypothesis,

$$v(S) = 0 \text{ for all } S \subseteq X \text{ such that } S \neq X. \quad (\text{B.3})$$

By (B.2), it remains to show that $v(X) = 0$. Fix any subset $Y \subseteq X$ with $\#Y = \#X - 3$, and define $\tilde{v} : 2^{(X \setminus Y)} \rightarrow \mathbf{R}$ by

$$\tilde{v}(S) := v(S \cup Y) - v(Y).$$

By (B.3), $\tilde{v}^{\mathcal{B}(X)} = 0$. Since $\tilde{v}(S) = \lambda(\{A : A \cap S \neq \emptyset, A \subseteq X \setminus Y\})$, the support $\tilde{\Lambda}$ of the conjugate Moebius inverse of \tilde{v} has no free triples. By validity of the claim for all three-element subsets, $\tilde{v}(X \setminus Y) = 0$, i.e. $v(X) = v(Y) = 0$.

Proof of Proposition 5.1 Fix any $w \in X$; we have to show that $v(X) = v(\{w\}) + \sum_{(x,y) \in \eta_w} d(x,y)$, i.e.

$$\sum_{A \in \Lambda} \lambda_A = \sum_{A: w \in A} \lambda_A + \sum_{(x,y) \in \eta_w} \left(\sum_{A: x \in A, y \notin A} \lambda_A \right), \quad (\text{B.4})$$

if and only if Λ is compatible with τ . Clearly, any λ_A with $A \in \Lambda$ and $w \in A$ occurs exactly once as a summand on both sides of the equality (B.4). Now consider $A \in \Lambda$ with $w \notin A$; choose $x_0 \in A$ such that the path connecting w and x_0 has the minimal number of edges among all paths that connect w with some element in A . If A is τ -connected, such x_0 is uniquely determined; let y_0 be the unique element such that $(x_0, y_0) \in \eta_w$. Then, λ_A occurs exactly once on both sides of (B.4), on the r.h.s. as a summand in the term $d(x_0, y_0)$. This shows that (B.4) holds if all $A \in \Lambda$ are τ -connected.

On the other hand, if $A^0 \in \Lambda$ is not τ -connected, it is clear that λ_{A^0} occurs at least twice on the r.h.s. but only once on the l.h.s. of (B.4). Hence, $v(X) < l_d(\tau)$ since $\lambda_{A^0} > 0$. The argument in particular shows that the length of any tree on X is always an upper bound for $v(X)$ no matter whether Λ is compatible with that tree.

Proof of Proposition 5.2 We show that \mathcal{T}_{phy} has no free triples; this implies bideterminacy by Theorem 4.1. Thus, consider a triple $S = \{x, y, z\}$ and assume, by way of contradiction, that all three pairs $\{x, y\}$, $\{y, z\}$ and $\{x, z\}$ are elements of the relativization $\mathcal{T}_{\text{phy}}|_S$, i.e. $\{x, y\} = A \cap A_{tax} \cap S$, $\{y, z\} = A' \cap A'_{tax} \cap S$ and $\{x, z\} = A'' \cap A''_{tax} \cap S$ for some $A, A', A'' \in \mathcal{H}_{ev}$ and $A_{tax}, A'_{tax}, A''_{tax} \in \mathcal{H}_{tax}$. First, assume that $z \notin A$; since \mathcal{H}_{ev} is a hierarchy, this implies $x \in A'$ and $y \in A''$. Therefore, $x \notin A'_{tax}$ and $y \notin A''_{tax}$, but this contradicts the hierarchical structure of \mathcal{H}_{tax} . A completely symmetric argument shows that $z \notin A_{tax}$ leads to a contradiction as well.

Proof of Proposition 6.1 Applying the argument in the proof of Proposition 5.1 to the restriction $v|_S$, one obtains $v(S) \leq l_d(\tau)$ for any tree τ on S , and $v(S) = l_d(\tau_S)$ if and only if $\Lambda|_S$ is compatible with τ_S (recall that $\Lambda|_S$ is the support of the attribute weighting function corresponding to $v|_S$). This implies the stated equivalence at once.

Proof of Proposition 6.2 Suppose $\Lambda|_S$ is compatible with the tree τ_S on S . Then, $d(x, S \setminus \{x\}) = d(x, y)$ holds for all $x, y \in S$ such that y is the only element in $S \setminus \{x\}$ to which x is connected by an edge of τ_S . Thus, tree-consistency implies (6.2).

To verify the converse statement, construct a tree τ_S on any given S recursively as follows. Choose an “outer point” $x \in S$ and its “gate” $y \in S \setminus \{x\}$ according to (6.2) and connect x and y by an edge. Then apply the same procedure to the set $S \setminus \{x\}$, and so forth. Clearly, the resulting graph τ_S is connected and acyclic, hence a tree on S . By construction, $v(S)$ equals the length of τ_S (for instance, take as starting point w one of the two last elements that were connected by an edge). Hence, by Proposition 6.1, Λ is tree-consistent.

Proof of Proposition 6.3 The proof is by contraposition. Thus, suppose that \mathcal{A} is regular and m -cyclic for some $m \geq 3$; let x_1, \dots, x_m and $A_1, \dots, A_m \in \mathcal{A}$ be as required in the definition of m -cyclicity. Define a diversity function v_1 from its attribute weighting function λ_1 by setting $\lambda_1 = 1$ at the universal attribute X , all singletons and all sets A_1, \dots, A_m ; for all other attributes $A \in \mathcal{A}$, set $\lambda_1(A) = 0$. For $0 < \epsilon < 1$, define another diversity function v_2 by setting $\lambda_2(X) = 1 + \epsilon$, $\lambda_2(\{x\}) = 1 + \epsilon$ for all singletons, $\lambda_2(A_i) = 1 - \epsilon$ for all $i = 1, \dots, m$, and $\lambda_2 = 0$ otherwise. It is now easily verified that $v_1^{S(S)} = v_2^{S(S)}$, $d_1|_S \geq d_2|_S$, yet $v_1(S) < v_2(S)$ in violation of monotonicity in dissimilarity. Indeed, for all $x \in S$, $v_1(\{x\}) = v_2(\{x\}) = 4$; furthermore, $d_1(x_i, x_j) = d_2(x_i, x_j) = 2$ whenever $|i - j| = 1$ or $|i - j| = m - 1$, and $d_1(x_i, x_j) = d_2(x_i, x_j) + \epsilon = 3$ whenever $1 < |i - j| < m - 1$. Yet, $v_1(S) = 2m + 1 < 2m + 1 + \epsilon = v_2(S)$.

For the proof of Theorem 6.1, we need the following definitions. Let γ be a *circular graph* of the form $x_1 \gamma x_2 \gamma \dots \gamma x_m \gamma x_1$ on the set $S = \{x_1, \dots, x_m\}$. Denote by \mathcal{K}_γ the family of all γ -connected subsets of S . A family \mathcal{C} is called a *k-cycle* on S if there exists a circular graph γ such that (i) $\mathcal{C} \subseteq \mathcal{K}_\gamma$, (ii) for every edge $\{x_i, x_{i+1}\}$ of γ ,⁴⁰ there exists $A \in \mathcal{C}$ such that $A \supseteq \{x_i, x_{i+1}\}$, (iii) for no two $A, A' \in \mathcal{C}$, $A \subseteq A'$, and (iv) $\#\mathcal{C} = k$.

⁴⁰For notational convenience, here and in the following proofs, addition of indices is to be understood modulo m , so that $x_{m+1} = x_1$.

Lemma B.1 *A model \mathcal{A} on X is acyclic if and only if for no $S \subseteq X$ there exists $\mathcal{C} \subseteq \mathcal{A}$ such that $\mathcal{C}|_S$ is a k -cycle on S for some $k \geq 3$.*

Proof Let \mathcal{A} be m -cyclic for some $m \geq 3$, and choose $S = \{x_1, \dots, x_m\}$ and $\mathcal{C} = \{A_1, \dots, A_m\} \subseteq \mathcal{A}$ satisfying (6.3). Clearly, in this case $\mathcal{C}|_S$ constitutes an m -cycle on S in the sense of the above definition.

Conversely, suppose that \mathcal{C} and $S = \{x_1, \dots, x_m\}$ are such that $\mathcal{C}|_S$ is a k -cycle with respect to the circular graph $x_1 \gamma x_2 \gamma \dots \gamma x_m \gamma x_1$ for some $k \geq 3$. Suppose first that there exist A_j and $A_{j'}$ in \mathcal{C} such that $A_j \cup A_{j'} \supseteq S$. Then, there must exist $A_l \in \mathcal{C} \setminus \{A_j, A_{j'}\}$ that intersects both A_j and $A_{j'}$. It is easily verified that in this case one can choose a three-element set $W = \{y_1, y_2, y_3\} \subseteq S$ such that $A_j \cap W = \{y_1, y_2\}$, $A_l \cap W = \{y_2, y_3\}$ and $A_{j'} \cap W = \{y_1, y_3\}$. Hence, \mathcal{A} is 3-cyclic, in particular not acyclic.

Now suppose that S is not included in the union of two elements of \mathcal{C} . Then, a set $W = \{y_1, \dots, y_l\}$ and a family $\{A_1, \dots, A_l\}$ satisfying (6.3) can be recursively determined as follows. Let $A_1 \in \mathcal{C}$; choose as y_1 the “endpoint” of A_1 in clockwise orientation, i.e. $y_1 = x_j$ where $x_j \in A_1$ and $x_{j+1} \notin A_1$. Now consider the set of all $A \in \mathcal{C}$ with $y_1 \in A$; among these, choose as A_2 the set whose “endpoint” y_2 (in clockwise orientation) is farthest away from y_1 . Continue recursively until for the first time $y_l \in A_1$. By construction, \mathcal{A} is l -cyclic with $l \geq 3$.

Proof of Theorem 6.1 First, we show that tree-consistency implies acyclicity. Suppose by way of contradiction, that \mathcal{A} is m -cyclic for some $m \geq 3$, and consider the set $S = \{x_1, \dots, x_m\}$ in (6.3). By tree-consistency, there exists a tree τ_S on S such that for all $A \in \mathcal{A}$, $A \cap S$ is τ_S -connected. But by (6.3), each $\{x_i, x_{i+1}\}$ must correspond to an edge of τ_S which contradicts the acyclicity of τ_S .

The basic idea for the proof of the converse implication is to ask, for given $S \subseteq X$, which connected graphs γ have the property that all attributes in the intersection closure⁴¹ $(\mathcal{A}|_S)^*$ of $\mathcal{A}|_S$ are γ -connected, and to show that acyclicity of \mathcal{A} forces the minimal such graphs to be acyclic, i.e. to be trees. To make this work, one needs to derive from the existence of a cycle in a minimal graph γ the existence of a k -cycle \mathcal{C} on an appropriate subset $S' \subseteq S$; the key difficulty is to ensure the γ -connectedness of the elements of $(\mathcal{A}|_S)^*$ in the subset S' . To overcome the difficulty, an induction argument is used.

Thus, we have to show that acyclicity implies tree-compatibility on every subset $S \subseteq X$. This is done by induction on the cardinality of S . The claim is evident for $\#S \leq 3$. Assuming validity of the claim for S , we need to show its validity for $W := S \cup \{x\}$ for any $x \in X \setminus S$. Thus, by assumption there exists a tree τ_S on S such that, for all A ,

$$A \in (\mathcal{A}|_S)^* \Rightarrow A \text{ is } \tau_S\text{-connected in } S. \quad (\text{B.5})$$

Note that τ_S -connectedness of any $A \in \mathcal{A}|_S$ implies τ_S -connectedness of any A in the intersection closure $(\mathcal{A}|_S)^*$ by acyclicity of τ_S . Let Ω denote the family of all connected graphs $\gamma \subseteq \tau_S \cup \{\{x, y\} : y \in S\}$ on W such that, for all A ,

$$A \in (\mathcal{A}|_W)^* \Rightarrow A \text{ is } \gamma\text{-connected in } W.$$

Note that Ω is trivially non-empty, as it contains the graph $\tau_S \cup \{\{x, y\} : y \in W\}$. Let γ be any \subseteq -minimal graph in Ω . We will show by contradiction that γ is acyclic, i.e. a

⁴¹By definition, the *intersection closure* \mathcal{A}^* of a model \mathcal{A} is the smallest family that contains \mathcal{A} and is closed under intersections.

tree. Suppose not. Clearly, any cycle in γ must be of the form $x\gamma y_1\gamma\dots\gamma y_m\gamma x$, where $S' := \{y_1, \dots, y_m\}$ is a τ_S -path in S ; let $y_0 := x$ and $W' := S' \cup \{y_0\}$.

Claim 1. For any edge $\{y, y'\}$ in γ , there exists $A \in \mathcal{A}|_W$ with $\{y, y'\} \subseteq A$ such that $\gamma \cap (A \times A)$ contains no cycle through $\{y, y'\}$ (say that A “ratifies” the edge $\{y, y'\}$). Indeed, suppose that this was false; then, consider $\gamma' = \gamma \setminus \{\{y, y'\}\}$. By assumption, every $A \in \mathcal{A}|_W$ with $\{y, y'\} \subseteq A$ contains a γ -cycle through $\{y, y'\}$; thus, every $A \in (\mathcal{A}|_W)^*$ remains connected with respect to γ' (by travelling the cycle in the opposite direction, if necessary). However, this contradicts the assumed minimality of γ .

Claim 2. For any $A \in (\mathcal{A}|_W)^*$, $A \cap S'$ is γ -connected in S .

This follows from (B.5) and τ_S -connectedness of S' , since the intersection of any path in a tree with a connected subset of it yields a τ_S -connected subset of the path, which by the form of S' is also γ -connected.

By Claim 1, all cycles W' possess a “covering,” i.e. a family $\mathcal{C} \subseteq \mathcal{A}|_W$ such that, for all edges $\{y_i, y_{i+1}\}$, there exists $A \in \mathcal{C}$ ratifying $\{y_i, y_{i+1}\}$ in the sense that⁴²

$$\{y_i, y_{i+1}\} \subseteq A \cap W' \text{ and } A \neq W'.$$

Pick a cycle W' and a covering \mathcal{C} of minimal cardinality (over all cycles and all coverings). Observe that by minimality, the elements of \mathcal{C} are non-nested. We note the following properties of \mathcal{C} .

Claim 3. There exist $A^1, A^m \in \mathcal{C}$ such that $A^1 \cap W' = \{x, y_1, \dots, y_k\}$ and $A^m \cap W' = \{y_l, \dots, y_m, x\}$.

Clearly, for any $A \in \mathcal{C}$, $A \not\supseteq \{x, y_1, y_m\}$, since otherwise by Claim 2, $A \supseteq W'$, whence $A \notin \mathcal{C}$. The sets A^1 and A^m ratifying the edges $\{x, y_1\}$ and $\{x, y_m\}$, respectively, must thus differ, and have the asserted form by Claim 2 again.

Claim 4. $A^1 \cap A^m = \{x\}$, i.e. $k < l$.

Suppose not, i.e. that $l \leq k$. By the construction of γ , $A^1 \cap A^m$ is γ -connected in W , i.e. there exists a path from y_l to x in W that crosses neither y_1 nor y_m . It follows that A^1 contains a γ -cycle in W through the edge $\{x, y_1\}$ in contradiction to the fact that A^1 ratifies $\{x, y_1\}$.

By the same token, the sets A^k and A^l (not necessarily distinct) in \mathcal{C} ratifying the edges $\{y_k, y_{k+1}\}$ and $\{y_{l-1}, y_l\}$, respectively, do not contain x . In particular, $\#\mathcal{C}|_{W'} \geq 3$. Let κ denote the circular graph on W' given by $x\kappa y_1\kappa\dots\kappa y_m\kappa x$. We will show that \mathcal{C} induces a cycle on W' with respect to κ . For this, it remains to show that every $A' \in \mathcal{C}|_{W'}$ is κ -connected in W' . By Claim 2, this is clear for all $A' \not\ni x$ since on S' the graphs κ and γ coincide. Hence, if $A' \in \mathcal{C}|_{W'}$ is to be disconnected in W' , it must be of the form $A' = B \cup \{x\}$ for some non-empty $B \subseteq S'$. By Claim 2, B is γ -connected in S' , hence by the argument in the proof of Claim 4, $B \subseteq \{y_{k+1}, \dots, y_{l-1}\}$. But this contradicts minimality of \mathcal{C} , as follows. Let j be the minimal index such that $y_j \in B$ and let $A \in \mathcal{C}|_W$ be such that $A' = A \cap W'$. By the γ -connectedness of A in W , there exists a γ -path $\{y_j, z_1, \dots, z_q, x\}$ in A . Then consider the γ -cycle $x\gamma y_1\gamma\dots\gamma y_j\gamma z_1\gamma\dots\gamma z_q\gamma x$ and construct a covering \mathcal{C}' on this cycle of lesser cardinality replacing A^m by A and removing at least A^l . We have thus established that $\mathcal{C}|_{W'}$ is a k -cycle on W' with $k \geq 3$. By Lemma B.1, \mathcal{A} can thus not be acyclic, the desired contradiction.

Proof of Proposition 6.4 The Metric Link Property (6.4) requires, for any set S , the

⁴²Recall that addition of indices is to be understood modulo m .

existence of an “outer point” $x \in S$ and its “gate” y to $S \setminus \{x\}$ such that $d(x, y) \leq d(x, z)$ and $d(z, x) \geq d(z, y)$ for all $z \in S \setminus \{x\}$. As already noted in the main text, part (i) follows at once from Theorem 6.1. Hence, consider part (ii), and assume that the support Λ of v has a minimal cycle of length $m \geq 4$. Specifically, suppose that for some $m \geq 4$, $S = \{x_1, \dots, x_m\}$ and $\{A_1, \dots, A_m\} \subseteq \Lambda$ satisfy (6.3) while for no proper subset of S there exist relevant attributes such that (6.3) is satisfied. We verify the violation of (6.4) by contradiction. Thus, suppose that $x_j \in S$ is an outer point in S . Consider any $l \notin \{j-1, j, j+1\}$; by minimality of S , any attribute $A \in \Lambda$ that contains x_j and x_l must also contain x_{j-1} and x_{j+1} . This implies that, for any such index l , $d(x_j, x_l) > d(x_j, x_{j-1})$ and $d(x_j, x_l) > d(x_j, x_{j+1})$. Hence, the gate of x_j to $S \setminus \{x_j\}$ must be either x_{j-1} or x_{j+1} . However, one also has $d(x_{j-1}, x_j) < d(x_{j-1}, x_{j+1})$ and $d(x_{j+1}, x_j) < d(x_{j+1}, x_{j-1})$, again exploiting minimality of S . Hence neither x_{j+1} nor x_{j-1} can serve as the gate. Since x_j was arbitrary, the Metric Link Property must thus be violated.

Part (iii) follows at once from part (ii) by contraposition: given an m -cyclic model with $m \geq 4$, simply choose a compatible diversity function v such that its support Λ has a minimal cycle of length m .

References

- [1] CHATEAUNEUF, A. and J.-Y. JAFFRAY (1989), Some Characterizations of Lower Probabilities and Other Monotonic Capacities through the Use of Moebius Inversion, *Mathematical Social Sciences* **17**, 263-283.
- [2] DEMANGE, G. (1982), Single Peaked Orders on a Tree, *Mathematical Social Sciences* **3**, 389-396.
- [3] DEMPSTER, A. (1967), Upper and Lower Probabilities Induced by a Multi-Valued Mapping, *Annals of Mathematical Statistics* **38**, 325-339.
- [4] DOBSON, A. (1996), *Conservation and Biodiversity*, Scientific American Library, New York.
- [5] ESTEBAN, J.-M. and D. RAY (1994), On the Measurement of Polarization, *Econometrica* **62**, 819-851.
- [6] FAITH, D.P. (1994), Phylogenetic Pattern and the Quantification of Organismal Biodiversity, *Phil. Trans. Roy. Soc. Lond. B* **345**, 45-58.
- [7] FAITH, D.P. and P.A. WALKER (1994), *Diversity: Reference and User's Guide*, Canberra: CSIRO Division of Wildlife and Ecology.
- [8] GASTON, K.J. (ed.) (1996), *Biodiversity. A Biology of Numbers and Difference*, Blackwell, Oxford.
- [9] GILBOA, I. and D. SCHMEIDLER (1995), Case-Based Decision Theory, *Quarterly Journal of Economics* **110**, 605-639.

- [10] HENDON, E., H.J. JACOBSEN, B. SLOTH and T. TRANAES (1996), The Product of Capacities and Belief Functions, *Mathematical Social Sciences* **32**, 95-108.
- [11] HURKA, T. (1983), Value and Population Size, *Ethics* **93**, 496-507.
- [12] JOHNSON, S. (1967), Hierarchical Clustering Schemes, *Psychometrika* **32**, 241-254.
- [13] KEIRSEY, D. and M. BATES (1984), *Please Understand Me – Character and Temperament Types*, Prometheus Nemesis, Del Mar (CA).
- [14] KREPS, D.M. (1979), A Representation Theorem for “Preference for Flexibility,” *Econometrica* **47**, 565-577.
- [15] LANCASTER, K. (1966), A New Approach to Consumer Theory, *Journal of Political Economy* **74**, 132-157.
- [16] LANCASTER, K. (1979), *Variety, Equity, and Efficiency*, Basil Blackwell, Oxford.
- [17] LAZEAR, E.P. (1999), Globalization and the Market for Team-Mates, *Economic Journal* **109**, C15-C40.
- [18] NEHRING, K. (1999a), Preference for Flexibility in a Savage Framework, *Econometrica* **67**, 101-120.
- [19] NEHRING, K. (1999b), Diversity and the Geometry of Similarity, *mimeo*, UC Davis.
- [20] NEHRING, K. and C. PUPPE (1999a), A Theory of Diversity, SFB-303 Discussion Paper A-605, University of Bonn.
- [21] NEHRING, K. and C. PUPPE (1999b), Modelling Economies of Scope, SFB-303 Discussion Paper A-607, University of Bonn.
- [22] NEHRING, K. and C. PUPPE (2001), Modelling Phylogenetic Diversity, *mimeo*, University of Bonn.
- [23] PATTANAİK, P. and Y. XU (1990), On Ranking Opportunity Sets in Terms of Freedom of Choice, *Recherches Economiques de Louvain* **56**, 383-390.
- [24] ROTA, G.C. (1964), Theory of Moebius Functions, *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **2**, 340-368.
- [25] SHAFER, G. (1976), *A Mathematical Theory of Evidence*, Princeton University Press, Princeton.
- [26] SHAPLEY, L.S. (1953), A Value for n -Person Games, in: *Contributions to the Theory of Games II (Annals of Mathematics Studies 28)*, ed. by H.W. Kuhn and A.W. Tucker, Princeton University Press, Princeton.
- [27] SUGDEN, R. (1998), The Metric of Opportunity, *Economics and Philosophy* **14**, 307-337.
- [28] WEITZMAN, M. (1992), On Diversity, *Quarterly Journal of Economics* **107**, 363-405.
- [29] WEITZMAN, M. (1998), The Noah’s Ark Problem, *Econometrica* **66**, 1279-1298.