Oecologia (2007) 154:75–84
DOI 10.1007/s00442-007-0810-3

POPULATION ECOLOGY

Quantifying change in distributions: a new departure index that detects, measures and describes change in distributions from population structures, size-classes and other ordered data

K. M. Menning · J. J. Battles · T. L. Benning

Received: 3 March 2006 / Accepted: 26 June 2007 / Published online: 11 August 2007

Abstract Many statistics are available to compare distributions. Some are limited to nominal data while others, such as skew, Kullback–Leibler, Kolmogorov–Smirnov and the Gini coefficient, are useful for providing information about ordered distributions. While many of these tests are useful for determining properties of data in histograms, there has not been a test until now that allows for the detection of differences between distributions, describes the difference and is sensitive to the location of the departures. Such a test could be critical for comparing pre-and post-event distributions, such as a change in the distribution of biomass due to fire, for example, or for comparing data from different locations, such as soil size distributions, and even for evaluating economic disparity or examining differences in age demographics. We present a new statistic, a departure index, which allows a test distribution to be compared with any reference distribution. The resulting index contains information about the location, magnitude and direction of departure from the reference distribution to the test distribution. The departure index in turn provides a standardized response range that allows for a comparison of results from different analyses. A case study of actual fire data demonstrates the sensitivity and range of the test.

Keywords Change detection · Quantitative analysis · Size class distributions · Structure · Statistics

Introduction

In science in general and ecology in particular, frequency distributions are commonly used to summarize the underlying structure of populations or resources (Begon et al. 1996). Data – whether they be size, age, performance ratings or other measures – are organized into defined categories or bins, and the numbers of organisms (or objects) per bin are examined in a histogram. With nominal data, the order of categories does not matter: for example, the order of species does not matter when calculating diversity indices such as the Shannon-Wiener Diversity Index (Zar 1999). In contrast, order does matter with ordinal, interval or ratio data (i.e. ordered data). With tree size classes, for example, order is inherent in the categories being considered. Often, the challenge in an ecological analysis is to quantify the changes in ordered distributions that result from a natural or experimental process. The ideal comparison would not only detect a difference between distributions but also describe the nature of the difference.

A number of statistical measures test differences between distributions (Wiegand et al. 2000; Yang et al. 2004). Some analyses are limited to nominal data, such as the Shannon-Wiener. Some, such as the chi-square goodness-of-fit test (Payette et al. 2000; Gaymer et al. 2001), the

Communicated by Marc Mangel.

Electronic supplementary material The online version of this article (doi:10.1007/s00442-007-0810-3) contains supplementary material, which is available to authorized users.

K. M. Menning (✉) · J. J. Battles
Division of Ecosystem Sciences,
Department of Environmental Science,
Policy and Management, University of California,
Berkeley, 137 Mulford Hall #3114,
Berkeley, CA 94720, USA
e-mail: kmenning@nature.berkeley.edu

T. L. Benning
Department of Environmental Science,
Harney Science Center, University of San Francisco,
2130 Fulton St, San Francisco, CA 94117, USA

paired-sample *t*-test (Jimenez and Decaens 2000), are useful for detecting differences between sets of nominal data but cannot be applied to ordered data.

The non-parametric Mann–Whitney test is appropriate for comparing non-parametric ordered distributions. Because Mann–Whitney is non-parametric, however, and only considers ranks instead of actual values, it discounts the size of real and important outliers that in biological systems may dominate the distribution, such as a single large tree. Hence, it is inadequate for testing differences between distributions in which outliers are important.

The Kolmogorov–Smirnov goodness-of-fit test robustly detects a difference between ordered distributions (Savage 1994; Payette et al. 2000; Yang et al. 2004). This test identifies the greatest difference in any bin between two cumulative histograms. The greater the difference, the more likely the distributions are statistically different, but it does not describe wherein the difference lies. A difference between the middle of two distributions is measured the same as a difference found at one of the tails of the distributions.

The Kullback–Leibler Criterion (Burnham et al. 2000; Burnham and Anderson 2002) is similar to Kolmogorov–Smirnov in its ability to detect differences. It is a useful measure of the absolute difference, or distance, between two distributions. The criterion, however, does not describe where the difference between the two distributions occurs. Lateral shifts of the mode left to right are not differentiated, for example. All differences are considered to be positive and measured in absolute terms regardless of where they occur. Also, the criterion does not produce symmetric results; comparing distribution *a* to distribution *b* yields a different result than comparing *b* to *a*.

A few other tests, such as skewness and the Gini coefficient, offer a greater range of results; they can be used to describe some of the differences between distributions. Skewness is a measure of the directional lopsidedness of a distribution (Weiner and Solbrig 1984; Bergqvist 1999). A positive value indicates that the distribution is skewed to the right, toward larger numbers, whereas a negative value indicates the distribution is skewed to the left, toward smaller numbers. Skewness, however, is insensitive to the distance of outliers from the mean (Weiner and Solbrig 1984). A group of small trees under a medium-sized tree is very different than the same group of small trees under one very large tree, yet the distributions' skews will be identical. If knowing the location and magnitude of a difference between distributions is important, then skewness is not adequate.

The Gini coefficient can be used to assess hierarchies, or inequalities in evenness, in distributions (Weiner and Solbrig 1984; Weiner 1990). If all trees were clumped in the same size class, for example, there would be no hierarchy because all elements (trees) would be equal. In the case of a single large tree with many small trees, hierarchy is high because one element dominates all others (Weiner and Solbrig 1984; Weiner 1990). While the Gini coefficient is effective at finding differences in hierarchy, it is ineffective at differentiating between distributions with similar hierarchies. Consider three distributions in which all individuals of each population are in a single size class. The Gini statistic for each of these homogeneous populations is zero. We know they have equal hierarchy – none – but we do not know from the statistic in which size classes the trees are clustered. In short, the Gini coefficient indicates whether the distribution is clumped, but not where that distribution is clumped. Also, the Gini coefficient has very limited ability to describe the direction of a departure between distributions (Menning 2003).

None of these metrics does a complete job of both detecting and characterizing the difference between one ordered distribution and another. In order to identify meaningful differences between two distributions, a metric should measure the magnitude and direction of the differences as well as maintain sensitivity to the location of departure from one distribution to another. For example, if a fire burns in a forest and kills trees in all size classes, we know that biomass has been reduced. But has it been reduced equally in all size classes? Has there been a shift in forest structure toward larger or smaller trees? It would be useful to have a metric that can answer these questions.

Trends in ordered histograms are often best observed in cumulative histograms. In Fig. 1a, two distributions with the same number of elements have only slight differences in the values in each bin and, consequently, determining a trend requires a keen eye. The cumulative data in Fig. 1b, however, illustrate that there is a clear overall trend: the gray distribution lags behind the black distribution but eventually catches up to it (since they have the same number of elements). Compared to the black distribution, the gray distribution is "right-shifted", or weighted toward larger size classes. Thus, it is not the difference in any one bin that makes a difference, but the accumulated trend of differences. For the sake of clarity, we use the term "direction" to indicate whether one distribution is shifted toward the larger (right) or smaller (left) end of a distribution.

"Location" describes the location on a distribution's horizontal axis where the major difference occurs. Consider two normal soil texture distributions that are very similar, with the exception that the second distribution has several additional large stones in it. The location of the difference between the two distributions is then toward the right end of the axis, where the large stones make one distribution different from the other. If the only difference between the two distributions were a few pieces of gravel

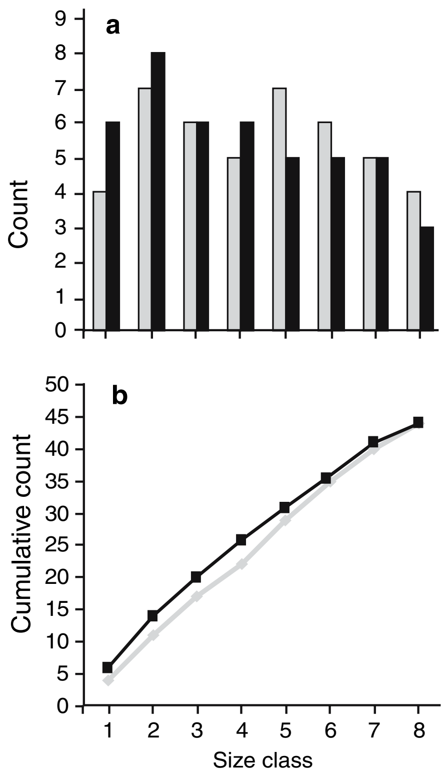

Fig. 1 **a** Two distributions (*black* vs. *gray bars*) contain the same total number of elements. As compared to the *black* distribution, the *gray* set of column bars has a few elements that occur later in the distribution, in larger size classes. **b** Cumulative distributions of the data in Fig. 1. The gray distribution is right-shifted, toward larger elements

in a mid-size class, then the location of the difference would be in the middle.

"Magnitude" is used to describe the extent of the difference between two distributions. In the case of the soil texture example, the magnitude of the difference is much greater if one set has 100 extra large stones than if there are only three extra. To describe the relative vertical height of a distribution, we use the standard term "amplitude". Consider two distributions of trees. The second set has tenfold more trees in each size class; the proportions are the same, but the totals different. The difference between these two distributions is that the amplitude of the second distribution is tenfold the first.

Any metric used to describe the difference between distributions should have the following properties. First, it should be well behaved; we should be able to look at the result and know what the answer means compared to related analyses. Second, the result should be standardized. If the metric has a standard response range, regardless of the distributions being examined, we can readily compare differences among analyses. Third, the metric should be relatively insensitive to the number of bins in the histogram. If we change the number of histogram bins in which the data are sorted, the metric should not change markedly.

In this paper we present a new statistic for consideration. The departure index meets each of the criteria described and offers promise as a tool for examining a wide variety of ordered data. Distributions may be compared with standardized and comparable results. A case study using actual fire data, as well as modeled results, is presented to demonstrate the sensitivity and range of the metric.

Methods

Derivation of the departure index

The departure index measures both the direction and magnitude of a departure from one distribution to another. The calculation compares two cumulative distributions: a reference distribution and a test distribution. The reference distribution may be a null distribution, a theoretical distribution or a measured distribution. The distribution that is tested against the reference distribution may be empirical or theoretical. For a full derivation, refer to Menning (2003).

The generalized version of the departure index is assigned the letter M and is written as:

$$M = \left(\frac{2}{k-1}\right) \frac{\left[\left(\hat{F}_1 - F_1\right) + \left(\hat{F}_2 - F_2\right) + \left(\hat{F}_3 - F_3\right) + \cdots + \left(\hat{F}_k - F_k\right)\right]}{n} \quad (1)$$

In this equation, $\hat{F}_1$ is the cumulative reference distribution for the first size class, and F_1 is the cumulative number of trees from the test distribution in that size class. The numerator contains the sum of differences between the cumulative histogram bins divided by the total number of elements n. The $k-1$ denominator is included to correct for the degrees of freedom as the number of bins in a histogram affects the calculation. One degree of freedom is lost because the index uses proportional data that sum to 1. If there are k bins, and therefore k comparisons, there are $k-1$ degrees of freedom, and the result must be divided by this quantity. Second, a scaling factor, 2, is included to adjust the result so that the departure index always has a range of two from minimum to maximum (−1 to +1, for example).

In order to simplify, the equation must first be expanded and values for the cumulative distribution variables, $\hat{F}_i$ and F_i must be replaced by their non-cumulative mathematical equivalents. The number of trees in any cumulative reference bin i, $\hat{F}_i$, is:

$$\hat{F}_i = n(\hat{p}_1 + \hat{p}_2 + \cdots + \hat{p}_i) \quad (2)$$

In this equation, $\hat{p}_i$ is the proportion of trees in the ith size class of the reference distribution. Also, each value of F represents a cumulative total including smaller size classes:

$$F_k = f_1 + f_2 + f_3 + \cdots + f_k \quad (3)$$

Substituting Eqs. 2 and 3 into Eq. 1 produces the generalized, expanded form of the departure index:

$$M = \left(\frac{2}{k-1}\right) \frac{[(n\hat{p}_1 - f_1) + (n(\hat{p}_1 + \hat{p}_2) - (f_1 + f_2)) + \ldots + (n(\hat{p}_1 + \cdots + \hat{p}_k) - (f_1 + \cdots + f_k))]}{n} \quad (4)$$

The equation may then be reorganized and simplified:

$$M = \left(\frac{2}{k-1}\right)\left(k\left(\hat{p}_1 - \frac{f_1}{n}\right) + (k-1)\left(\hat{p}_2 - \frac{f_2}{n}\right) + (k-2)\left(\hat{p}_3 - \frac{f_3}{n}\right) + \cdots + \left(\hat{p}_k - \frac{f_k}{n}\right)\right) \quad (5)$$

Now that the equation has reached its full expression, it may be rewritten for ease of use.

The resulting formulation of the statistic is:

$$M = \left(\frac{2}{k-1}\right)\sum_{i=1}^{k}\left[\left(\hat{p}_i - \frac{f_i}{n_f}\right)(k+1-i)\right] \quad (6)$$

In this equation, k is the number of bins in the histogram (e.g. small, medium, large: k = 3 bins); f_i is the count of elements in bin i of the test distribution (e.g. how many trees are in size class i?); n_f is the total number of elements in all bins of the test distribution (e.g. how many trees are in the plot?); $\hat{p}_i$ denotes the proportion of elements in bin i of the reference distribution.

If one wishes to calculate M using actual numbers from the reference distribution, rather than proportions, the following substitution into Eq. 6 is suggested:

$$\hat{p}_i = \frac{\hat{f}_i}{n_{\hat{f}}} \quad (7)$$

In this equation, $\hat{f}_i$ represents the count of elements in bin i of the reference histogram and $n_{\hat{f}}$ equals the number of elements in the reference distribution:

$$M = \left(\frac{2}{k-1}\right)\sum_{i=1}^{k}\left[\left(\frac{\hat{f}_i}{n_{\hat{f}}} - \frac{f_i}{n_f}\right)(k+1-i)\right] \quad (8)$$

Properties of the departure index

The departure index has a number of distinct properties that require elaboration in order to understand how it can be applied in the case study. First, the statistic is well-behaved. A positive value will always indicate the test distribution is right shifted compared to the reference distribution. A negative value will always indicate the test distribution is left-shifted. The magnitude of the departure index will indicate how far the test distribution is shifted compared to the reference distribution. There are no unusual behaviors as the value of the index approaches its maximum or minimum.

Second, because the departure index compares relative frequencies, the index's response is standardized. The range of the departure index always has an absolute value of 2. As a result, a departure index value of +0.2 from one analysis is the same as a +0.2 value from a second analysis as long as they use bins defined the same way.

Third, the maximum and minimum values of the departure index depend on the type of reference distribution used and so contain important information about that distribution. If a symmetric reference distribution is used (e.g. uniform or normal), possible departure index values will range from −1 to +1. This is because compared to the evenly-distributed reference distribution, the test distribution could depart equally left (−1) or right (+1).

If the reference distribution itself is already weighted heavily in one direction, however, the degree to which the test distribution could depart is asymmetric. This principle may be best illustrated with a common inverse-*j* shaped (exponentially declining) distribution as reference (Fig. 2). This distribution might describe tree diameter distributions, particle size abundance, foraging distance or age classes – any distribution in which there are many more small elements than large. Compare this reference distribution to two test distributions that are entirely right- or left-shifted; all the elements are in the tails of the distributions (Fig. 2). Because the reference distribution is already strongly left shifted, a complete shift to the left in the test distribution

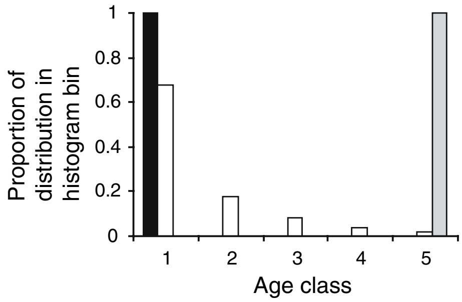

Fig. 2 Departures from a reference distribution may not be symmetric, as shown in this figure of an inverse-*j* shaped reference distribution (*white*) and two test distributions. Each test distribution, left-shifted (*black*) and right-shifted (*gray*), departs from the reference distribution by having all elements in a single bin at one end of the histogram. Because the reference distribution is already left-shifted, the magnitude of the left-shifted distribution's departure is less than that of the right-shifted distribution

results in a departure index value with a modest magnitude (−0.27). In contrast, the test distribution comprising large elements departs markedly, and this is reflected in the magnitude of its departure index value: +1.73. Hence, the departure index should be reported by stating its value and range endpoints: "1.73 (−0.27, 1.73)".

The unusual range endpoints of the departure index (e.g. −0.27, 1.73) is functional and should not be normalized. As currently described, the departure index always has an absolute range of 2. If the result is divided by the maximum value (1.73), or otherwise adjusted to attempt normalization, the range is no longer consistent. One analysis' range between endpoints might be 1.4 and another 1.8, instead of the standard 2.0. Departure index values from such analyses cannot be compared with each other. Further, after normalization, a +0.3 change in gravel texture size due to a flood, for example, would not be equivalent to a +0.3 change in a different soil setting. The ability to draw comparisons is lost with such normalization. In addition, information on the original range (which indicates something about the lopsidedness of the reference distribution) is lost.

Case study: methods

Field data

The study area is located in Mineral King Valley, Sequoia and Kings Canyon National Parks, California (USA). Data were collected in five 1-ha plots located near the Tar Gap Trail at an elevation of approximately 2700 m (D. Newburn and J. Battles, unpublished data). These plots are located within 2 km of each other in a montane conifer forest dominated by red fir (*Abies magnifica*). Pitcher (1987) conducted a detailed fire history of the red fir forest along the Tar Gap Trail. See his paper for a more comprehensive description of the study area.

In 1997, the five study plots were designated for prescribed burning. Before the fire, each plot was gridded into 5×5-m quadrats, and all patches of young regenerating red fir were identified. Regeneration patches were defined as areas with (1) no live canopy-sized trees and (2) an average density of more than 2000 understory trees per hectare (more than five trees per quadrat). Trees less than 20 cm in diameter at breast height (DBH; measured at 1.37 m height) were considered understory trees. Seedlings were defined as any tree between 30 and 130 cm in height. The DBH was measured to within 1 cm for any understory tree taller than 1.37 m. The exterior corners of the patches were marked with stainless steel rods that would endure fire. In the fall of 1998, a prescribed fire burned in all plots. All trees in the quadrats were re-inventoried in 2000 to determine mortality due to fire.

Statistical evaluation

Histograms for this analysis were created from tree size measurements. Five size classes were created (Table 1). All seedlings were placed into one seedling class, and the understory trees were grouped into 5-cm DBH bins.

We used the departure index to compare the observed pre-fire data with four different post-fire tree distribution models: one real and three theoretical. The three distributions were selected to demonstrate how the departure index can be used to evaluate the magnitude, direction and statistical confidence of changes in the size-distribution given widely different scenarios.

Table 1 Size-specific mortality rates for observed and simulated conditions for a prescribed fire in the red fir forests in Mineral King Canyon, Sequoia-Kings Canyon National Park (1997)

Size class	Initial density pre-fire (stems/ha)	Observed post-fire mortality (%)	Modeled mortality of trees		
			Null (%)	Small-size-biased (%)	Large-size-biased (%)
Seedlings	2,587	65	65	0	80
0–5 cm	1,641	70	65	20	60
5–10 cm	708	61	65	40	40
10–15 cm	249	53	65	60	20
15–20 cm	72	24	65	80	0

In each of the three modeled cases, the overall mortality rate (65%) was kept constant to match the actual overall mortality rate. The simulated results from the three modeled mortality scenarios – null mortality (n = 999 simulations), small- and large-size-biased mortality – were compared to observed results using the departure index

Our null model assumed that the probability of fire-caused mortality is independent with respect to tree size. For this model, 65% of the trees were randomly removed independent of size-class (this is the same percentage of trees in all size classes killed by the actual fire). Theoretically, a collection of such size-independent post-fire distributions, regardless of how many trees survived the fire, should have a mean departure index value of 0 when compared with the reference pre-fire distribution.

To illustrate the flexibility of the departure index, we also generated two theoretical post-fire distributions in which we assumed fire preferentially killed trees in different size classes. For the large-size-biased model, we decreased the death rate for each increase in tree size class at a steady rate. In this model, 80% of the seedlings were killed, 60% of the trees in the 5- to 10-cm range and so on until none of the large trees were killed. The resulting distribution had a disproportionate number of large trees after fire (Table 1). We inverted this bias to create a post-fire distribution where larger trees were preferentially killed and all seedlings survived (small-size-biased model; Table 1). Such a distribution should be even more heavily skewed toward the small size classes when compared to a typical inverse-*j* shaped population curve.

Absolute measures of "distance" or the difference between distributions are provided by applying the Kullback–Leibler criterion (Burnham et al. 2000; Burnham and Anderson 2002). Significance was determined using Monte Carlo simulations to quantify the range of departure indices that would be expected under the assumption of size-independent mortality (Manly 2007). To do so, we ran 999 simulations where 65% of the trees were randomly removed (the null model). The departure index was calculated for the distribution of the surviving trees after each iteration. From the 999 departure index values we calculated a 95% confidence interval from the randomizations. The result constitutes a conservative test of statistical significance using these confidence intervals: any observed departure index that occurs outside of this interval would be significantly different from the pre-fire distribution (Manly 2007).

Results

The prescribed fire in the red fir forests reduced the density of small trees from 5257 trees/ha to 1842 trees/ha post-fire, a 65% decrease. Actual mortality rates due to fire for small trees and seedlings are shown along with modeled mortality rates presented for comparison (Table 1). These sets of mortality rates result in differential mortality, and the proportional survivorship is shown in Table 2 and Fig. 3. Given the pre-fire distribution, the departure index has a possible range from −0.39 to +1.61, indicating that the pre-fire, or reference, distribution is already strongly skewed to the left (Fig. 3a, b): there are more seedlings and small trees, and a post-fire distribution can shift further right (to +1.61) than left (−0.39).

The Monte Carlo simulation provided the basis for the tests of significance. As hypothesized, the mean difference between the pre-fire distribution and the null model's 999 random distributions was close to zero (Table 2). The 95% confidence interval ranged from −0.032 to 0.030 (Fig. 4). Any departure index value outside these bounds indicates a

Table 2 Surviving trees shown in each size class as a percentage of the total number surviving in that scenario. Survival of trees post-fire was determined 1 year after the fire

Size class	Initial pre-fire percentage of trees in size class	Observed post-fire percentage of surviving trees	Modeled survival of trees[a]		
			Null (%)	Small-size-biased (%)	Large-size-biased (%)
Seedlings	49.2	49.1	49.2	58.3	27.7
0–5 cm	31.2	26.7	31.2	29.6	35.1
5–10 cm	13.5	14.9	13.5	9.6	22.7
10–15 cm	4.7	6.3	4.7	2.2	10.6
15–20 cm	1.4	3.0	1.4	0.3	3.9
Totals	100.0	100.0	100.0	100.0	100.0
Departure Index value[b]	Range: −0.39 to +1.61	+0.05	−0.0017 (mean)	−0.11	+0.25
Kullback–Leibler criterion[c] (range: 0–1)	na	0.0052	0	0.016	0.054

na not available

[a] In the three modeled scenarios – null mortality (n = 999 simulations), small-, and large-size-biased mortality – survival is the result of applying the mortality rates listed in Table 1

[b] Range refers to the endpoints for the departure index

[c] The Kullback–Leibler criterion is provided as an indicator of the magnitude of the difference between the pre-fire condition and the observed post-fire or modeled distribution

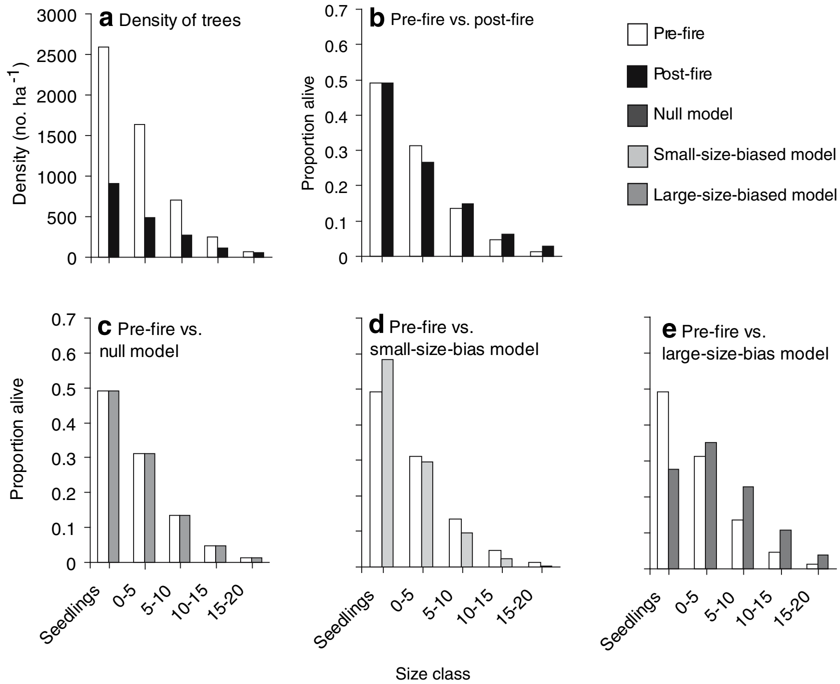

Fig. 3 **a** Counts of living both red fir (*Abies magnifica*) and western white pine (*Pinus monticola*) inventoried by size class in 1997 pre-fire and in 1998 post-fire, Sequoia National Park, California, USA. A reduction in the total count in each size class results from fire. **b** Proportions of the total number of living trees in each size class shown both before and after fire. After the fire, a slight reduction, in proportion, occurs in the first two size classes; a slight increase in proportion occurs in the three larger size classes. **c** Proportional composition comparing the pre-fire distribution and the results of applying the null mortality model. There is no difference between the two distributions. **d** Proportional composition comparing the pre-fire distribution and small-size-biased model. A small increase in the smallest size class precedes a small reduction in the four larger size classes **e** Proportional composition pre-fire distribution and after applying small-size-biased model. Applying the large-size-biased model results in a large reduction in the living trees in the smallest size class and an increase in all other size classes

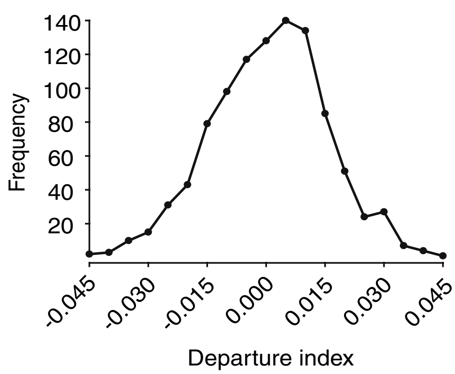

Fig. 4 In a Monte Carlo simulation, a null reference distribution was tested against 999 randomly simulated distributions (J. Battles, unpublished data). The hypothesis tested was that the mean departure index value should be zero (indicating no net difference) and the index should not show any systematic bias. The mean difference between the null distributions and the random distributions was zero; the 95% confidence interval ranged from −0.032 to 0.030

significant difference in the distributions as measured. Similarly, a visual comparison of pre-fire data with the results of the null-model mortality rates in which all size classes experience the same mortality rates shows no difference in the proportional distributions (Fig. 3c).

Comparing the empirical pre-fire and post-fire tree counts indicate biomass loss in every size class (Fig. 3a), but from this graph alone it is unclear whether the fire reduced tree counts across size classes equally. Did fire kill proportionately more large or small trees? Figure 3b depicts the proportional composition of these trees. These data indicate that there has been a slight loss in the first two size classes and a slight increase in the last three (black and white columns). The post-fire distribution is slightly right-shifted toward a larger proportion of large trees. Correspondingly, the departure index indicates that there is a slight right-shift from the pre- to post-fire distribution: +0.05 (−0.39 to 1.61; Table 2). This departure index result implies a small but significant shift in the tree size distribution post-fire as +0.05 falls outside the Monte Carlo confidence interval range (−0.032 to 0.030).

The small-size-biased model results in a moderate left-shift (Fig. 3d). With a higher proportion of small trees surviving than large trees, the forest proportionately shifts toward a composition of smaller trees and seedlings. Correspondingly, the departure index value is small and negative at −0.11 (−0.39 to 1.61). This result is significant as −0.11 falls outside the bounds of the confidence intervals (−0.032 to 0.030).

In the large-size-biased mortality model, proportionately more small trees are killed than large trees (Fig. 3e); consequently, more large trees survive the fire in this model. The departure index reports a moderate right shift toward large trees of +0.25 (−0.39 to 1.61), which is significant (+0.25 is far outside the 95% confidence range of −0.032 to 0.030).

Corroborating results are yielded by the Kullback–Leibler criterion, a measure of the absolute difference between the distributions, which indicates that compared to the pre-fire tree distribution, the null model does not differ at all (Table 2). Actual post-fire data differ less (0.0052) than the small-tree-biased model (0.016), and the large-tree-biased model departed the most (0.054). For comparison of the scale of these values, a distribution that is exactly the reverse of the pre-fire data (switching column 1 with 5, 2 with 4) yields a large difference between the two, as measured by the Kullback–Leibler criterion (0.96) and departure index (+1.22, range −0.39 to 1.61).

Discussion

The departure index successfully detects and describes differences between all of the real and modeled distributions. Further, Monte Carlo simulations indicate that these perceived differences are statistically significant. Likewise, the magnitudes of the departure indices are corroborated by the Kullback–Leibler criterion that was used to detect differences between the distributions.

With regard to the actual post-fire data, there were small but significant shifts in the size distributions of trees after the fire (Table 2). Critically, these differences are also biologically important. Trees in the largest size class of 15–20 cm died at a rate of 24% compared to a rate of 65% for the seedlings. That is evident in the distribution despite the numerical dominance of seedlings in the pre-fire and post-fire populations. Having a way to measure and describe this biologically significant change is important and has, up until now, been difficult or impossible due to constraints of available statistics.

Existing metrics and the departure index

Current distribution metrics are limited in their ability to describe the magnitude and direction of the difference between two distributions (Menning 2003). Some are efficient at detecting differences between distributions but not in describing the difference in terms of magnitude and direction. Many of the statistics have additional limitations in examining the differences between two distributions. With the exception of the Gini coefficient, which scales from 0 to 1, none of the statistics described are standardized. The magnitude and range of the results vary from comparison to comparison (Menning 2003). Each of the metrics considered failed to measure the variables of interest: the magnitude and direction of the departure from one distribution to another (Menning 2003).

In contrast to these measures, the departure index measures the magnitude and direction of shift between distributions. It has a number of mathematical properties that are useful. First, results are standardized with a dynamic response range of two. A standard index range makes comparisons of different analyses possible. A 0.2 shift found in one forest after fire may be compared with a 0.6 value found in a second forest (as long as the same histogram bins were used).

Second, the index is well-behaved. Shifts left are always negative, shifts left are always positive, and neutrally balances changes in total magnitude are always zero. In contrast, the Kullback–Leibler test does not differentiate wherein a difference occurs and does not indicate direction with a sign.

Third, the inclusion of the maximum and minimum endpoints of the range embeds information about the reference distribution itself. Test distributions have a limited ability to depart to the left or right and the index indicates the relative extent of this range.

Fourth, the index is insensitive to changes in absolute amplitude of distributions and is insensitive to sample size. The index will provide the same result with 10 or 1000 elements as long as the two distributions have the same relative distribution (Menning 2003).

Fifth, unlike some of the statistics that are very sensitive to the number of the bins in the histogram, the departure index is relatively insensitive. The departure index always compares values within a bin, and the bin size and count for the reference distribution and the measured distribution are identical. If there are four trees in the 40–50 size class, or if one partitions that into two trees in the 40–45 size class and two more in the 45–50 one, nearly the same departure index value results (Menning 2003).

Sixth, any shape of distribution may be a reference distribution. This property is illustrated when departure index values produced using differently shaped reference distributions are compared. Test data of tree size distributions from 140 actual forest plots were tested with an even reference distribution. Departure index values range from −1 to +1. The process was repeated using an inverse-*j* shaped reference distribution. Values from the inverse-*j* reference distribution (horizontal axis) range from −0.7 to 1.3. A comparison of the two departure index values for each plot for both analyses shows a linear relationship between all

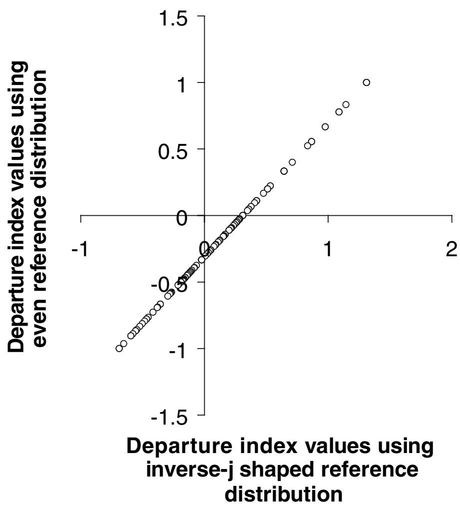

Fig. 5 Departure index values were calculated for tree distributions in 140 plots in a study (Menning 2003). Departure index values were calculated with two different reference distributions. Departure index values on the *x*-axis were calculated using an inverse-*j* shaped distribution. The corresponding departure index value for each plot, calculated with an even distribution, is displayed on the *y*-axis. The linear relationship implies that a linear relationship applies between departure index values calculated using different reference distributions at all locations

plots (Fig. 5). This result demonstrates that a simple linear relationship exists between tests conducted with the two differently shaped departure index analyses at all points. In other words, any reference distribution may be used, and the distribution of results will be identical. As a result, it is not necessary to state the properties of a reference distribution other than its endpoints when making comparisons.

The benefits of using a standard or normal reference distribution are that the index scales nicely between -1 and +1. As a disadvantage, an even reference distribution may be far from the realistic range of the test distribution. This makes assessments of individual distributions harder to understand. If a forest-wide size-class distribution averaged -0.6, for example, a departure index value of “0” would actually indicate a strong rightward departure from the average – even though 0 intuitively indicates no change.

In contrast, there are benefits to using a more realistic reference distribution, such as an overall average. First, each distribution is compared directly with the mean distribution. A slight shift left or right from the mean overall distribution is indicated by a sign difference. Zero indicates no departure from the overall average. Second, the asymmetric range of the index demonstrates that the test distribution could depart further in one direction than the other. Used with a non-symmetric reference distribution, however, the index can range between irregular endpoints, such as -0.6 to +1.4.

Seventh, the departure index does not require data to be normally distributed. The results of multiple departure index analyses, however, are normally distributed, as per the Central Limit Theorem, and this property allows the departure index to be tested for sensitivity and significance with a Monte Carlo approach. In other words, the departure index is a relatively sensitive test able to resolve small differences in distributions.

Limitations of the departure index

As with any statistic, there are limitations to what the departure index is able to measure. While the departure index is effective at finding and measuring horizontal inequalities (departures toward the left or right ends of a distribution), it is not meant to detect differences when a difference occurs equally on either side of the mean. This trait provides an additional argument for using a realistic reference distribution. If the reference distribution were different, for example, such as an inverse-*j* shaped distribution, the departure index would be likely to respond to a change on either side of the mean. The Kullbac–Leibler criterion is helpful as a complementary test, as it provides a measure of the absolute difference between any two distributions regardless of their symmetry. A departure index value of 0 with a positive Kullback–Leibler criterion would indicate that there was a difference between the reference and test distributions but that there was no asymmetric shift from the reference to the test distribution.

Similarly, the departure index is not meant to detect linear changes in amplitude. In certain cases, this is not just a limitation but also a benefit. The departure index is particularly useful in situations in which one wishes to determine relative change between ordered classes, not absolute change. In the case of a fire in a forest, we know that there is a loss of biomass in all or most size classes. We do not know, before using the departure index, if the biomass reduction was equal in all size classes, or if it affected some classes more than others. We can state that yes, there was a net loss of biomass. Using the departure index, however, we can quantify a shift toward larger size classes of trees. In other words, while many trees burned, proportionately more small trees were burned than large trees. This is a very important result for understanding forest stand structure as well as the effects of fire. Further, this property of being insensitive to changes in amplitude demonstrates how the index is insensitive to sample size as long as two distributions are distributed the same way (Menning 2003).

Conclusion

We believe the departure index to be a useful measure of the magnitude and direction of the departure of a test distribution from any reference distribution. The metric has useful mathematical properties not provided by other statistics currently available. We anticipate it having a wide range of applications, ranging from topics as diverse as forest ecology and neighborhood economic balances to population demography, variation from equilibrium states and even variability in group competition. Fisheries scientists, for example, might use the departure index to measure the change in gravel texture in a spawning bed due to a flood event. Economists wishing to characterize differences in income distributions could use the index to measure the differences between two cities or ethnic groups. Readers may test their own data using the departure index calculator published on-line by the journal (see Electronic Supplementary Material).

Acknowledgements This research has been supported by the University of California, Berkeley; Sequoia and Kings Canyon National Parks of the National Park Service (NPS), Biological Resources Division (BRD) of the USGS; Duke University; National Interagency Fire Center (NIFC); and the Canon National Parks Science Scholars Program (CNPSS). David Newburn and field crews funded by Dr. John Battles collected many of the data used in the fire-mortality case study. Dr. Wayne Sousa, Dr. Nathan Stephenson and three anonymous reviewers provided valuable reviews, insights and suggestions. This research complies with all legal requirements of the United States where the research was conducted.

References

Begon M, Harper JL, Townsend CR (1996) Ecology: individuals, populations, and communities, 3rd edn. Blackwell Science, Oxford

Bergqvist G (1999) Wood volume yield and stand structure in Norway spruce understorey depending on birch shelterwood density. For Ecol Manage 122:221–229

Burnham KP, Anderson DR (2002) Model selection and multimodel inference: a practical information–theoretic approach, 2nd edn. Springer, Heidelberg

Burnham KP, Anderson DR, Thompson WL (2000) Null hypothesis testing: problems, prevalence, and an alternative. J Wildl Manage 64:912–923

Gaymer CF, Himmelman JH, Johnson LE (2001) Distribution and feeding ecology of the seastars *Leptasterias polaris* and *Asterias vulgaris* in the northern Gulf of St Lawrence, Canada. J Mar Biol Assoc UK 81:827–843

Jimenez JJ, Decaens T (2000) Vertical distribution of earthworms in grassland soils of the Colombian Llanos. Biol Fertil Soils 32:463–473

Manly BFJ (2007) Randomization, bootstrap and Monte Carlo methods in biology, 3rd edn. Chapman & Hall/CRC, Boca Raton

Menning KM (2003) Forest heterogeneity: methods and measurements from extensive field plots, fire modeling, and remote sensing of the mixed conifer forest of the Southern Sierra Nevada, USA. PhD thesis. University of California, Berkeley

Pitcher DC (1987) Fire history and age structure of red fir forests in Sequoia National Park, California. Can J For Res 17:582–587

Payette S, Bhiry N, Delwaide A, Simard M (2000) Origin of the lichen woodland at its southern range limit in eastern Canada: the catastrophic impact of insect defoliators and fire on the spruce–moss forest. Can J For Res 30:288–305

Savage M (1994) Anthropogenic and natural disturbance and patterns of mortality in a mixed conifer forest in California. Can J Forest Res 24:1149–1159

Weiner J (1990) Asymmetric competition in plant populations. Trends Ecol Evol 5:360–364

Weiner J, Solbrig OT (1984) The meaning and measurement of size hierarchies in plant populations. Oecologia 61:334–336

Wiegand K, Ward D, Thulke H-H, Jeltsch F (2000) From snapshot information to long-term population dynamics of Acacias by a simulation model. Plant Ecol 150:97–114

Yang Y, Monserud RA, Huang S (2004) An evaluation of diagnostic tests and their roles in validating forest biometric models. Can J Forest Res 34:619–629

Zar JH (1999) Biostatistical analysis. Prentice Hall, Upper Saddle River