Article

Composite Estimators for Forest Growth Derived from Symmetric, Varying-Length Observation Intervals

Francis A. Roesch

Southern Research Station, USDA Forest Service, 200 WT Weaver Blvd., Asheville, NC 28804, USA;
Francis.A.Roesch@USDA.gov; Tel.: +01-828-257-4871

Received: 22 March 2019; Accepted: 9 May 2019; Published: 11 May 2019

Abstract: Estimates of growth or change in a forest population parameter for a specific length of time, such as cubic meters of wood per hectare per year, are often made from sample observation intervals of different lengths of time. For instance, a basic building block of growth estimators in forest inventory systems is often the annual mean of the first differences of all observations for a particular year, regardless of observation interval length. The aggregate differences between successive observations on re-measured forest sample plots can be viewed as a linear combination, while forest growth is usually assumed to be non-linear. Bias can be assumed to exist whenever a linear combination is used to estimate a specific segment of an underlying non-linear trend. The amount of bias will depend upon the relationship of the intended estimation interval relative to the set of observation intervals. Here, three specific segments, relative to each year of interest, form the bases for a standard set of three estimands. Bias-ratio-adjusted composite estimators for use with observations made on alternative sets of symmetric interval lengths are compared in a simulation against this standard set of estimands. The first estimand has a one-year basis, the second has a five-year mid-interval basis, and the third has a five-year end-of-period basis. For the first and second bases, the initial results clearly show a logical ordering of bias and mean-squared error by observation interval length relative to the target interval length. As expected, some deviance from these clear trends are shown for the end-of-period basis. In the presence of three simple distributions of symmetric measurement intervals, the bias-ratio adjustments and subsequent composite estimators are shown to usually be effective in reducing bias and mean-squared error, while being most obviously effective for the most disparate distribution of intervals and for the end-of-period basis.

Keywords: national forest inventory (NFI); simulation; annual growth; estimation

1. Introduction

National forest inventories and monitoring systems provide a valuable service by collecting and maintaining data, and by presenting summaries and interpretations of those data for public use. The data are obtained through well-defined observations made over specific land areas and through specific periods of time. The accompanying analyses of, and even the very presentation of these data often have underlying assumptions that may not be immediately obvious to a new user. Additionally, even though these assumptions are usually well documented and accessible to experienced users, their full implications with respect to the users' expectations often need further exploration. One such set of assumptions concerns the use of the annual mean of the first differences of all observations for a particular year, regardless of the length of time that passed since the previous observation. To this end, a standard set of estimands for forest growth was defined in Roesch [1] to address a varying set of potential user expectations for the United States Department of Agriculture (USDA)'s Forest

Inventory and Analysis Program (FIA), which is the national forest inventory for the United States of America (USA). The estimands differ in their basis, with each basis being defined by its period length and viewpoint orientation (i.e., the reference year within the period.) Estimand 1 has a one-year basis, Estimand 2 has a five-year mid-interval basis, and Estimand 3 has a five-year end-of-period basis. Also, in that work, a number of candidate methods of bias adjustment were tested for their effectiveness in the presence of positively asymmetric measurement intervals, such as those that might arise when a budget reduction leads to increasing interval lengths. That is, the interval lengths were positively asymmetric from the intended length. Here, I evaluate bias-ratio adjustments in a simulation for effectiveness under three simple distributions of symmetric measurement intervals. The first of these distributions is a discrete simplification of one that might arise when the re-measurement of forest plots is randomly reordered. For all three bases, the intent is to determine the circumstances under which bias-ratio adjustments are prudent and likely to be effective in substantially reducing both bias and mean-squared error.

The observation of broad-area forest growth by any national forest inventory (NFI) constitutes an observation of a mixture of more refined forest growth distributions. The components of the mixture could be defined at different scales, such as one growth curve for each general forest type, or, more specifically, such as one growth curve for each tree species within each narrowly defined forest type. However one defines the components of the mixture, it will invariably be true for an NFI that some of the components will be well understood while others will be poorly understood. It should be recognized that the most useful of procedures for an NFI will be those that work well regardless of the underlying distribution of components. For it is not only the goal of most NFI efforts to estimate forest growth for the entire country, but to also be able to estimate forest growth over any sub-area and time period of interest. This recognition is the impetus for the approach taken here. The approach simply assumes that there exists a final joint distribution or growth curve resulting from the mixture of underlying components. It is safest to assume that the underlying components themselves are non-linear and that the resulting mixture distribution is non-linear. Couple this with the observation that the aggregate differences between successive observations on re-measured forest plots for each of the different interval lengths constitute different linear combinations, and we realize that a unique bias can be assumed to exist when each of these linear combinations is used to estimate the mean of a specific segment (or basis) of the underlying non-linear growth trend. It is the evaluation of this bias and the amelioration of its effects that are of primary interest in this paper.

2. Materials and Methods

2.1. Estimators

As noted, the amount of bias for a specific interval length is annually specific, rendering this a particularly difficult problem to model in an unsupervised inventory processing system, while solutions that work well within an unsupervised system are highly desirable for large-scale efforts such as national forest inventory systems. Although multiple estimation systems were discussed in Roesch [1], to which the interested reader is referred, here, the focus is on the bias-ratio estimators discussed there, as they are expected to be the most robust in an unsupervised system.

The assumptions and theory in this section follow the assumptions and theory discussed in Roesch [1]. To review, let $\delta_{bi(m)}$ represent an estimator of growth or change, Δ_{bi} in a population parameter over a specific temporal interval length (or basis b), such as basal area per hectare per year, derived from an observation interval of length m , either centered on year i (the mid-year viewpoint) or ending in year i (the end-of-period (EoP) viewpoint).

A naïve estimator, referred to below as the sample mean, is the annual mean of the first differences of all observations for year i regardless of observation interval length. In the case of Estimands 1 and 2, this means all intervals centered on year i , while, in the case of Estimand 3, this means all intervals

ending in year i . The premise is that an estimator that usually dominates the naïve estimator in terms of having a smaller bias or mean-squared error (MSE) is worthy of further consideration.

Roesch [1] gave a few examples of composite estimators to combine the $\delta_{bi(m)}$ estimators and pointed out that only when the basis interval length b is equal to the observation interval length m , and growth is linear under the mid-year viewpoint will a possibility exist for $\delta_{bi(m)}$ to be unbiased.

Composite estimators can be used to combine multiple estimators by assigning a weight to each of the component estimators. The weights (which sum to 1) correspond to an investigator's "degree of belief" in each estimator's relative applicability to an estimand of interest. When the component estimators are biased, or might be biased, as in the case at hand, the usual practice is to base one's degree of belief on the inverse of an estimate (mse) of the mean-squared error (MSE) of each estimator (Green and Strawderman [2]). For any basis b , the mse -weighted composite estimator would be

$$C_{bi} = \sum_{all\ m} \hat{\alpha}_{bi(m)} \delta_{bi(m)}, \quad (1)$$

where

$$\hat{\alpha}_{bi(m)} = \frac{1/mse(\delta_{bi(m)})}{\sum_{all\ m} 1/mse(\delta_{bi(m)})}. \quad (2)$$

Estimates of the bias and variance, the components of MSE, could be obtained either from an earlier sample or from the current sample.

2.2. Bias-Adjusted MSE-Weighted Composite Estimators

Highly biased component estimators can be problematic in composite estimation. Roesch [1] pointed out that, if prior information is available on the magnitude of the expected bias, then one could attempt to adjust for the bias prior to forming the composite estimator; thus, bias adjustments for alternative measurement intervals could be estimated under a wide range of models. The strategy developed in that paper, utilizing a simple ratio model, is also used here.

Assume that, at year i , there exists a difference (or bias) for each observation interval of length m relative to the basis interval length:

$$B_{bi(m)} = \Delta_{bi} - \Delta_{mi}, \quad (3)$$

where, as above, m and b are either centered on year i (the mid-year viewpoint) or end in year i (the EoP viewpoint), and some estimator of $B_{bi(m)}$ is available and denoted as $\hat{b}_{bi(m)}$.

The bias ratio for each m , relative to the basis, would then be

$$\hat{r}_{bi(m)} = \frac{\bar{d}_{bi}}{\bar{d}_{mi} - \hat{b}_{bi(m)}} \quad (4)$$

where $\bar{d}_{bi}$ is the sample mean annual cubic meter volume growth per hectare for a period of length b , centered on year i , and $\bar{d}_{mi}$ is the sample mean annual cubic meter volume growth per hectare for all observations from intervals of length m , centered on year i .

Once $\hat{r}_{bi(m)}$ is estimated from an adequately sized sample, it could be used in subsequent samples to adjust results for measurement intervals of size m to its basis length equivalent:

$$\hat{\delta}_{bi(m)} = \hat{r}_{bi(m)} \hat{d}_{mi}$$

where $\hat{d}_{mi}$ is the subsequent (or current) sample mean annual cubic meter volume growth per hectare for all observations from intervals of length m , centered on year i .

Substituting $\hat{\delta}_{bi(m)}$ into Equation (1) results in

$$\hat{C}_{bi} = \sum_{all\ m} \hat{\alpha}_{bi(m)} \hat{\delta}_{bi(m)}. \quad (6)$$

When the weights are developed a priori (from an earlier sample or “fit” data), I use the designation of $\hat{\alpha}_{bi(m,f)}$, and, when they are developed from the current sample (or “test” data), I use the designation of $\hat{\alpha}_{bi(m,t)}$. This can lead to any number of potential estimators. The results from three potential estimators, which utilize bias-ratio adjustments from the fit data applied to the test data, are given below. The first estimator is the unweighted bias-adjusted mean:

$$BRA_{bi(t)} = \sum_{all\ m} \left(\frac{1}{n(m)} \right) \hat{\delta}_{bi(m)}, \quad (7)$$

where $n(m)$ indicates the number of interval length categories. Then, the a-priori weighted composite estimator is given as

$$wBRA_{bi(f)} = \sum_{all\ m} \hat{\alpha}_{bi(m,f)} \hat{\delta}_{bi(m)}, \quad (8)$$

and, finally, the contemporaneously weighted composite estimator is given as

$$wBRA_{bi(t)} = \sum_{all\ m} \hat{\alpha}_{bi(m,t)} \hat{\delta}_{bi(m)}. \quad (9)$$

2.3. Simulation

The data from re-measured FIA forested plots from 125 FIA survey units in 34 states in the United States were used to build 125 simulated populations. These survey units are usually groups of counties within a state, and vary widely both in total land area and in the proportion of forest land. The interested reader will find the data available in a public database known as the FIADB [3]. For the purposes of this study, to build the sampled populations, I limited the data to those collected under the annual panel design, as described in Bechtold and Patterson [4], in states with a target cycle length of five years or less, and for which adequate data were collected for the target estimation years of 2004 through 2008. Figure 1 is a map showing these states, within the United States, in blue. These data are the same as those used in Roesch [1] and had ending year measurements made between 1998 and 2016, inclusive. As in the previous citation, data from measurement intervals of less than one year were eliminated because they were not originally intended to be analyzed as growth intervals. The data from each of the 125 populations were used to create a compatible set of five matrices, one for each of the four components of forest growth and one for initial annual volume, as defined below. All or any subset of these five matrices could be considered a population for simulation purposes.

Figure 1. The states (in blue), within the United States of America (USA), from which the data were collected by the United States Department of Agriculture (USDA) Forest Service Forest Inventory Program (FIA) and used in this study. These data are publicly available in the Forest Inventory and Analysis Database (FIADB) [3].

Because the survey units were so diverse, each was treated as a separate population in the simulation, rather than using a method to construct multiple populations from a single dataset. That is, each population is assumed to represent the set of forested conditions within its associated survey unit. These estimation units cover a range in latitude in the northern hemisphere from approximately 25° north (N) to 47° N.

The methods used here to construct the annual populations were previously described in Roesch [5], for a smaller dataset, and later used in Roesch et al. [6]. The data from each forest sample plot in each survey unit were used to create a row with a sequence of 21 successive, annual values, for each of five compatible variables of interest, commonly referred to as the components of growth, although only one of those variables (live growth) was used for the work reported in this paper. The five variables of interest consist of the volume of wood ($\text{m}^3 \cdot \text{ha}^{-1}$) at the beginning of each year t (V_t), the growth of wood on all living trees (L_t) during the year ($\text{m}^3 \cdot \text{ha}^{-1} \cdot \text{year}^{-1}$), the volume of all living trees entering the population (E_t) during the year ($\text{m}^3 \cdot \text{ha}^{-1}$), the volume of all trees harvested (H_t) during the year, and the volume of all trees dying (M_t) during the year. Compatibility would then require that

$$V_2 = V_1 + L_1 + E_1 - H_1 - M_1. \quad (10)$$

The data from all measurements of each forest plot were converted into a 21-year series of values, for each variable of interest, as described below. Most of the ground plots were measured three or more times during the 21-year period, resulting in the observation of at least two growth periods during the 21-year time period. The allocation of the observed values proceeded as follows: firstly, any harvests were allocated to a particular year in the observation interval. Because exact harvest times were unknown, harvested volume was randomly allocated to a year within each observation interval. Linear interpolation between observations and extrapolation beyond the limits of observation were used to obtain an initial value for the live growth, entry, and mortality components for each year, as well as a starting cubic meter per hectare value in the first year, with annual adjustments made as necessary when high levels of harvest and mortality reduced all volume to zero. This completion of this process resulted in Set 1.

The underlying assumption for the matrices in Set 1 is that each 21-year series in Set 1 represents the mean of a variable of interest for a forested condition class that is composed of similar but unique land segments with similar developmental characteristics through the 21-year period. To achieve this diverse set of land segments, random variance was applied 100 times at two levels to each row of Set 1. That is, each line in Set 1 was used to create 100 lines in Set 2, the population of annual values for each survey unit. In level 1, in order to add variance but maintain trend, all values for each growth component in each row were multiplied by a random variate, unique for the row, drawn from an N

(1, 0.025) distribution. A second level of variance was effected annually by multiplying the result of level 1 for each annual value in each row by a random variate, unique for each value, drawn from an $N(1, 0.0025)$ distribution. To maintain compatibility, the initial annual volume matrix was then re-calculated, starting at the first year of non-zero volume observation for each line and applying the growth components recursively. As mentioned above, this study uses only the live growth matrix from each of the resulting populations. This process can be viewed as transforming the observations for each remeasured plot into a 100-row by 21-year condition class in which the 21-year series of values in each row is randomly different from the other rows in the condition class. The characteristics of the originating FIA survey units are, for the most part, retained in these populations, and some of these populations may be similar to specific areas in other national forest inventories. The populations are diverse in size and composition, as illustrated in Figure 2, which gives the distribution of sizes for the 125 populations in the top graph and the standard deviation of the within-population annual variance of cubic meter volume growth of living trees.

Figure 2. The size in hectares for the 125 populations (a), and the standard deviation of the within-population annual variance of cubic meter volume growth (b).

To create a set of populations to test the effects of interval length on the estimates of wood growth ($\text{m}^3 \cdot \text{ha}^{-1} \cdot \text{year}^{-1}$) for an intended target length of five years, I expanded the 125 live growth populations, by distributing the annual growth throughout the year in a manner intended to mimic intra-annual growth. This was done to partially account for two factors. The first is that, in much of the study area, forest plot measurements are taken throughout the year, while forest growth occurs for only part of the year, and the growth season varies through the study area. The second is that the rate of growth varies throughout the growing season. Unfortunately, our available data are not nearly refined enough to allow the reliable estimation of these factors at the scale at which we know they vary. Therefore, in the simulation, a simple approximation is used to distribute the annual growth values throughout an approximate growing season spanning 80% of the year. To accomplish this, each year was divided into 10 equal-length segments or decimal-years (dy). Because the actual apportionment of growth to dy within each of the underlying original populations for each year is unknown, the annual growth was apportioned to dy 1 through 10 of each year in the proportions of 0, 0.05, 0.1, 0.2, 0.2, 0.18, 0.12, 0.1, 0.05, and 0, respectively. This created 210 columns within each row from the 21 years available (from 1995 through 2015).

The decimal-year populations were then used to examine the effects of different distributions of interval lengths on the different estimands of interest. The distributions each combine five of the nine interval lengths of 1, 3, 4, 4.5, 5, 5.5, 6, 7, and 9 years, which are denoted in the text as Y1, Y3, Y4, Y4.5, Y5, Y5.5, Y6, Y7, and Y9, respectively. The sampling simulation consisted of taking 100 samples of size 1000 for each interval length from each of the 125 populations, and calculating each estimator for each sample. Initially, I present the simulation results that are obtained from the treatment of the sample of each interval length as a stand-alone sample. Following that, I combine the samples of three different combinations (or distributions) of five of the nine interval lengths within each population into a single sample in each iteration. I then compare the results with respect to three estimands, for each of the three distributions. The entire sampling simulation, including data and R-code, is available from the corresponding author on digital versatile disc (DVD) upon request.

3. Results

Figure 3, Figure 4, and Figure 5 show, for Estimands 1, 2, and 3, respectively, the simulation mean volume growth (a), the simulation mean bias (b), and the simulation mean of mean-squared errors (c) for the sample means from each of the nine interval lengths, for the years 2004–2008, over 100 samples of size 1000 from each of the 125 populations.

Figure 3. The simulation mean volume growth (a), the simulation mean bias (b), and the simulation mean of mean-squared errors (c) for Estimand 1 (the annual mean) for the means from each of the nine interval lengths of 1, 3, 4, 4.5, 5, 5.5, 6, 7, and 9 years, applied to the years 2004–2008, from 100 samples of size 1000 from each of the 125 populations.

Figure 4. The simulation mean volume growth (a), the simulation mean bias (b), and the simulation mean of mean-squared errors (c) for Estimand 2 (the five-year centralized moving average) for the means from each of the nine interval lengths of 1, 3, 4, 4.5, 5, 5.5, 6, 7, and 9 years, applied to the years 2004–2008, from 100 samples of size 1000 from each of the 125 populations.

Figure 5. The simulation mean volume growth (a), the simulation mean bias (b), and the mean of the simulation mean-squared errors (c) for Estimaand 2 (the five-year end-of-point (EoP) moving average) for the means from each of the nine interval lengths of 1, 3, 4, 4.5, 5, 5.5, 6, 7, and 9 years, applied to the years 2004–2008, from 100 samples of size 1000 from each of the 125 populations.

Figures 3b and 4b clearly show the bias that is incurred as the length of the measurement interval moves away from the basis interval length. Note that subplots Figures 3a and 4a are identical because they are simply the simulated sample realization of all intervals of each length. Figure 3b shows the plot for Y1 to be the most symmetric around zero, while Figure 4b shows the plot for Y5 to be the most symmetric around zero. This symmetry around zero is reduced in both of these subfigures as the interval length moves away from the basis interval length. Because growth is not linear, these subfigures also show that there are cases for individual years when the value from one or more of these alternative plots will show less bias; however, the overall bias is still greater.

The following important question remains: Is that increase in bias large enough to have a significant impact on mean-squared error (MSE)? Figures 3c and 4c show that, for the one-year basis and the five-year mid-year basis, the answer is no. In both of these subfigures, the MSE plots are clearly ordered for each year in inverse proportion to interval length. That is, the variance reduction achieved through the longer interval lengths outweighs the increase in bias (squared). Therefore,

although attempts at bias reduction when combining these intervals may be somewhat successful, the subsequent effect on overall MSE would be expected to be minimal.

Figure 5, for the five-year EoP basis, tells a somewhat more complex story than Figures 3 and 4. In this figure, all sample interval lengths are applied to the final year of the five-year basis. For that reason, the plot for a one-year observation interval (Y1) in Figure 5a matches the plots for Y1 in Figures 3a and 4a, but the plots for all other intervals are different. Owing to the effect of lag bias, the plots for the four-year observation interval (Y4) and the 4.5-year observation interval (Y4.5) show slightly less overall bias than the plot for the five-year interval (Y5), and the plot for Y4.5 shows the closest match to the basis-matching symmetry noted in Figures 3 and 4. That is, the sample contributing to the Y5 plot for year y contains intervals occurring throughout year y to predict the values for the five-year period ending at the end of year y . Therefore, a small part of the growth observation occurs outside of the basis of interest. Also, the overall bias is greater for the non-basis intervals (i.e., those other than Y4, Y4.5, and Y5) in Figure 5b than in Figures 3b and 4b. Although it seems that successful bias adjustment might be beneficial overall, it is much less clear what the model for bias adjustment should be. The symmetry observation suggests that it might be reasonable to “prefer” an adjustment to Y4.5, Y4, and Y5, in that order, for the five-year EoP basis, when sufficient data are available at those alternative (to Y5) interval lengths. If one rather preferred that the estimates be drawn from within the basis of interest, then one should prefer an adjustment to Y4, Y4.5, and Y5, in that order.

Consider the three following symmetric distributions, each consisting of five of the nine interval lengths given above in equal proportions:

Distribution 1 (d1) uses the five interval lengths of 4, 4.5, 5, 5.5, and 6 years;

Distribution 2 (d2) uses the five interval lengths of 3, 4, 5, 6, and 7 years;

Distribution 3 (d3) uses the five interval lengths of 1, 3, 5, 7, and 9 years.

These distributions are symmetric around the target observation interval length of five years. For each distribution, 1000 rows were sampled in equal proportions for each of the five intervals, in each of the 125 populations, for 100 iterations. Additionally a “fit” dataset was developed using the same methods, which comprised 5000 rows from each population (five iterations of samples without replacement of 1000 rows each).

Above, I presented the simulation results that were obtained from the treatment of the sample of each interval length as a stand-alone sample. Below, I combine the samples of the five interval lengths within each population into a single sample in each iteration. I then compare the results with respect to three estimands, for each of the three distributions.

Figure 6, Figure 7, and Figure 8 show, for Estimands 1, 2, and 3, respectively, the simulation mean bias (left) and the simulation mean of mean-squared errors (right) for each of four estimators (the sample mean, $BRA_{bi(t)}$, $wBRA_{bi(f)}$ and $wBRA_{bi(t)}$ for each of the three distributions. These estimators are plotted with $MeanPPer$, which is simply the mean of the sample interval most closely corresponding to the basis interval for each distribution. $MeanPPer$ would not usually be considered, because normally it would have a small sample size, and the interval length from which it is drawn might not be one of a small, discrete set of intervals, as it is in this simulation. Therefore, in this simulation, favorable results in terms of MSE for $MeanPPer$ simply indicate that filtering out the observations arising from more diverse interval lengths might be a reasonable approach.

1-Yr Basis

Observation Distribution d1

Observation Distribution d2

Observation Distribution d3

Figure 6. The simulation mean bias in the left graphs (a, c, and e) and the simulation mean of mean-squared errors in the right graphs (b, d, and f) for Estimand 1 (with a one-year basis), for each distribution d1 through d3, from top to bottom, respectively. These summary statistics were compiled after 100 iterations of composite samples of size 1000, for each of five sample observation lengths, within each of the 125 populations.

5-Yr MidYear Basis

Observation Distribution d1

Observation Distribution d2

Observation Distribution d3

Figure 7. The simulation mean bias in the left graphs (a, c, and e) and the simulation mean of mean-squared errors in the right graphs (b, d, and f) for Estimand 2 (with a five-year mid-year basis), for each distribution d1 through d3, from top to bottom, respectively. These summary statistics were compiled after 100 iterations of composite samples of size 1000, for each of five sample observation lengths, within each of the 125 populations.

5-Yr EoP Basis

Observation Distribution d1

Observation Distribution d2

Observation Distribution d3

Figure 8. The simulation mean bias in the left graphs (a, c, and e) and the simulation mean of mean-squared errors in the right graphs (b, d, and f) for Estimand 3 (with a five-year EoP basis), for each distribution d1 through d3, from top to bottom, respectively. These summary statistics were compiled after 100 iterations of composite samples of size 1000, for each of five sample observation lengths, within each of the 125 populations.

These summary statistics were compiled after 100 iterations of composite samples of size 1000, for each of five sample observation lengths, within each of the 125 populations, for each of the three distributions. The estimand in Figure 6 is Estimand 1, the population's annual average for each year of interest, while the estimand in Figure 7 is Estimand 2, the population's five-year annualized mean

centered on each year of interest, and the estimand in Figure 8 is Estimand 3, the population's five-year annualized mean attributed to the end year in each series.

4. Discussion

Figure 6, for Estimand 1 with a one-year basis, shows that there is usually some advantage to bias adjustment, as well as some additional benefit to inverse-MSE-weighted composite estimation. All bias graphs (Figure 6a,c,e) show the alternative estimators to be more symmetrical around zero than the sample mean (or naïve estimator). This is what one would expect. Additionally all of the MSE graphs (Figure 6b,d,f) show the alternative estimators to almost always have a smaller MSE than the sample mean. This is true even for *MeanPPer*, which uses one-fifth of the data being used by the other estimators. *MeanPPer* is usually greater in MSE than the other alternative estimators ($BRA_{bi(t)}$, $wBRA_{bi(f)}$, and $wBRA_{bi(t)}$) in Figure 6b,d, for the two least dispersed distributions (d1 and d2). This is also true in Figure 6f for the most dispersed distribution (d3), although the distinction is less clear. Note also that, in the corresponding bias graph (Figure 6e), there is virtually no distinction between the alternative estimators. At least two factors contribute to this. The first is that the closest interval in d3 to the estimand of interest is the one-year interval, meaning that all observation intervals contributing to *MeanPPer* are very close to the annual interval of interest. The second factor is that, even though these distributions are all symmetrical around the intended observation interval of five years, they are increasingly asymmetrical (from d1 to d3) from the one-year basis for Estimand 1.

Figure 7, for Estimand 2, shows that there is not much need for bias adjustment or inverse-MSE weighting for d1 and d2. Recall that d1 uses the interval lengths of 4, 4.5, 5, 5.5, and 6 years, and d2 uses lengths of 3, 4, 5, 6, and 7 years. This would seem to indicate that these interval lengths, when symmetrically distributed, are not sufficiently far from the five-year mid-year basis to have a negative impact on estimation of Estimand 2. In fact, when examining Figure 7a, one can see that attempts at bias adjustment resulted in bias plots which are less symmetrical around zero, which could indicate a tendency to increase overall bias. For d3, the bias adjustment appears to have had a small advantage, resulting in more symmetrical plots around the zero line than the sample mean in Figure 7e, which, for four of the five years, contributed to the small reduction in MSE seen in Figure 7f.

In considering Estimand 3 with a five-year EoP basis, Figure 8 shows some advantage to bias adjustments given sampling interval distributions d2 and d3, and a clear additional benefit to the use of inverse-MSE-weighted composite estimation is seen for all three distributions d1, d2, and d3 over use of the naïve sample mean.

These simulation results show that, if the estimand of interest and the center of the target observation interval are in close correspondence, then a symmetrical distribution of actual observation intervals can have a level of deviation from the target interval length that would arise from logistically driven changes in the order of field plot observation without adversely affecting growth estimates. This is arguably the case when observation distribution d1 is applied to the five-year mid-year basis and closely represented by the simulation results in Figure 7a,b. The advantages afforded by observation intervals which are approximately symmetrical to the basis of interest are further reinforced in the remainder of Figure 7. Figures 6 and 8 show substantial improvement in terms of MSE for the alternative estimators over the naïve estimator, while that improvement is less obvious in Figure 7. This improvement was not necessary for Estimand 2, while it was necessary for Estimands 1 and 3. This is due solely to the fact that only for Estimand 2 is the distribution of the varying growth intervals symmetrical with respect to the estimand.

Forest growth is an inherently non-linear phenomenon, while many forest inventory algorithms in NFI systems rely implicitly on linearity assumptions. These linearity assumptions are often useful and pragmatic; however, their limitations often go untested. For instance, the class of sample designs being considered in this paper (temporally overlapping panelized designs) does not have an inherent assumption of linearity, but both the implementation of the design and the choice of data aggregation methods almost always do have linearity assumptions. This work recognizes that the “realized”

sample is often different from what was intended by the sample design, while the theory justifying or underlying most forest monitoring estimation systems relies solely on a description of the sample design rather than on the “actual” or “realized” sample. The work reported here is a further exploration of the bias that can arise as a result of these differences and a mitigating bias adjustment procedure. Note that, if forest growth actually was a linear phenomenon and all growth intervals were actually symmetric with respect to the intended interval, then the bias investigated in this series of simulations with respect to all three estimands would not exist.

As in previous work, the idea of an established set of target estimands proved useful in that it provided a forum for discussion about how estimators can be modified to address various users’ expectations when they encounter the term “average annual growth”.

5. Conclusions

The simulation results demonstrate that, even if the distribution of the varying-length observation intervals is symmetric with respect to the intended observation interval length, the degree to which the procedures described here will be necessary or helpful depends both upon the degree of local curvilinearity of the underlying growth function and upon the degree of symmetry between the estimand of interest and the actual observation intervals. The relative advantages of bias adjustment and composite estimation depend upon at least three factors: the temporal length of the estimand of interest, the locus of the observation intervals relative to the locus of the estimand of interest, and the variability in observation interval lengths. Specifically, we see that there can be substantial advantage to bias adjustment and inverse-MSE-weighted composite estimation when the mean observational interval length is substantially different from the length of the basis (as in Estimand 1) and when the basis interval center does not correspond to the differing-length observation interval centers (as in Estimand 3).

Funding: Funding for this research was provided by the author’s employer, the USDA Forest Service, Southern Research Station, Forest Inventory and Analysis Unit, Knoxville, TN, USA.

Acknowledgments: This manuscript was improved through the incorporation of review comments offered by Kathryn Roesch and two anonymous reviewers.

Conflicts of Interest: The author declares no conflicts of interest. The funders had no role in the design of the study; in the analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Roesch, F.A. Composite Estimators for Growth Derived from Repeated Plot Measurements of Positively-Asymmetric Interval Lengths. *Forests* **2018**, *9*, 427. [CrossRef]
2. Green, E.J.; Strawderman, W.E. Combining inventory estimates with possibly biased auxiliary information. *For. Sci.* **1990**, *36*, 693–704.
3. FIADB. Available online: https://apps.fs.usda.gov/fia/datamart/CSV/datamart_csv.html (accessed on 7 May 2019).
4. Bechtold, W.A.; Patterson, P.L. *The Enhanced Forest Inventory and Analysis Program-National Sampling Design and Estimation Procedures*; U.S. Department of Agriculture Forest Service, Southern Research Station: Asheville, NC, USA, 2005; General Technical Report SRS-80. Available online: <http://www.srs.fs.fed.us/pubs/20371> (accessed on 18 April 2019).
5. Roesch, F.A. Toward robust estimation of the components of forest population change. *For. Sci.* **2014**, *60*, 1029–1049. [CrossRef]
6. Roesch, F.A.; Schroeder, T.C.; Vogt, J.T. Effects of Cycle Length and Plot Density on Estimators for a National-Scale Forest Monitoring Sample Design. *Forests* **2017**, *8*, 325. [CrossRef]

© 2019 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).