UNDERSTANDING STATISTICS, 1(3), 191–203

TEACHING ARTICLE

Calculating a Factorial ANOVA From Means and Standard Deviations

Barry H. Cohen
Department of Psychology
New York University

In this article, I present an alternative formula for the calculation of a factorial analysis of variance (ANOVA), which requires only the mean, standard deviation, and size for each cell of the design, rather than the individual scores. This new method allows a modern hand-held calculator to do most of the work, while still giving students the educational experience of working directly with data. An example is given, in which the new method is applied to a published table of data from a two-way unbalanced ANOVA design. I argue that the new formula is not just easier to use than the traditional (raw-score) calculation formula (especially when dealing with higher order factorial designs), but that it is a better teaching tool, and it conveniently allows estimation of effect sizes from means and standard deviations even when the original authors do not present the corresponding *F* ratios. Although it seems that instructors are moving away from teaching ANOVA calculation in favor of focusing on the interpretation of computer output, I propose that students will learn more about the structure of ANOVA from using my new method than by avoiding hand calculation entirely (using only statistical software), or by performing calculations with the traditional method.

Keywords: analysis of variance, computational formula, hand-held calculator, teaching statistics

Requests for reprints should be sent to Barry H. Cohen, Department of Psychology, New York University, 6 Washington Place, New York, NY 10003.

In the field of psychology, all of the popular introductory (e.g., Gravetter & Wallnau, 2000; Pagano, 2001) and most of the graduate-level (e.g., Howell, 1997; Keppel, 1991) statistics texts I have seen emphasize the calculation of analysis of variance (ANOVA) by means of the raw-score or computational formula for the sum of squared deviations (*SS*) from the mean. To eliminate the need for double summation signs, the formula for the between-groups *SS* in a one-way ANOVA, for instance, is often expressed in the following fashion:

$$SS_{between} = \sum \frac{T_i^2}{n_i} - \frac{\left(\sum X\right)^2}{N}, \tag{1}$$

where T_i is the sum of the scores in the ith group, n_i is the number of scores in the ith group, ΣX is the sum of all scores, and N is the total number of scores (this notational system is not standard—many different symbols are used by various text authors). The widespread use of this formula and its companions is a holdover from the days when only the simplest four-function calculating machines were available, and such formulas were needed to reduce calculation effort and assure accuracy. Originally, these formulas were taught not because they best reveal the structure of ANOVA, but rather because students would ultimately need to use them to analyze their own data. In fact, definitional formulas were, and still are, often used to explicate the concepts of ANOVA; the definitional formula that corresponds to Formula 1 is the following:

$$SS_{between} = \sum n_i (\bar{X}_i - \bar{X}_G)^2, \tag{2}$$

where $\bar{X}_G$ equals the grand mean. However, in most statistics texts for the behavioral sciences, the definitional formulas are quickly dismissed as not suitable for calculation purposes, and the raw-score formulas are introduced along with a small data set to illustrate the calculation of a one-way ANOVA.

My argument is that the use of raw-score formulas for teaching purposes has several drawbacks, and that its chief advantages have become irrelevant due to the ready availability of cheap electronic calculators and statistical software. Given that the disadvantages now outweigh the advantages of the raw-score formula, I propose that it no longer be taught (except for historical purposes), and be replaced by an approach that I describe in the next section.

Perhaps, the largest drawback to the raw-score formula is that students cannot see why a greater spread of group means, for instance, leads to a larger value for Formula 1, whereas this is fairly easy to see from Formula 2. Another drawback to teaching ANOVA in terms of raw-score formulas is that an artificial and confusing discontinuity is created between teaching the t test and teaching ANOVA. The t-test formula is usually presented in terms of the means and variances of the two

samples. However, when the one-way ANOVA is then presented in terms of raw-score formulas, it is often not made clear that it is even possible to calculate the ANOVA if given only the means, variances, and sample sizes of the groups; this serves to obscure the connection between ANOVA and the *t* test. A change in symbols and terminology inadvertently adds to the confusion, so that there appears to be little relation between the *t* test and ANOVA (e.g., it is not often shown that the denominator of the *F* ratio in an independent-samples, one-way ANOVA is just a weighted average of the sample variances, which is just a simple extension of the pooled variance in the *t* test). The chief advantage of the raw-score approach, a reduction in computational effort and tedium, is no longer relevant, given the potential use of modern hand-held calculators in conjunction with more instructive definitional formulas, as shown next.

CALCULATING THE ONE-WAY ANOVA FROM MEANS AND STANDARD DEVIATIONS

When dealing with a one-way ANOVA in which all of the samples are the same size, it is easier to calculate the variance estimates (the mean squares) directly than to deal with the *SS* components, if you use the definitional formulas. The definitional formula for the between-groups mean square is:

$$MS_{between} = \frac{n\sum(\bar{X}_i - \bar{X}_G)^2}{k-1}, \tag{3}$$

where n is the common sample size, and k is the number of groups. Without the n, the above formula is the unbiased variance of the group means (it is unbiased because the denominator is not the number of groups, but the number of groups minus one). Thus, Formula 3 can be rewritten as n times the unbiased variance of the sample means. This is where the modern hand-held calculator makes the difference. After entering the sample means into a calculator, the standard deviation (*SD*) of these means can be obtained with the press of another key or two. In this case, it is the unbiased *SD* that we want. (On many calculators, the biased *SD* is designated by the symbol σ with n as a subscript to indicate that n is being used in the denominator of the formula. The unbiased *SD* may be indicated by either σ or s, but with $n - 1$ as a subscript, to indicate that $n - 1$ is being used as the denominator. If you are not sure which key is which on your calculator, enter the following three numbers: 3, 5, 9. The biased *SD* of these numbers is 2.494, and the unbiased *SD* is 3.055.) Squaring this value and then multiplying by n yields $MS_{between}$, the numerator of the *F* ratio. The denominator of the *F* ratio is the simple arithmetic mean of the (unbiased) sample variances, each of which is easily obtained by entering the data for that group and squaring the unbiased *SD*.

The procedure just described can be summarized in terms of the following formula for the one-way ANOVA with equal-sized groups.

$$F = \frac{n s^2_{\bar{X}}}{\frac{\Sigma s^2}{k}}, \tag{4}$$

where n is the common sample size, $s^2_{\bar{X}}$ is the unbiased variance of the sample means (i.e., Formula 3 without the n), Σs^2 is the sum of the sample variances, and k is the number of groups.

Calculators that have a key for the standard deviation are now the low end of the calculator market, and are usually called scientific calculators, and only occasionally, statistical calculators. These calculators currently cost about $10 to $15, so it is quite reasonable to require that each student purchase one (in fact, most students will already have one because they are so readily available).

It can be argued that the raw-score formulas are more efficient in that they do not require first calculating the means and variances of the samples. However, it is hard to imagine a case in which you would not want to have these descriptive statistics, anyway. It can also be argued that the raw-score formulas generally lead to greater accuracy because it is likely that too few digits beyond the decimal point will be retained when the means and variances are used as intermediate steps. This is true, but extreme accuracy is not required for pedagogical purposes, and researchers almost always use statistical software to compute statistics they intend to report. Unless it is expected that students will not have access to statistical software for data analysis, and therefore will need the increased accuracy of the raw-score formulas in their future work, there is simply no need to present these formulas—they reveal nothing comprehensible (at least to the beginning student) about the structure of ANOVA.

A more compelling argument against the teaching of ANOVA in terms of Formula 4 is that it cannot be used when the sample sizes are not equal, unless the unweighted-means method is used, and this method is almost never taught in the case of one-way ANOVA (if the unweighted-means solution is desired, n in Formula 4 is replaced by the harmonic mean of the sample sizes). According to the weighted-means solution, the usual method for a one-way ANOVA with unequal sample sizes, the numerator of the F ratio is found by dividing Formula 2 by its number of degrees of freedom. This approach requires that the grand mean be subtracted from each sample mean and the difference squared before being added. I still believe that this procedure is no more tedious than using the raw-score formula, and far more instructive (students can see that the more a sample mean varies from the grand mean, the more it contributes to the numerator of the F ratio, and

they can see that a larger n for a given sample also increases the size of the result), but the means and variances approach, coupled with the modern calculator, does not make a truly dramatic difference in the calculation of ANOVAs until you reach the complexity of a factorial ANOVA.

A NEW SYSTEM FOR CALCULATING FACTORIAL ANOVAS

Calculating a Balanced Two-Way ANOVA

Although Formula 3 provides the most instructive and least tedious (if you use a calculator to find the variance of the sample means) way to calculate the numerator for an equal-n, one-way ANOVA, this approach does not translate simply to the two-way ANOVA. Formula 3 does work for the main effects when applied to the marginal means for each factor (k becomes the number of levels for the factor and n the number of subjects included at each level of the factor), but to obtain the numerator for the interaction it is more convenient to work with the SS components, which are additive. That brings us back to a choice between Formulas 1 and 2, both of which are tedious, and neither of which capitalize on the ease with which standard deviations and variances can be found with modern calculators. The new computational system I present below extends the simplicity of Formula 4, and the power of the modern calculator, to the calculation of the two-way and higher order factorial ANOVAs.

The computational system I propose for factorial ANOVA is based chiefly on one simple formula that can be applied repeatedly with minor variations. In the case of the one-way ANOVA with equal-sized samples, this formula can be written as follows:

$$SS_{between} = N\sigma^2(\text{means}), \tag{5}$$

where N represents the total number of observations, and the term σ^2 (means) represents the biased variance of the sample means (i.e., the variance of sample means that has k, the number of means, in its denominator, rather than $k - 1$). This term could be written more compactly as $\sigma^2_{\bar{x}}$, but my use of parentheses makes it easier to modify the formula for the different components of the two-way ANOVA.

Numerators of the F *ratios.* To calculate a balanced two-way ANOVA with Formula 5, first create a matrix of cell means, and then average across and

down to find the marginal (i.e., row and column) means. Then, apply Formula 5 to the cell means as follows: $SS_{between\text{-}cells} = N\sigma^2$ (cell means). Next, apply the formula to both the row and column means: $SS_{rows} = N\sigma^2$ (row means); $SS_{columns} = N\sigma^2$ (column means). Finally, the *SS* for interaction is found by subtraction, as usual: $SS_{inter} = SS_{between\text{-}cells} - SS_{rows} - SS_{column}$. Then, SS_{inter}, SS_{rows}, and $SS_{columns}$ are divided by their respective degrees of freedom to form the numerators of the three *F* ratios.

Denominators of the F *ratios.* The denominator (or error term) of all three *F* ratios is the same, and can be found by simply averaging all of the cell variances. First, the unbiased variance is calculated separately for each cell; then, if the design is balanced (all of the cells have the same number of observations), the mean of the cell variances is found. Averaging the cell variances is clearly the most instructive way to find the ANOVA error term, but if you have the raw scores, and reducing computational effort is your goal, you can avoid having to calculate the variance of each cell by employing one additional application of the formula to find the total *SS*: $SS_{total} = N\sigma^2$ (scores), where σ^2 (scores) refers to the biased variance of all of the observations. Then $SS_{within\text{-}cells}$ can be found by subtracting $SS_{between\text{-}cells}$ (already found) from SS_{total}. Finally $SS_{within\text{-}cells}$ is divided by the appropriate degrees of freedom to form the denominator of the three *F* ratios.

Advantages of the new system. What makes this system instructive is that it can be readily seen that the size of the numerator for the main effect for rows, for instance, depends on the variance of the row means multiplied by the total *N*, and is therefore increased by a greater spread of the row means; also, a given spread of means leads to a larger numerator when multiplied by a larger number of subjects. It is easy to lose sight of how the formulas relate to your data when dealing with cell, row, and column sums as required by the usual raw-score formulas. The sums are not meaningful quantities as they depend on cell sizes, but the means of cells, rows, and columns are easy to interpret. What makes this system so much less tedious than the use of raw-score formulas, is the fact that the calculator does most of the work. In terms of calculator functions, here are the steps of my system for computing a balanced two-way ANOVA.

Calculator steps. First, enter the observations separately for each cell of the design and press the appropriate keys to find the mean and *unbiased* standard deviation for each cell (it is useful to obtain the *SD*s for descriptive purposes). Average the cell means across rows and down columns to find the marginal means. Then, enter all of the cell means, press the key for the *biased* standard deviation, square the result, and multiply by the total *N* to find $SS_{between\text{-}cells}$. Also, press the key for the

mean; this gives the grand mean, which will be useful for checking other calculations. Repeat the procedure contained in the previous two sentences for both the row and column means. The (grand) mean found after entering the row or column means should be the same as the (grand) mean found for the cell means, and serves as a check that the data have been entered correctly into the calculator. Finally, this procedure can be repeated one more time by entering all of the observations to find SS_{total} (again finding the same grand mean as a check), to find $SS_{within\text{-}cells}$ by subtraction, or $MS_{within\text{-}cells}$ can be found directly by averaging the unbiased cell variances. A complete step-by-step numerical example will be presented in a later section of this article; additional worked-out examples (including ANOVAs with repeated measures) can be found in Cohen (2000).

Although the calculator does most of the work, the student still gets the experience of working directly with his or her data. The greatest drawback of the system just described is the number of digits past the decimal point that must be retained for intermediate results to preserve accuracy. However, for classroom purposes this is not a serious concern; considerable error due to rounding off can be tolerated. Moreover, the simplicity of the system is a strong advantage. The one basic formula required is easy to memorize or look up, and the student does not have to keep track of the number of observations in each row or column (or even the number of rows and columns, until it is time to divide each *SS* component by the appropriate degrees of freedom) as is necessary with the usual raw-score formulas.

Calculating an Unbalanced Two-Way ANOVA

The chief disadvantage of the calculation system I am proposing when applied to the one-way ANOVA is its inability to accommodate the weighted-means analysis of a design with unequal groups. Fortunately, this drawback does not arise in the case of an unbalanced factorial ANOVA; the weighted-means analysis is not usually considered appropriate in that case. It is the regression approach (called Type 3 sums of squares in SPSS) that is now the most widely used analysis system for unbalanced ANOVAs (assuming the lack of balance is fairly small and random); this is a calculation process not feasible to perform with hand-held calculators. However, it can be instructive to teach students the once-popular unweighted-means analysis for unbalanced two-way ANOVAs because this procedure produces exactly the same results as the regression approach in the common 2 × 2 case. This analysis is made both less tedious and more educational by employing the computational system I am proposing. The steps described in the previous section need be modified only slightly for the unbalanced case.

First, note that the row and column means are found by taking the simple average of the cell means in a row or column, ignoring any differences in cell sizes. A

slight modification of Formula 5 is then applied to the cell, row, and column means. The adjusted formula is as follows:

$$SS_{between} = N_h\sigma^2 \text{ (means)}, \tag{6}$$

where N_h is the harmonic mean of all of the cell sizes multiplied by the number of cells. The following is a compact formula for N_h:

$$N_h = \frac{c^2}{\sum \frac{1}{n_i}}, \tag{7}$$

where c represents the number of cells, and n_i is the size of the ith cell. The denominator for all of the F ratios is a weighted average of the unbiased cell variances, just as in the weighted-means analysis of the one-way ANOVA with unequal ns (this term is called the *pooled-variance estimate* when performing a two-group t test).

CALCULATING A THREE-WAY ANOVA

The system described above for the two-way ANOVA easily generalizes to higher order factorial designs, and designs with repeated measures, and the advantages of its simplicity become more obvious (and more welcome) as the design grows more complex. The raw-score approach to the three-way ANOVA, for instance, is based mainly on repeated applications of Formula 1, whereas my approach substitutes Formula 5 in each case, leading to the following advantages:

1. My approach is applied to tables of means (collapsed across none, one, or two factors at a time). An inspection of these means gives a more accurate feeling for the size of effects that may be occurring at different levels in the design than does an inspection of tables of sums; the latter are required by the raw-score approach.
2. The n_i term in Formula 1 keeps changing between layers of the analysis (different numbers of factors being summed across), and sometimes within a layer (if factors have different numbers of levels), but N in Formula 5 remains a constant.
3. In the raw-score approach many sums have to be squared and kept track of, but in my system the calculator does nearly all of the computation, which is a considerable savings in even the simplest of three-way ANOVAs.

I suspect that very few of the instructors who cover three-way ANOVAs in their courses actually discuss the details of its computation by hand (or assign such exercises), preferring instead to emphasize the use of statistical software, and the interpretation of the results. Although my system makes it more reasonable to consider teaching the steps of these complex computations, I would not argue strongly in favor of teaching computational details for ANOVAs more complex than the two-way. However, to those instructors who see no value in teaching computations for ANOVA at any level of complexity, I wish to point out that there can occur occasions when the original data are not available, but a table of means and standard deviations is. One may come across a published journal article that contains such a table, but does not include an analysis of all of the ANOVA effects that one would like to see. I will illustrate my system with an example of such a case next. If one does not have access to statistical software that accepts means and standard deviations for a factorial ANOVA as input (I suspect that is often the case), my computational system can supply the missing F ratios easily, and these F ratios can then be used in conjunction with the sample sizes to estimate the effect sizes of the original study, perhaps to aid the planning of a future study.

AN EXAMPLE FROM THE PSYCHOLOGICAL LITERATURE

In their Study 2, Lee and Robbins (1998) published the results of a two-way ANOVA exploring the effects of group communication during the performance of a manual task (low cohesion = no communication allowed; high cohesion = some written communication), and the participant's degree of social connectedness (as determined by low or high ratings on a Social Connectedness Scale) on state anxiety, state self-esteem, and social identity. Some significant effects were reported for self-esteem and social identity, but for state anxiety it was stated that "there were no significant main or interaction effects" (p. 343). However, suppose that your main interest is state anxiety, and you are planning a similar study with a considerably larger sample (the Lee and Robbins study involved only 44 participants in the four cells combined). Even though the results for this variable were not significant in the Lee and Robbins study, you may want to see just how close to significance the results came, and use the F ratios and cell sizes to estimate the effect sizes that may apply to your replication. Fortunately for my purposes, Lee and Robbins provided a table of means and standard deviations for all three dependent variables (cell sizes were included as well because their factorial design was not balanced). Although the published table had arranged the four cells in one row for each dependent variable, I have se-

lected only the state anxiety variable, and rearranged the data to display the marginal means in Table 1. Next, I use my new approach to compute a two-way ANOVA on these data.

Because this design is not balanced, Formula 6 should be used instead of Formula 5. The N_h that must be used in Formula 6 is four (the number of cells) times the harmonic mean of the four cell sizes. Formula 7 does this for us in one step:

$$N_h = \frac{c^2}{\sum \frac{1}{n_i}} = \frac{4^2}{\frac{1}{9}+\frac{1}{8}+\frac{1}{13}+\frac{1}{14}}$$
$$= \frac{16}{.111+.125+.077+.0714} = \frac{16}{.3844} = 41.6.$$

The error term for the three F ratios is the weighted average of the four cell variances. The SD entries in Table 1 are unbiased, so they just have to be squared to produce unbiased variances; then, each is multiplied by $n_i - 1$ before being added. Finally, the total, which is SS_W, is divided by df_W. If the design were balanced, we would be taking a simple average of the variances, instead.

$$MS_W = \frac{(9-1)13.82^2+(8-1)9.43^2+(13-1)9.42^2+(14-1)9.62^2}{40}$$
$$= \frac{8(191)+7(88.9)+12(88.7)+13(92.5)}{40} = \frac{4417.2}{40} = 110.43.$$

TABLE 1
Level of State Anxiety As a Function of the Imposed Level of Group Cohesion and the Reported Level of Social Connectedness

	Cohesion		
Connectedness	*Low*	*High*	*Row Ms*
Low			
M	37.13	39.22	38.175
SD	13.82	9.43	
n	9	8	
High			
M	39.31	32.71	36.01
SD	9.42	9.62	
n	13	14	
Column *Ms*	38.22	35.965	37.0925

Note. From "The Relationship Between Social Connectedness and Anxiety, Self-esteem, and Social Identity," by R. M. Lee and S. B. Robbins, 1998, *Journal of Counseling Psychology, 45*, p. 343. Copyright 1998 by the American Psychological Association. Adapted with permission of the authors.

Next, I apply Formula 6 to the cell, row, and column means:

$$\begin{aligned} SS_{between\text{-}cells} &= N_h\sigma^2 \text{ (cell means)} = N_h\sigma^2(37.13,39.22,39.31,32.71) \\ &= 41.6\ (7.163) = 298.0 \\ SS_{rows} &= N_h\sigma^2 \text{ (row means)} = N_h\sigma^2(38.175,36.01) \\ &= 41.6\ (1.172) = 48.76 \\ SS_{columns} &= N_h\sigma^2 \text{ (column means)} = N_h\sigma^2(38.22,35.965) \\ &= 41.6\ (1.271) = 52.87. \end{aligned}$$

The *SS* for the interaction of the two factors can now be found by subtraction: $SS_{inter} = 298.0 - 48.76 - 52.87 = 196.37$. Because the numerator *df*s are all equal to 1, each of the *SS*s just calculated can be divided by the error term to obtain the *F* ratios: $F_{connectedness} = 48.76/110.43 = .44$; $F_{cohesion} = 52.87/110.43 = .48$; $F_{interaction} = 196.37/110.43 = 1.78$.

If you did not have a calculator with the *SD* function, you could still use the previously mentioned formulas, but the variances would have to be calculated the long way. For example, using the definitional formula for the biased variance, σ^2 (37.13, 39.22, 39.31, 32.71) would be calculated the following way: $[(37.13 - 37.0925)^2 + (39.22 - 37.0925)^2 + (39.31 - 37.0925)^2 + (32.71 - 37.0925)^2\ ]/4 = (.0014 + 4.5263 + 4.9173 + 19.2063)/4 = 28.6513/4 = 7.163$. You can see how much work the calculator saves even for the smallest (i.e., 2 × 2) factorial ANOVA.

There is no indication that either of the main effects calculated above would ever become significant in a larger study, but it can also be seen that if the cell sizes were made three times larger without changing the apparent effect size for the interaction, this effect could easily attain significance. Of course, in the 2 × 2 case, the interaction can be readily calculated as a linear contrast of cell means, but my system easily generalizes to any number of levels and factors.

CONCLUSIONS

Now that computers and statistical programs are so readily available for calculating complex ANOVAs and are so much faster and more accurate than people using raw-score formulas and hand-held calculators, many statistics instructors no longer require their students to calculate factorial ANOVAs by hand. Other instructors feel that some calculations are important for solidifying statistical concepts, but are finding it increasingly difficult to justify the time and tedium involved with the hand calculation of two-, and especially three-way ANOVAs. Modern hand-held calculators, in conjunction with the use of Formula 5 or 6, now offer a compromise between the extreme tedium of applying raw-score formulas to a factorial ANOVA, and the

totally "hands free" approach of using statistical software. Calculating factorial ANOVAs with the system I have proposed gives students a feeling for why a greater spread of means increases the size of the *F* ratio, as does a greater sample size.

Recently, Lovett and Greenhouse (2000) laid out a series of principles that affect how easily students learn difficult, abstract material, such as found in the usual statistics course. Several of these principles help to explain the advantages of my proposed calculation system. Specifically, students can be expected to learn ANOVA more quickly and retain this learning longer because they are practicing calculations by themselves (Principle 1), their understanding of ANOVA will be deeper because it will be anchored by and integrated with what they learned about the *t* test (Principle 4), and their learning will be more efficient because the new formula imposes less of a mental load as it is used (Principle 5). Nonetheless, I do not strongly recommend my approach for an introductory undergraduate course. I teach students at the master's and doctoral level, and virtually all of these students have already been introduced to factorial ANOVA (almost always in the context of the raw-score method of ANOVA), and can better appreciate the new approach without confusion. It is potentially confusing to start using the biased variance after students have learned to use the unbiased variance in the context of statistical inference. However, if students do not use SS_{total} to avoid calculating SS_{within}, they can simply remember to apply the biased variance only to group or cell means, and never to individual scores when inference is the goal.

Finally, students with this training will know how to analyze ANOVA designs from published results and be able to find *F* ratios that were not presented. This can even be done with designs involving repeated measures, if the error term you need for an *F* ratio not presented is shared by an *F* ratio that is. I require my students to find tables of means and *SD*s in the psychological literature (especially tables not completely analyzed in the article they come from), and then calculate two- or even three-way ANOVAs based on those tables. This helps my students see the utility of the new approach, and they seem to appreciate the simplicity of and the reduced effort required by the new formula. Some instructors will still feel that teaching the hand calculation of ANOVA, even with the system proposed here, is not a good use of class time, but there is no longer much justification for teaching raw-score procedures that are wasteful of time and effort, while providing little educational value.

REFERENCES

Cohen, B. H. (2000). *Explaining psychological statistics* (2nd ed.). New York: Wiley.

Gravetter, F. J., & Wallnau, L. B. (2000). *Statistics for the behavioral sciences* (5th ed.). Belmont, CA: Wadsworth.

Howell, D. C. (1997). *Statistical methods for psychology* (4th ed.). Belmont, CA: Wadsworth.

Keppel, G. (1991). *Design and analysis: A researcher's handbook* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.

Lee, R. M., & Robbins, S. B. (1998). The relationship between social connectedness and anxiety, self-esteem, and social identity. *Journal of Counseling Psychology, 45,* 338–345.

Lovett, M. C., & Greenhouse, J. B. (2000). Applying cognitive theory to statistics instruction. *The American Statistician, 54,* 196–206.

Pagano, R. R. (2001). *Understanding statistics in the behavioral sciences* (6th ed.). Belmont, CA: Wadsworth.