A User's Guide to RECIPE: A FORTRAN Program for Determining One-Sided Tolerance Limits for Mixed Models With Two Components of Variance Version 1.0

Mark G. Vangel
National Institute of Standards and Technology
Statistical Engineering Division, 101/A-337
Gaithersburg, MD 20899-0001

July 15, 1994

1 Introduction

This document provides guidance to users of the computer program RECIPE (REGression Confidence Intervals on PErcentiles). This program can provide approximate one-sided tolerance limits (or, equivalently, confidence intervals on percentiles) for a wide range of situations where one is able to assume a normal probability model. Arbitrary regression models with or without a random effect can be analyzed using this program, and in this ability to accommodate between-batch variability RECIPE is perhaps unique.

RECIPE is a general program for one-sided mixed model tolerance limits for any mixed model having one or two components of variance, with no interaction between fixed and random effects. However, this work was motivated by the need for statistical methodology for use in determining design allowables for composite materials in aircraft applications, particularly in the presence of between-batch variability. Most readers of this document will be users of composite materials, so we will explain this program using the examples of A-basis and B-basis material properties (or design allowables). An A-basis value is a $(.99, .95)$ lower tolerance limit, and a B-basis value is a $(.90, .95)$ lower tolerance limit. Alternatively, A- and B- basis values can be interpreted to be 95% lower confidence bounds on the 1st and 10th population percentiles, respectively. For more information on statistically-based design allowables and their relation to tolerance limits, see Mil-Handbook-17D (1994, Volume 1, Chapter 8) and Vangel (1996). This user's manual is intended to be usable by engineers with little statistical training; consequently (as in Vangel 1996) the examples are discussed with these engineers in mind, and some background material is included which can be omitted by the experienced statistician.

The theory underlying the method used is documented in Vangel (1995a, 1994) and will not be discussed here. Instead, we will illustrate the use of this computer program through a series of representative examples.

2 One-Sided Tolerance Limits and Confidence Intervals on Percentiles

The present section is divided into two parts: the first part provides a precise definition of one-sided tolerance limits and their relationship to confidence intervals on percentiles, the second part attempts to explain tolerance limits in terms of the engineering application to design allowables. Obviously, the statistician may want to read the first part carefully and skim the remainder, while the opposite will likely be true for the engineer.

2.1 A Mathematical Definition

Let U be a random variable, and assume that we are interested in interval estimates of quantiles of U. Let T be a statistic based on a random sample from U. A (β, γ) lower tolerance limit is a random variable T such that a proportion of at least β of the population of U is covered by the interval

(T, ∞) with probability γ. That is,

$$\Pr\left[\Pr\left(U > T | T\right) \geq \beta\right] = \gamma. \tag{1}$$

One-sided upper tolerance limits are defined similarly (see the AZppendix). We refer to β as the *content* and γ as the *confidence* .

Now, assume in addition that U denotes a $\mathrm{N}(\mu, \sigma^2)$ random variable, where $\sigma^2 = \sigma_b^2 + \sigma_e^2$ is the sum of a between-group component of variance σ_b^2 and a within-group component of variance σ_e^2. For example, U might represent the strength of a random specimen chosen from a random batch of a material, where $\mu = w^T\theta$ may depend on covariates. The program RECIPE calculates approximate (β, γ) lower tolerance limits for U.

2.2 Statistically-Based Design Allowables

A *design allowable* for a material is the maximum value of stress or strain at which one can be reasonably certain that failure will not occur. For the design of structures for which weight is not a primary consideration, allowables are typically calculated by dividing a stress level at which failure is known to often occur by a sufficiently large constant (a *safety factor*) (Gere and Timoshenko, 1984, p. 29). The structure is then designed so as to ensure that the stresses (or strains, etc.) do not exceed the allowables for the materials. No use is made of statistics in this general definition of allowable. The strength of a material is regarded as a known constant, and this value is divided by a safety factor, which may reflect extensive engineering experience in similar applications.

Material properties often exhibit considerable scatter, however, and this is particularly apparent for many composite materials. Also, the use of deterministic safety factors can result in structures which are heavier than they need be, an obvious drawback to their use in aerospace applications. Consequently, the concept of design allowable has a precise statistical definition in the aircraft industry, established long ago in Mil-HDBK-5 (1987) for metals, and carried over to the corresponding handbook for composites, Mil-HDBK-17 (Volume 1, 1994). (What we refer to as 'allowables' here are called 'material basis properties' in Mil-HDBK-17; for the present discussion these terms can be regarded as equivalent.)

A B-basis design allowable, or material basis value, is defined to be a 95% lower confidence limit on the 10th percentile of the population of a material property (usually strength, but sometimes ultimate strain). An A-basis value is defined as a 95% lower confidence limit on the 1st percentile of a population; this more stringent value is typically used in situations where failure of a component would cause structural failure.

It is helpful to begin by considering the simplest case. Based on strength measurements for n identical specimens of a material, tested under identical experimental conditions, a B-basis is calculated. This calculation can be done using RECIPE (see Example 1), though this scenario is so simple that the calculations could actually be done by hand (e.g., Mil-HDBK-17, 1994, Volume 1, Section 8.5). We could plot a histogram of these values, and we might imagine what this histogram would look like if we had many (even infinitely many) strength values. The histogram would likely approach a smooth curve, which we call the *population* of strength values. This population has a 10th percentile, the value of which we will never know, since we only have n specimens, and n is presumably quite small. Our B-basis value tells us something about this percentile, however. A B-basis value has the property that if we were to obtain n specimens over and over again, and calculate many of these basis values, 95% of the time these (hypothetical) B-basis values would be less than the unknown 10th population percentile. The B-basis value thus provides a conservative estimate of the 10th percentile; we say that it is less than this percentile *with 95% confidence*.

Complications of at least three kinds can be introduced into these calculations. First, we don't know the functional form of this 'smooth curve' which we think of as being the population of material property values. We will make the assumption that this curve is Gaussian (or 'bell-shaped'); statisticians would say that it has a 'normal distribution'. Other models, such as the Weibull distribution, are sometimes used, but at present only the normal distribution can be applied to the complicated 'messy data' scenarios, involving several batches, which are typical in applications.

Another common difficulty occurs when the population mean varies with factors such as temperature, layup, and humidity, forcing the experimenter to spread his testing resources rather thinly by testing only a very few specimens at any one set of experimental conditions. Allowable curves (or surfaces) may be required, for example with temperature as an independent variable. The computations are likely to be too difficult for hand or calculator calculations, though RECIPE can provide the desired results fairly easily (Examples 2-5).

The third complication arises when we have data from several batches, and we are concerned with variability among these batches, and among future batches which we will see during production. Our population here may consist of the strength of a random specimen chosen from a random batch, with the population variance being the sum of between-batch and within-batch components. If we ignore the batches and pool the data, then we will be acting as if the between-batch component of this variance is zero, and hence underestimating the population variance. As a result, our allowables could be too high.

The second and third difficulty discussed above often occur together when, for example, data from several batches are obtained at several temperatures. RECIPE can be used for these problems, as we will see in Example 4.

3 Regression Analysis

Testing is expensive, so it is not surprising that in industry one usually obtains only a small amount of data for a fixed set of experimental conditions (e.g. five room temperature dry unidirectional tensile strength measurements on data from a single batch). If one chooses to determine material basis properties for given conditions using only data obtained at these conditions, then one will often be faced with prohibitively low values because of the limited data. However, if one is willing to assume that other sets of data on the same material come from populations which differ in their means, but have the same variance, then *regression analysis* is an extremely powerful general statistical technique which uses *all* of the data to determine material basis values at *each* condition. In addition, if some of the conditions are continuous variables such as temperature, one can interpolate or (with caution) extrapolate to estimate basis values for conditions for which *no* test data is available. We relate the various datasets corresponding to different conditions by a *regression model* and make assumptions of independence, constant variance, and normality. In return for these assumptions, we are able to make much more efficient use of the data than if we analyzed each condition separately; however, if the assumptions do not approximately hold, regression methods can lead to misleading results. So the use of any regression analysis program, in particular RECIPE, must be accompanied by careful inspection of the data to check the validity of assumptions. RECIPE is not a replacement for general regression analysis software: it is a program which only calculates one-sided tolerance limits (in particular, material basis values), and the use of RECIPE should be supplemented by data analysis using general-purpose statistical software. A discussion of the practice of regression analysis, including the interpretation and criticism of the results and the diagnosis and treatment of the violation of assumptions is beyond the scope of this user's guide. There are many textbooks on this topic at all levels (e.g., Box, et. al. (1978), Chapter 14; Weisberg (1980)).

4 Regression Models

The objective of a regression analysis for material basis properties is to obtain basis values for a particular response (e.g., tensile strength) as functions of fixed factors (e.g., temperature, layup, and humidity). We will refer to the measured response values as *observations*, and the values which describe the conditions corresponding to these observations as *covariates*. For example, if we assume a linear relationship between tensile strength and temperature, then the mean strength at a temperature T_i is, in the limit of infinitely many observations at this temperature, equal to $\theta_0 + \theta_1 T_i$. The constants θ_0 and θ_1 are generally unknown, and must be estimated from the data. The values that these constants multiply, here 1 and T_i, are covariates; together they describe the fixed conditions under which the ith strength observation was made.

Assume that the data being analyzed consist of n observations at l fixed conditions (or *levels*), and number these conditions $1, 2, \ldots, l$. In our example of linear regression on temperature, we have l temperatures, and l corresponding sets of covariates: (1, T_1), (1, T_2),..., (1, T_l). We need to indicate which fixed condition corresponds to each observation, so let the fixed condition for the sth observation be $p(s)$. We will also allow for the fact that each observation is made on a specimen from one of m batches. These batches are numbered $1, 2, \ldots, m$, and $q(s)$ indicates the batch corresponding to the sth observation. We will denote the observations by y_s, for $s = 1, 2, \ldots, n$, where the sth value comes from fixed level $p(s)$ and from batch $q(s)$.

We assume that the $\{y_s\}$ are a sample from a normal distribution with mean

$$\mu_{p(s)} = \theta_1 z_{p(s),1} + \theta_2 z_{p(s),2} + \ldots + \theta_r z_{p(s),r}, \tag{2}$$

where the $\{z_{p(s),u}\}$,for $1 \leq p(s) \leq l$ and $u = 1, \ldots, r$, are known constants and the $\{\theta_u\}$ are parameters to be estimated. For example, if we are assuming that mean strength varies linearly with temperature, and if condition $p(s) = 1$ corresponds to 75 degrees, then

$$\mu_1 = \theta_1 + \theta_2 75, \tag{3}$$

so $r = 2$, $z_{11} = 1$, and $z_{12} = 75$.

We can never observe the means $\mu_{p(s)}$. Each data value consists of the sum of $\mu_{p(s)}$ plus a random quantity $b_{q(s)} + e_s$, where $b_{q(s)}$ takes on a different value for each batch $q(s)$, and e_s takes on a different value for each observation. We assume that the $\{b_{q(s)}\}$ and $\{e_s\}$ are random samples from normal populations with means zero and variances σ_b^2 and σ_e^2, respectively. We will refer to σ_b^2 as the *between-batch variance*, and to σ_e^2 as the *within-batch* (or *error*) variance. We can now express the data as

$$y_s = \mu_{p(s)} + b_{q(s)} + e_s = \theta_1 z_{p(s),1} + \ldots + \theta_r z_{p(s),r} + b_{q(s)} + e_s, \tag{4}$$

where the $\{z_{p(s),u}\}$ are known, the $\{\theta_u\}$ are unknown fixed quantities, and the $\{b_{q(s)}\}$ and $\{e_s\}$ are random quantities with unknown variances. The specification (4) is called a *regression model.* Every regression analysis begins with the choice of a regression model. In the remainder of this section, we illustrate the construction of regression models with five examples. In the following section we provide analyses for particular cases of each of these examples, using actual graphite/epoxy strength data.

4.1 Example 1: Simple Random Sample

We begin with the simplest case of all: a simple random sample of n observations from a single batch at a fixed set of conditions. For this case, we have $l = 1$ condition and $m = 1$ batch, so $p(s) = q(s) = 1$ for each s. We write this model as

$$y_s = \theta_1 + e_s. \tag{5}$$

Note that $b_{q(s)}$ does not appear in (5). We cannot estimate between-batch variability with fewer than two batches, just as we cannot estimate a variance with fewer than two observations.

4.2 Example 2: Random Effects ANOVA Model

Now assume that we have data on several batches, each tested under the same set of fixed conditions. Since we have only one set of fixed conditions, the model for this example has a constant mean, but now we have both between-batch and within-batch components of variance. So $l = 1$, and

$$y_s = \theta_1 + b_{q(s)} + e_s. \tag{6}$$

Equation (6) is the usual random-effects ANOVA (or simply 'ANOVA') model of Mil-HDBK-17D (Volume 1, Section 8.5.4).

4.3 Example 3: Simple Linear Regression With Data From a Single Batch

We return now to the situation where we have data from a single batch, so that $m = 1$; but now we allow for several conditions, so that $l > 1$. To fix ideas, assume that we have several sets of unidirectional tensile strength data from a single batch, with each set being tested at a different temperature, and with all other conditions held constant. Assume further that the strength for this material is believed to vary linearly with temperature, at least for temperatures within the range of the data. As in (5), we cannot estimate between-batch variability. The regression model appropriate for this situation is:

$$y_s = \theta_1 z_{p(s),1} + \theta_2 z_{p(s),2} + e_s, \tag{7}$$

where $z_{p(s),1} = 1$ and $z_{p(s),2} = t_i$, the ith test temperature.

4.4 Example 4: Simple Linear Regression With a Random Effect

If we have the same situation as in Section 4.3, except that we have data from more than one batch, then we can introduce the $b_{q(s)}$ random batch effect in the model, to get

$$y_s = \theta_1 z_{p(s),1} + \theta_2 z_{p(s),2} + b_{q(s)} + e_s. \tag{8}$$

4.5 Example 5: One-Way Mixed Model ANOVA: Basis Values With Data From Multiple Sources

Suppose that we have several batches of data from each of several manufacturers, and that they wish to combine their resources to determine basis values. If we are absolutely certain that the manufacturing and testing are identical for all of the data, then we can ignore the fact that the data came from multiple sources. Often, however, there will be slight differences among the manufacturers in the way the material was fabricated and/or tested. In such cases, if we are not willing to assume that the variability between and within batches are close to being the same for all manufacturers, then there is no alternative to applying the usual ANOVA method (as in Section 4.2) separately to each manufacturer's data. But if we are willing to assume that each set of data exhibits the same variability (with a possibly different mean for each manufacturer), then *all* of the batches can be used to determine a basis value for *each* manufacturer. These basis values will often be substantially higher, and closer together, than if each manufacturer had acted alone.

To develop a regression model for this example, let the mean for the ith manufacturer be μ_i. If there are l manufacturers, we have $r = l$ unknown fixed parameters $\mu_1, \mu_2, \ldots, \mu_l$ – in addition to the components of variance σ_b^2 and σ_e^2. Hence, the regression model is of the form

$$\begin{aligned} y_s &= \theta_1 z_{p(s),1} + \ldots + \theta_r z_{p(s),r} + b_{q(s)} + e_s \\ &= \mu_{p(s)} + b_{q(s)} + e_s. \end{aligned} \tag{9}$$

We have taken the zs to be $z_{p(s),u} = \delta_{p(s),u}$, where $\delta_{p(s),u}$ (the Kronecker-δ) equals one when $p(s) = u$, and zero otherwise. The fixed parameters are $\theta_i = \mu_i$.

5 Examples

To illustrate the use of the program, we will now present an analysis for each of the five examples of the previous section, using actual graphite/epoxy tensile strength data. The models (5) and (6) can also be analyzed using methods described in Mil-HDBK-17D (Volume 1, Sections 8.5.5 and 8.5.4, respectively). The present approach reduces exactly to the Mil-HDBK-17D methods for these cases. The linear regression model without batch effects (7) is an example of the regression model discussed in the Handbook (Volume 1, Section 8.5.8). The one-way mixed model (9) is also discussed in Mil-HDBK-17D (Volume 1, Section 8.5.9), and RECIPE again essentially agrees with the Handbook procedure for this special case. The regression model with random effects (8) cannot be handled using the statistical methods presently in Mil-HDBK-17D, though it is easily treated using RECIPE.

In order to use RECIPE for a particular problem, it is necessary to create a file which contains the data, the information necessary for the program to construct the regression model, and a list of the covariate values at which the basis value is to be evaluated. The files for the five examples discussed in this documentation are included with the software, and are named 'ex1.dat', 'ex2.dat', ..., 'ex5.dat'. The format of these files will become clear as we discuss the examples.

5.1 Analysis for Example 1

For the simple random sample (5), the example dataset has observations on five specimens from a single batch. The file 'ex1.dat' is

```
#
#   RECIPE Example #1: Simple random sample
#
#   -- For this example, we have 5 observations: all at the same
#      fixed level and from the same batch.  RECIPE is a very
#      general program which is here used for a very simple
#      example. This example might seem confusing because it
#      is so special.  If so, consider the more complicated
#      examples, particularly Example #4. Ironically, the
#      simpler examples may then be easier to understand.
#
#   --  ntot, nlvl, nbch, npar, npts, prob, conf
#
     5 1 1 1 1 .9d0 .95d0
#
#  --  Fixed levels. Here nlvl=1 and npar=1; that is there is only
#      one fixed level and one regression parameter (a constant mean),
```

```
#      so this part of the input consists of one row and one column,
#      containing just the number '1'.
#
   1
#
#  --  Fixed level, batch number, response value.  Note that there
#      is only one level (nlvl=1) and one batch (nbch=1).
#
   1    1  328.1174
   1    1  334.7674
   1    1  347.7833
# (this just shows that comments can be put anywhere: even among
#  the data values. This is useful, for example, if a data value
#  is to be removed from the analysis. Simply put a '#' at the
#  beginning of the appropriate line, and decrease 'ntot' by 1
#  in the first noncomment line)
   1    1  346.2661
   1    1  338.7314
#
#  --  Points at which to evaluate tolerance limit. Here the only fixed
#      effect is a constant mean, so this part of the input is trivial.
  1
```

Lines which begin with a '#' are *comment lines* which are ignored by the program. Comment lines can be inserted anywhere, and are intended to make RECIPE data files self-documenting. The input to this program is free-format, so it doesn't matter which column values are in, so long as they are in the correct order and separated by spaces. The sole exception to this is that comment lines must have a '#' in column 1.

The first non-comment line of any RECIPE file has seven constants, to which we give the mnemonics 'ntot', 'nlvl', 'nbch', 'npar', 'npts', 'prob', and 'conf'. The total number of observations (n) is 'ntot', the number of fixed levels (l) is 'nlvl', the number of batches (m) is 'nbch', and the number of fixed parameters (r) is 'npar'. It is necessary to specify the number of points at which the basis values are to be determined. For example, if a linear regression model relates strength to temperature, then a basis value can be calculated at any number of temperatures, i.e. the temperatures at which basis values are determined need not correspond to values for which data is available. So the fifth number on this line, 'npts', specifies the number of basis values which are to be calculated. The sixth and seventh values, 'prob' and 'conf', give the *content* and *confidence* which are to be used. For purposes of allowable calculations, one need only remember that 'prob' should be .99d0 or .90d0, for A- and B-basis values, respectively, and that 'conf' should be .95d0.

In this example, we see that there are $n = 5$ observations, at $l = 1$ fixed level, from $m = 1$ batch, with $r = 1$ fixed parameters, and that a single B-basis value is to be calculated. (Since this corresponds to a simple random sample, there is only one basis value which it makes sense to calculate.)

The next $l = 1$ noncomment lines specify the fixed levels; for this example there is only one fixed level, and it is just the mean, so this part of the file has only one line with a '1' in it.

The following $n = 5$ uncommented lines each gives, from left to right, a fixed level ($p(s)$, here $p(s) = 1$), batch ($q(s)$, here $q(s) = 1$), and observation (strength y_s), for $s = 1, \ldots, 5$.

The next npts=1 uncommented line gives the zs corresponding to each point at which a basis value is to be calculated. Again, because this example is a simple random sample, this part of the file consists of only a single line with a '1'.

We run RECIPE as follows:

```
recipe
 Filename (without .dat extension) ?
ex1

  RECIPE : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex1.crt not found.
     Satterthwaite approximation will be used.

  regini : Warning: between-batch variance cannot
           be estimated from these data. Results
           will be based on the assumption that the
           between-batch variability is negligible.
```

```
        Probability    Confidence   Regression  Tolerance Limit
           0.90            0.95     339.133120     311.338667
```

The first two columns of the output indicate that a B-basis value has been calculated. The third column gives the value of a point on the least squares regression line (here just the sample mean), and the fourth column gives the corresponding basis value (here the usual normal B-basis value for a single sample of five specimens). The warning message reminds us that we cannot estimate between-batch variability with data from a single batch, and that consequently this basis value has been calculated under the assumption that there is no between-batch variability.

There are two methods which RECIPE can use to calculate allowables. One involves the use of a Satterthwaite approximation (Satterthwaite, 1946), and the other requires using an auxilliary program SIMPVT in order to obtain a quantile of a pivotal random variable for which the probability distribution cannot be determined in analytical form. Usually, these two methods will give very nearly the same answers, at least for material basis value calculations. The simpler Satterthwaite approximation is therefore recommended for general use. For more information, see Vangel (1995a) and the Appendix.

5.2 Analysis for Example 2

For the one-way ANOVA model (6), the example data file 'ex2.dat' is:

```
#
#  RECIPE Example #2: Basis value from a one-way ANOVA model
#
#  -- This example has 31 observations in 6 batches, for which
#     an ANOVA B-basis value is to be determined
#
#   --  ntot, nlvl, nbch, npar, npts, prob, conf
   31 1 6 1 1 .9d0 .95d0
#
#  -- Fixed levels. Here we are fitting a one-way ANOVA model, so there
#     is only one fixed level, and only one fixed parameter (the mean)
```

```
#     to estimate.
   1
#
#  -- Fixed level number, batch number, strength. Since we have
#     only one fixed level, the first column is all ones. The
#     second column gives the batch number, and the third column
#     gives the strength values.
    1   1  328.1174
    1   1  334.7674
    1   1  347.7833
    1   1  346.2661
    1   1  338.7314
    1   2  297.0387
    1   2  293.4595
    1   2  308.0419
    1   2  326.4864
    1   2  318.1297
    1   2  309.0487
    1   3  337.0930
    1   3  317.7319
    1   3  321.4292
    1   3  317.2652
    1   3  291.8881
    1   4  297.6943
    1   4  327.3973
    1   4  303.8629
    1   4  313.0984
    1   4  323.2769
    1   5  312.9743
    1   5  324.5192
    1   5  334.5965
    1   5  314.9458
    1   5  322.7194
    1   6  291.1215
    1   6  309.7852
    1   6  304.8499
    1   6  288.0184
```

```
    1   6  294.1995
#
#   -- Points at which to evaluate tolerance limit. For the one-way
#      ANOVA model used here, there is only one point at which the
#      evaluation can be done: it corresponds to the one fixed
#      level of the model.
    1
```

The output is similar in form to the example discussed above:

```
recipe
  Filename (without .dat extension) ?
ex2

  RECIPE : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex2.crt not found.
     Satterthwaite approximation will be used.

        Probability    Confidence   Regression  Tolerance Limit
           0.90           0.95     316.010884     271.672860
```

Note, however, that the warning message that was output for Example 1 doesn't appear here. We are able to estimate the between-batch variability since we have six batches, and the fourth column gives the one-way random effects ANOVA basis value, instead of the single sample basis value of Example 1.

5.3 Analysis for Example 3

For an example of a regression model with data from a single batch, we have data on tensile strength obtained at -67 and 75 degrees Fahrenheit. The file 'ex3.dat' is:

```
#
#  RECIPE Example #3: Regression model with data from a single batch
```

```
#
#  -- This dataset has 11 observations at two fixed levels. The
#     data come from 1 batch, there are two fixed parameters to
#     estimate (the slope and intercept of a straight line), and
#     a B-basis value is to be calculated at 7 points on this line.
#
#   --  ntot, nlvl, nbch, npar, npts, prob, conf
   11 2 1 2 7 .9d0 .95d0
#
#  -- We are fitting a model y=a+bT at two levels: T=75 degrees and
#     T=-67 degrees.  The first column corresponds to 'a' in this
#     linear equation; the second column corresponds to 'b'. Note
#     that these values need not be given in any special order,
#     for example (1, -67) need not come before (1, 75).  The
#     important thing is that the order of the rows given here
#     must correspond to the level indicator, p(s), given with each
#     response value.
   1 75
   1 -67
#
#  -- Now we have the 11 observations.  The first column is the
#     level (=1 for 75 degrees, =2 for -67 degrees), the second
#     column is the batch (always 1), and in the third column are
#     the strength observations.
#
   1    1   328.1174
   1    1   334.7674
   1    1   347.7833
   1    1   346.2661
   1    1   338.7314
   1    1   340.8146
   2    1   343.5855
   2    1   334.1746
   2    1   348.6610
   2    1   356.3232
   2    1   344.1524
#
```

```
#  -- Finally, we give the seven points at which basis
#     values are to be determined.  These correspond
#     to seven different temperatures -67,...,50. Note
#     that the first column of ones is required because
#     of the intercept in the regression model
  1 -67
  1 -50
  1 -25
  1  0
  1 25
  1 50
  1 75
```

Note that the first noncomment line of 'ex3.dat' indicates (in order, from left to right) that we have 11 observations in all, that the data are at 2 fixed levels, that all of the data are from a single batch, that the fixed part of the model involves 2 unknown parameters (actually, it turns out that we are fitting a straight line), that we will evaluate the basis value curve at 7 points, and that the tolerance limits to be calculated are B-basis values.

This example illustrates the common situation where a material basis value is required as a function of temperature. We have data at two fixed levels, corresponding to the temperatures -67 and 75 degrees, and we would like to determine basis values at the 7 temperatures -67, -50, -25, 0, 25, 50 and 75 degrees. The intercept of the linear function is constant for all temperatures, so the first column equals 1 for the 2 rows which give the levels of the fixed effect, as well as the 7 rows which give the points at which the basis values are to be evaluated. The output from running RECIPE on these data is

```
recipe
  Filename (without .dat extension) ?
ex3

  RECIPE : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex3.crt not found.
     Satterthwaite approximation will be used.
```

```
regini : Warning: between-batch variance cannot
         be estimated from these data. Results
         will be based on the assumption that the
         between-batch variability is negligible.

      Probability    Confidence   Regression  Tolerance Limit
         0.90           0.95      345.379340     325.887099
         0.90           0.95      344.665104     325.747683
         0.90           0.95      343.614756     325.338699
         0.90           0.95      342.564409     324.619436
         0.90           0.95      341.514062     323.538853
         0.90           0.95      340.463714     322.102027
         0.90           0.95      339.413367     320.366619
```

Each of the last seven lines gives a point on the regression line, and the corresponding point on the B-basis curve, for each of the seven sets of independent variables (seven temperatures) in the file 'ex3.dat'. Note that there is a warning message to remind us that we cannot estimate between batch variability using data from a single batch. The basis values calculated are valid under the assumption that the between-batch variability is zero (or at least negligible).

5.4 Analysis for Example 4

For the fourth example, we have data at the same two temperatures, but now with several batches at each temperature. The file 'ex4.dat' is

```
#
#     RECIPE Example #4: Regression model with data from several
#                        batches
#
#  -- In this example, we have 72 strength observations on data
#     from 8 batches.  A straight-line regression is fit with
#     two fixed levels (temperatures).  B-basis values are calculated
#     for 7 points along this curve.
#
```

```
#  -- ntot, nlvl, nbch, npar, npts, prob, conf
 72 2 8 2 7 .9d0 .95d0
#
#  -- There are two fixed levels, corresponding to
#     75 and -67 degrees.
 1  75
 1 -67
#
#  -- The following 72  rows give the fixed level in the
#     first column, the batch in the second column, and the
#     strength observation in the third column.
   1   1  328.1174
   1   1  334.7674
   1   1  347.7833
   1   1  346.2661
   1   1  338.7314
   1   2  297.0387
   1   2  293.4595
   1   2  308.0419
   1   2  326.4864
   1   2  318.1297
   1   2  309.0487
   1   3  337.0930
   1   3  317.7319
   1   3  321.4292
   1   3  317.2652
   1   3  291.8881
   1   4  297.6943
   1   4  327.3973
   1   4  303.8629
   1   4  313.0984
   1   4  323.2769
   1   5  312.9743
   1   5  324.5192
   1   5  334.5965
   1   5  314.9458
   1   5  322.7194
```

1	6	291.1215
1	6	309.7852
1	6	304.8499
1	6	288.0184
1	6	294.1995
2	1	340.8146
2	1	343.5855
2	1	334.1746
2	1	348.6610
2	1	356.3232
2	1	344.1524
2	2	308.6256
2	2	315.1819
2	2	317.6867
2	2	313.9832
2	2	309.3132
2	2	275.1758
2	3	321.4128
2	3	316.4652
2	3	331.3724
2	3	304.8643
2	3	309.6249
2	3	347.8449
2	4	331.5487
2	4	316.5891
2	4	303.7171
2	4	320.3625
2	4	315.2963
2	4	322.8280
2	5	340.0990
2	5	348.9354
2	5	331.2500
2	5	330.0000
2	5	340.9836
2	5	329.4393
2	7	330.9309
2	7	328.4553

```
   2   7  344.1026
   2   7  343.3584
   2   7  344.4717
   2   7  351.2776
   2   8  331.0259
   2   8  322.4052
   2   8  327.6699
   2   8  296.8215
   2   8  338.1995
#
#    -- The following 7 rows give the points at which
#       the B-basis value is to be calculated: these
#       correspond to 7 temperatures -67,-50,...,75.
  1 -67
  1 -50
  1 -25
  1  0
  1 25
  1 50
  1 75
```

A run of RECIPE produces the output:

```
recipe
  Filename (without .dat extension) ?
ex4

  RECIPE : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex4.crt not found.
     Satterthwaite approximation will be used.

        Probability    Confidence   Regression  Tolerance Limit
           0.90           0.95      327.537310     286.895095
           0.90           0.95      326.157386     285.580736
           0.90           0.95      324.128085     283.557672
           0.90           0.95      322.098785     281.470595
```

0.90	0.95	320.069485	279.335972
0.90	0.95	318.040184	277.119935
0.90	0.95	316.010884	274.783636

The input and output files have the same form as for Example 3. The important distinction between Example 3 and Example 4 is that the basis values in Example 4 account for between-batch variability, while in Example 3 we calculated basis values strictly valid for only a specific batch. Note also that the warning message that appeared in Example 3 does not show up here, since we have data from several batches.

5.5 Analysis for Example 5

In this example, we have data on several batches of the same material from each of two manufacturers. We assume that the variability is the same for each manufacturer, so that model (9) applies, with $l = r = 2$. However, there was an important difference in processing, so that one would expect the means to be different for each manufacturer. The data file 'ex5.dat' is

```
#
#  RECIPE Example #5: Basis values using data from multiple sources
#
#    -- In this example, we have five batches of data: three from
#       one source, and two from a second source.  We would like
#       to use all five batches of data to get a tolerance limit
#       for each source.
#
#    --  ntot, nlvl, nbch, npar, npts, prob, conf
#
  15 2 5 2 2 .9d0 .95d0
#
#    -- The fixed part of this model is a different mean for
#       each of the two sources
   1 0
   0 1
#
#    -- Here are the 15 data values.  Column 1 indicates the
```

```
#       fixed level (data source), and column 2 indicates the
#       number of the batch. The third column gives the strength
#       values.
  1  1  75.8
  1  1  78.4
  1  1  82.0
  1  2  68.8
  1  2  70.9
  1  2  73.5
  1  3  74.5
  1  3  74.8
  1  3  78.8
  2  4  81.3
  2  4  87.7
  2  4  89.0
  2  5  88.2
  2  5  91.2
  2  5  94.2
#
#    -- The tolerance limit are to be calculated at two
#        points, which correspond to the two sources. So
#        we just repeat the two lines for the fixed part
#        of the model here.
  1 0
  0 1
```

The file 'ex5.dat' tell us that there are 15 data values, and that we are using a regression model with $r = 2$. The first column of the 15 rows of 'ex5.dat' which contain data indicates the fixed level, the second column for these rows indicates the batches, and the third column gives the strength values. The fixed part of the model has two means, one for each data source. So the rows which give the fixed levels, and the rows which give the points at which basis values are to be evaluated, have a 1 in one column and a 0 in the other. Contrast this with Examples 1 and 2, where there is only 1 fixed level, and so the corresponding rows have just 1 column having a single value, 1.

The RECIPE output for this example is:

```
recipe
  Filename (without .dat extension) ?
ex5

  RECIPE : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex5.crt not found.
     Satterthwaite approximation will be used.

        Probability    Confidence   Regression  Tolerance Limit
           0.90           0.95       75.277778      59.401536
           0.90           0.95       88.600000      71.902179
```

The B-basis values are therefore 59.40 and 79.90 for the two manufacturers. As a simple exercise in using this program, one can show (following Example 2, using the data from Example 5) that the if each manufacturer had used their own data alone, then the basis values would be 52.8 and 34.6, respectively. Note that the mixed model (9) gives basis values which are higher and closer together. In particular, the very low value 34.6 is due to the fact that the second manufacturer only has data on two batches.

6 Batches, Panels, and Confounding

RECIPE is based on the assumption of at most two sources of variability; we have called these 'between-batch variability' and 'within-batch variability'. In the manufacturing of composites, however, there are typically at least three sources of variability. For composites made from prepreg, the additional source is due to the fact that several specimens are typically manufactured together as a 'panel', consequently we will refer to this third source as 'between-panel' variability.

When we have data on a material from several batches, but at only one set of fixed conditions (e.g., Example 2) we cannot estimate batch and panel variabilities separately. Whenever we obtain data from a new panel, that data also comes from a different batch. (In statistical terminology, we say the batch and panel variances are *confounded.*) So, what we call 'between-batch variability' in such cases is actually the sum of the between-batch and between-panel variances. Unless the between-panel variability is negligible, we will over-estimate the between-batch variance in such cases. This can result in material basis properties that are lower than they should be.

Next, consider the situation where data are available from several batches, at more than one set of fixed conditions (e.g., Example 4). If we assume also that data at different conditions from the same batch are from different panels, then we are able, in principle, to estimate the between-batch and between-panel variances separately. However, since we are not able to include both of these sources of variability in our regression models, the between-panel variance is confounded, not with the between-batch variance as above, but with the *within-batch* variance. This can result in material basis values that are somewhat *higher* than they should be, but this is likely to be a less serious problem than the case where panel and batch variances are confounded, for several reasons. Perhaps the most important of these is that of the sources of variability, that due to batches is our primary concern, and this is now being treated appropriately. Another reason is that there is typically considerable variability within panels, and if the between-panel variance is small with respect to this third source of variability, then the material basis properties will not be substantially higher than they should be.

7 Conclusion

In this document, we have illustrated the use of a computer program, RECIPE, for determining one-sided normal tolerance limits for mixed models having two variance components. Since our primary audience is users of composite materials, the focus has been on the application of this program to determining material basis properties, and ultimately design allowables, for composite materials. To some extent, these notes are self-contained; there is a brief discussion of regression models, and some discussion of the concept of confounding, for example. More background information is provided in the tutorial article Vangel (1996). However, the routine user of this program should acquire at least an elementary knowledge of statistics at the level of Box, et. al (1978), and Weisberg (1980), or else consult periodically with someone knowledgeable in this field. Also, there has been no discussion here of the theoretical foundation of the algorithm implemented in this program. Material on this topic can be found in technical articles (Vangel 1995a, 1994).

8 Appendix: Advanced Topics

In this section, some special topics are discussed which may be of interest, particularly to users of this program who are interested in applications other than material design allowables. Issues discussed here include checking the actual confidence level by simulation, an improvement over the Satterthwaite approximation for highly unbalanced datasets, upper tolerance limits, and two-sided confidence limits on quantiles. This section assumes a higher level of statistical expertise than most of the rest of this manual.

8.1 SIMCOV: Examining the Actual Confidence Level

If the between-batch variance is zero and one knows this to be the case, then tolerance limits provided by RECIPE will be exact. However, when the possibility of between-batch variability is allowed for, the actual confidence level will depend on the ratio of between to within batch variances σ_b^2/σ_e^2, or, equivalently, on the *intraclass correlation* $\rho = \sigma_b^2/(\sigma_b^2 + \sigma_e^2)$. The intraclass correlation is the correlation between observations from the same batch. It is more convenient to use ρ than the variance ratio, because it assumes values in the finite interval [0, 1].

The *nuisance parameter* ρ is unknown, and there are often too few batches to be able even to estimate it very well. We would like to have a tolerance limit procedure for which the actual confidence level equals the nominal level, whatever ρ might be. This goal is probably unattainable in general, although one can come extremely close for certain very simple regression models (see Vangel 1992). This difficulty is analogous to the well-known Behrens-Fisher problem concerning the two-sample test for equality of means in the presence of variances in unknown ratio. However, RECIPE provides tolerance limits for which the confidence levels usually do not depend strongly on ρ, and for which the actual confidence is generally fairly close to the nominal level.

In order to determine how close the actual confidence level corresponding to the RECIPE algorithm is to the nominal level, it is necessary to simulate. This is because the actual confidence level depends on the model matrix and on the points on the regression surface at which the tolerance limits are calculated, which will be different for different applications. The program SIMCOV is provided to simulate the actual confidence. It takes as input the same file which is used by RECIPE, and it provides confidence levels for various levels of ρ for each point at which tolerance limits are to be calculated.

One can expect SIMCOV to show that the RECIPE intervals are somewhat conservative when ρ is near zero, somewhat anticonservative for intermediate values of ρ, and nearly exact for $\rho = 1$. For highly unbalanced datasets, the confidence may differ substantially from the nominal level when $\rho = 1$ (for an example, see Vangel 1995b). This indicates that the Satterthwaite approximation is not adequate, and that improved performance can be obtained by replacing the Satterthwaite value with the appropriate quantile of a simulated *pivotal* random variable. By doing this, one can attain *exactly* the nominal confidence level when $\rho = 1$ (to within the accuracy of the simulated pivotal quantile), and this will typically improve performance for intermediate values of ρ as well.

As an example, if SIMCOV is applied to the input file for Example 5, something resembling the following output will result:

```
  Filename (without '.dat' extension) ?
ex5

  SIMCOV : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex5.crt not found.
     Satterthwaite approximation will be used.

  Number of simulation replicates ?
5000
  Integer seed ?
32
  Number of values for intraclass correlation ?
11
 1= use same random numbers for each rho
 0= use different random numbers for each rho ?
0


        rho       confidence
```

0.0000	0.9646
0.0000	0.9686
0.1000	0.9550
0.1000	0.9578
0.2000	0.9494
0.2000	0.9470
0.3000	0.9482
0.3000	0.9504
0.4000	0.9444
0.4000	0.9364
0.5000	0.9384
0.5000	0.9432
0.6000	0.9378
0.6000	0.9350
0.7000	0.9310
0.7000	0.9368
0.8000	0.9344
0.8000	0.9322
0.9000	0.9422
0.9000	0.9448
1.0000	0.9516
1.0000	0.9522

For $\rho = 0, .1, \ldots, 1$, the actual confidence was obtained from 5000 simulated regressions. The two values given for each ρ correspond to the two points at which the tolerance limit is to be calculated. This example data file is for a (.90, .95) lower tolerance limit, and it is clear that the nominal confidence of .95 is nearly attained for all ρ. Note that if N is the number of simulations and the actual and nominal confidences are indeed equal, then one would expect the simulation results to usually fall within the two-standard-deviation interval

$$\gamma \pm 2\sqrt{\frac{\gamma(1-\gamma)}{N}}, \tag{10}$$

where γ is the nominal confidence. For this example, $\gamma = .95$ and the interval (10) is $(.944, .956)$.

8.2 SIMPVT: An Improvement on Satterthwaite's Approximation for Highly Unbalanced Data

Usually, SIMCOV will demonstrate that RECIPE will provide confidence levels reasonably close to the nominal level. However, for unbalanced models we can improve on the Satterthwaite tolerance limits if we are willing to do more work. If RECIPE finds a file with a '.crt' extension, then it will read the critical values from that file, rather than using a Satterthwaite approximation. The program SIMPVT simulates the pivotal random variable for $\rho = 1$ and creates a '.crt' file for use by SIMCOV and by RECIPE. An example will help illustrate the use of SIMPVT.

An unbalanced dataset was created from Example 5 by deleting four values: two from batch 1 and one each from batches 4 and 5. The new input file, called 'ex5a.dat' follows:

```
#
#  RECIPE Example #5a: Basis values using data from multiple sources
# This is an 'unbalanced version' of Example #5 in which four
# values have been deleted: two from batch 1, and one each from
# batches 4 and 5. Note that 'ntot' has been changed from 15 to 11.
#
#   -- In this example, we have five batches of data: three from
#      one source, and two from a second source.  We would like
#      to use all five batches of data to get a tolerance limit
#      for each source.
#
#   --  ntot, nlvl, nbch, npar, npts, prob, conf
#
  11 2 5 2 2 .9d0 .95d0
#
#   -- The fixed part of this model is a different mean for
#      each of the two sources
   1 0
   0 1
#
#   -- Here are the 15 data values.  Column 1 indicates the
#      fixed level (data source), and column 2 indicates the
```

```
#       number of the batch. The third column gives the strength
#       values.
  1  1  75.8
# 1  1  78.4
# 1  1  82.0
  1  2  68.8
  1  2  70.9
  1  2  73.5
  1  3  74.5
  1  3  74.8
  1  3  78.8
  2  4  81.3
  2  4  87.7
# 2  4  89.0
  2  5  88.2
  2  5  91.2
# 2  5  94.2
#
#     -- The tolerance limit are to be calculated at two
#        points, which correspond to the two sources. So
#        we just repeat the two lines for the fixed part
#        of the model here.
  1 0
  0 1
```

The actual confidence that RECIPE will achieve for this dataset with the Satterthwaite approximation for $\rho = 0, .5, 1$ are determined by SIMCOV:

```
simcov
  Filename (without '.dat' extension) ?
ex5a

  SIMCOV : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical value file ex5a.crt not found.
     Satterthwaite approximation will be used.
```

```
  Number of simulation replicates ?
25000
  Integer seed ?
12
  Number of values for intraclass correlation ?
3
 1= use same random numbers for each rho
 0= use different random numbers for each rho ?
0


        rho        confidence

      0.0000       0.9681
      0.0000       0.9668
      0.5000       0.9421
      0.5000       0.9430
      1.0000       0.9574
      1.0000       0.9566
```

When $\rho = 1$ the actual confidence, although probably acceptably close to the nominal .95, is well outside the two-standard-deviation limit of $(.947, .953)$.

SIMPVT is now used to produce a critical value file 'ex5a.crt' which can be used instead of the Satterthwiate approximation. Since the SIMCOV has shown that the actual confidence level at $\rho = 1$ which we want to improve on is already close to .95, we must determine the pivotal quantile quite accurately in order to see any improvement in the confidence level. So we will have SIMPVT obtain the desired quantiles from 1,000,000 simulated values of the pivotal random variable.

```
simpvt
  Filename (without '.dat' extension) ?
ex5a
  Number of simulation replicates ?
1000000
  Integer seed ?
23
```

```
  SIMPVT : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0,  April 1995)

      Simulated critical values

  Number of values     =            2
  Number of replicates =      1000000
  Seed                 =           23
  Input file           = ex5a.dat
  Output file          = ex5a.crt

    5.293335957447922
    5.525013667521424
```

The two numbers printed out by SIMPVT are the critical values corresponding to the two points at which tolerance limits are to be calculated; they have been written to the new file 'ex5a.crt'.

Now we run SIMCOV again to see how much improvement we've realized. It helps to use the same seed as in the previous run of SIMCOV in order to make it easier to discern any improvement.

```
simcov
  Filename (without '.dat' extension) ?
ex5a

  SIMCOV : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical values from file ex5a.crt will be used.

  Number of simulation replicates ?
25000
  Integer seed ?
12
  Number of values for intraclass correlation ?
3
 1= use same random numbers for each rho
 0= use different random numbers for each rho ?
```

```
0


        rho       confidence

      0.0000      0.9678
      0.0000      0.9664
      0.5000      0.9405
      0.5000      0.9411
      1.0000      0.9499
      1.0000      0.9497
```

Note that SIMCOV uses the '.crt' file this time, and that the confidence when $\rho = 1$ is very nearly exactly the nominal level.

How much of a difference will this make in the actual $(.90, .95)$ lower tolerance limits? To see this, RECIPE was run with the '.crt' file, the last three characters in this file name were changed, and RECIPE was run again. The results are

```
recipe
  Filename (without .dat extension) ?
ex5a

  RECIPE : One-Sided Random-Effect Regression Tolerance Limits
  (Version 1.0, April 1995)

  *** Simulated pivot critical values from file ex5a.crt will be used.

        Probability    Confidence   Regression  Tolerance Limit
           0.90           0.95      73.871429      60.316295
           0.90           0.95      87.100000      72.952700
```

and

```
recipe
  Filename (without .dat extension) ?
ex5a
```

```
RECIPE : One-Sided Random-Effect Regression Tolerance Limits
(Version 1.0, April 1995)

 *** Simulated pivot critical value file ex5a.crt not found.
   Satterthwaite approximation will be used.

      Probability    Confidence   Regression  Tolerance Limit
         0.90           0.95      73.871429      59.714960
         0.90           0.95      87.100000      72.458949
```

respectively. For more theoretical information on this topic, see Vangel (1995a).

8.3 Upper Tolerance Limits

Because the application which motivated the development of this methodology is to lower tolerance limits, and because it was important to make the software as easy-to-use for non-statisticians as possible, the RECIPE software must be modified slightly in order to calculate upper tolerance limits.

An *upper tolerance limit*, as calculated by RECIPE, is a statistic of the form

$$L = w^T \hat{\theta} - KS, \tag{11}$$

where $\hat{\theta}$ is a vector of estimated regression coefficients, w is a vector of constants which determines the point on the regression surface at which a tolerance limit is to be determined, S is the residual standard deviation, and K is a statistic which depends on the estimated variance ratio. It is not hard to show that if (11) provides a (β, γ) *lower* tolerance limit, then

$$U = w^T \hat{\theta} + KS \tag{12}$$

will be a $(1-\beta, 1-\gamma)$ *upper* tolerance limit. Therefore, to calculate a (β, γ) upper tolerance limits, provide $1-\beta$ and $1-\gamma$ as the content and confidence in the input file, and modify the source code of RECIPE by changing the sign of the tolerance limit factor. Specifically, subroutine 'regdat' concludes with the lines

```
        t (i) = xm(i) -tfct*sqrt(rmsa)
 10   continue
```

```
      return
      end
```

The sign on 'tfct' should be changed to give

```
         t (i) = xm(i) +tfct*sqrt(rmsa)
 10   continue
      return
      end
```

Of course, it is trivial to modify the program to allow the use to specify either upper or lower tolerance limits in the input file, and this will probably be a feature in the next version of the software.

8.4 Confidence Limits on Quantiles

It is obvious from the definition of a one-sided tolerance limit that a (β, γ) lower tolerance limit is a 100γ percent lower confidence limit on the $100(1-\beta)$ percentile, and that a (β, γ) upper tolerance limit is a 100γ percent upper confidence limit on the 100β percentile of the population. It is easy to show that two one-sided tolerance limits can be constructed to provide any desired *two-sided* confidence limits on any population quantile.

To be precise, let B_1 and B_2 be lower tolerance limits with confidences $(1+\gamma)/2$ and $(1-\gamma)/2$, respectively. Then, since for tolerance limits of the form calculated by RECIPE B_1 is always less than B_2, the random interval $[B_1, B_2]$ provides a 100γ percent two-sided confidence interval on the $100(1-\beta)$th population percentile. Hence, one-sided tolerance limits can provide both one- and two-sided confidence intervals on quantiles.

References

[1] Box, G. E. P. , Hunter, W. G., and Hunter, J. S. (1978), *Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building*, John Wiley ans Sons, New York.

[2] Gere, J. M. and Timoshenko, S. P. (1984), *Mechanics of Materials*, Boston: Prindle, Weber & Schmidt.

[3] Mil Handbook 5E (1987), *Metallic Components for Aircraft Structures*, Philadelphia: Naval Publications and Forms Center.

[4] Mil Handbook 17D (1994), *Polymer Matrix Composites, Volume I: Guidelines*, Naval Publications and Forms Center, Philadelphia.

[5] Satterthwaite, F. E. (1946), "An Approximate Distribution of Estimates of Variance Components," *Biometrics Bulletin*, 2, 110-114.

[6] Vangel, M. G. (1992),"New Methods for One-Sided Tolerance Limits for a One-Way Balanced Random-Effects ANOVA Model," *Technometrics*,34, 176.

[7] Vangel, M. G.(1994), "ANOVA Estimates of Variance Component for 'Partially-Balanced' Mixed Models", submitted for publication.

[8] Vangel, M. G. (1995a), "One-Sided β-Content Tolerance Limits for Mixed Models With Two Components of Variance' , submitted for publication.

[9] Vangel, M. G. (1995b), "One-Sided β-Content Tolerance Limits for Mixed Models", *Proceedings of the Section on Physical and Engineering Sciences*, American Statistical Association, 200-206.

[10] Vangel, M. G. (1996), "Design Allowables From Regression Models Unsing Data From Multiple Batches", *Proceedings of the 12th ASTM Symposium on Composite Materials Testing and Design*, to appear.

[11] Weisberg, S. (1980), *Applied Regression Analysis*, Second Edition, John Wiley and Sons, New York.