Charles Annis [1]

Probabilistic Life Prediction *Isn't* as Easy as It Looks

Reference: Annis, C., **"Probabilistic Life Prediction *Isn't* as Easy as It Looks,"** *Probabilistic Aspects of Life Prediction, ASTM STP-1450,* W. S. Johnson and B.M. Hillberry, Eds., ASTM International, West Conshohocken, PA, 2003.

Abstract: Many engineers affect "Probabilistic life prediction" by replacing constants with probability distributions and carefully modeling the physical relationships among the parameters. Surprisingly, the *statistical* relationships among the "constants" are often given short shrift, if not ignored altogether. Few recognize that while this simple substitution of distributions for constants will indeed produce a non-deterministic result, the corresponding "probabilities" are often woefully inaccurate. In fact, even the "trend" can be wrong, so these results can't even be used for sensitivity studies. This paper explores the familiar Paris equation relating crack growth rate and applied stress intensity to illustrate many statistical realities that are often ignored by otherwise careful engineers.

Keywords: crack growth, Paris equation, probability, statistics, Monte Carlo, simulation, non-deterministic, probabilistic, joint, conditional, marginal, multivariate

Introduction:

There is more to Monte Carlo simulation than replacing constants with probability densities. The purpose of this study is to demonstrate this by comparing the observed distribution of lives of 68 nominally identical crack growth specimens with Monte Carlo (MC) simulations of lives based on the distributions of their Paris law parameters. It will be shown that several common MC sampling techniques produce wildly inaccurate results, one with a standard deviation that is *7X* larger than was exhibited by the specimen lives themselves. The cause of such aberrant behavior is explained. It is further observed that estimates of the Paris law parameters are jointly distributed as bivariate normal, and

[1] Principal, Charles Annis, P.E., Statistical Engineering, Palm Beach Gardens, FL 33418-7161. Charles.Annis@StatisticalEngineering.com

a MC simulation using this joint density reproduces the specimen mean and standard deviation to within a few percent.

The lessons here apply to *any* regression model, not just to these data, nor only to crack growth rate models.

The Data:

In the mid 1970s Dennis Virkler, then a Ph.D. student of Professor Ben Hillberry at Purdue, conducted 68 crack growth tests of 2024-T3 aluminum [1, 2]. These tests were unusual for several reasons. They were conducted expressly to observe random behavior in fatigue. While almost all crack growth tests measure crack length after some number of cycles, Virkler measured cycle count at 164 specific crack lengths. This provided a direct measure of variability in cycles, rather than the usually observed variability in crack length at arbitrary cyclic intervals. While two of the specimens appear to stand out from their brethren, the purpose of this investigation is not to play Monday Morning Quarterback 25 years after the game, and there is no reason not to consider all 68 specimens here. In any event their exclusion changes only the numeric details. The fundamental results are not affected, nor are they affected by using a normal, rather than lognormal density to describe them.

It is common practice to fit a single da/dN *vs.* ΔK curve through multiple specimens of the same material tested under the same conditions of temperature, stress ratio, and frequency. In the study reported here, however, 68 individual Paris models were used. Fitting a single curve describes the mean trend behavior very well, but it obscures random specimen-to-specimen differences. Since real applications are subjected to similar randomness, it is necessary to capture that effect as well.

Fatigue Lives Are Lognormal:

It has been long recognized that fatigue lives are satisfactorily modeled using the lognormal density. For these 68 specimens that model is less than optimal and there is some evidence that the probability density may be a mixture of two densities. It is not the purpose of this paper to repeat the earlier work by Virkler, Hillberry and Goel [2], and as it turns out, the actual form of the distribution of the specimen lives themselves only influences the numeric details of this study, since each specimen's crack growth rate curve was treated individually. (Treating the data as normal, however, results in a bias in the simulated mean of about 5%. The bias using the lognormal is negligible.)

Conventional Monte Carlo Simulation:

Unlike many engineering analytical results, probability estimates are difficult to verify experimentally. This unfortunate reality has perpetuated the misuse of a valid statistical tool, and the consequences may not be apparent for years to come.

Most engineering Monte Carlo simulations are performed this way.

1. Set up a conventional deterministic analysis;
2. Replace constants with probability distributions;
3. Sample once from each distribution;
4. Compute the deterministic result and store the answer;
5. Repeat steps 3 and 4 many times;
6. Compute the mean and standard deviation of the collected results.

Sadly, many engineers are unfamiliar with the implicit *statistical* assumptions that are at the foundation of Monte Carlo simulation, but as been observed elsewhere [3] "Simply not understanding the nature of the assumptions being made does not mean that they do not exist."

What possibly could be wrong with this paradigm? Luckily we (the engineering community) have a dataset that is nearly perfect for answering this question, viz. the data collected by Virkler and Hillberry, as part of Virkler's Ph.D. dissertation. Professor Hillberry graciously made these available for further study.

Monte Carlo Modeling Specifics:

After fitting individual Paris equations to each of the 68 specimens, the mean and standard deviation for the individual Paris parameters, intercept, *C*, and slope, *n*, were computed. The well-known Paris model for fatigue crack growth is given in equation 1.

$$da/dN = 10^{C}(\Delta K)^{n} \tag{1}$$

where *da/dN* is the crack growth rate, in mm per cycle, and ΔK is the applied stress intensity factor, in MPa$\sqrt{}$m, given by equation 2.

$$\Delta K = \Delta\sigma\sqrt{\pi a}\, f(a \mid geometry) \tag{2}$$

Here, $\Delta\sigma$ is the testing stress range, $\sigma_{\max} - \sigma_{\min}$, *a* is the crack length, and *f( )* is a function of the specimen (or component) geometry and crack length. Of course, when equation 1 is plotted on a *log-log* grid this is a straight line with intercept *C* and slope *n*.

Assuming for the sake of simplicity that there was no variation in the starting crack size, the final crack size, or the test stress, the calculated cyclic lifetime can be computed from the individual Paris fits using equation 3.

$$\begin{aligned} da/dN &= 10^{C}\left[\Delta\sigma\sqrt{\pi a}\, f(a \mid geometry)\right]^{n} \\ dN &= da \Big/ \left\{10^{C}\left[\Delta\sigma\sqrt{\pi a}\, f(a \mid geometry)\right]^{n}\right\} \\ N &= \int_{a_0}^{a_{final}} 10^{-C}\left[\Delta\sigma\sqrt{\pi a}\, f(a \mid geometry)\right]^{-n} da \end{aligned} \tag{3}$$

In practice this integration is usually carried out numerically.

To conduct the usual MC simulation N_i is computed from $h(C_i, n_i)$ where *h( )* is equation 3, and *i* ranges from 1 to say 1000 (or 10 000).

Many MC practitioners then calculate a mean and standard deviation for N, or $log_{10}(N)$, report the results and stop there, since there is nothing against which to compare the distribution of computed values for N_i. Virkler's data show the observed distribution of actual specimen lives and thus provide a direct comparison for these calculations.

The Paris Law is Adequate:

Before going further it is prudent to check the goodness-of-fit of the Paris equation itself. If the underlying model for crack growth rate is inadequate there is little hope for accurate life prediction based on it. The sigmoidal shape of the *da/dN ΔK* data (Figure 1) suggests a model such as the SINH [4] might do a better job than the straight line Paris model (and it does, increasing the ratio of standard deviations of calculated lives, 0.918 for Paris, to 0.957 for the SINH by reducing the disagreement between calculated and observed specimen lives from 8.2% to 4.3%). The added model complexity, however, obscures the real issue here, namely the abysmal performance of a rather common Monte Carlo simulation (700% error in predicted scatter). Since the Paris law is adequate it is used here for simplicity.

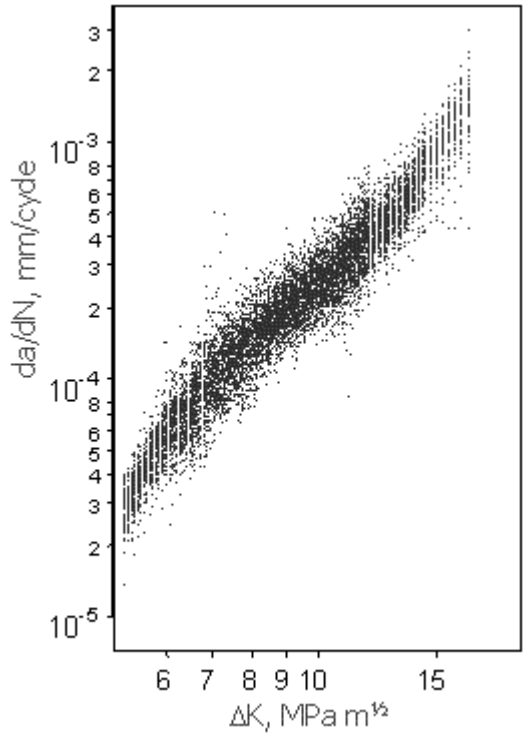

Figure 1 - *da/dN vs. DK are S-shaped.*

A Note on Modeling:

Statisticians often assess the efficacy of a mathematical model by decomposing the sums-of-squares of differences between the model and the observations. We, however, are less interested in the differences between the measured crack growth rates, da/dN_i, and their Paris model, than we are in their integrated collective behavior, as given by

equation 3. Such an integrated metric summarizes all sources of "error" - material variability, lack-of-fit, testing uncertainties – into the difference between the observed specimen life, and that provided by equation 3. We thus have traded the potential for better arithmetic diagnostics (scrutiny of the Paris model) for a more direct measure of what we are really interested in - life prediction performance.

How well does the conventional Monte Carlo algorithm perform?

The conventional MC simulation of 1000 samples, with independent model parameters, *C* and *n*, did an acceptable job predicting the *mean* lifetime, after the *log* transform. Because the data are skewed to the right, as all fatigue data are, the untransformed simulated results overestimate means of the symmetrical normal models slightly.

The simulated *standard deviations* were another matter: The actual observed standard deviation for 68 specimens is 0.03015 log_{10} units (18 447 cycles)[2]. The conventional MC simulation of 1000 samples, with independent model parameters, *C* and *n*, produced a standard deviation of 0.19778 log_{10} units (140 261 cycles), *6.6X too large!*

A closer look shows the situation gets even worse. To be fair, the best possible Paris model would use the 68 individual Paris fits, since no simulation could be expected to be better than the actual specimens' behavior. Using the 68 Paris equations in equation 3 produces a standard deviation of 0.02769 log_{10} units (16 332 cycles), which is *smaller* than the observed standard deviation by about 8%. Why?

Of the 68 specimens, two seemed to exhibit longer lives than what might have been inferred by from the behavior of the other 66. All 68 specimens were used here. Since the actual specimen life doesn't directly influence its *da/dN vs. ΔK* behavior, predicted lives based on these two Paris fits would be more like their sister specimens, resulting in the smaller standard deviation for the integrated Paris equations. So to provide a fair comparison with simulated Paris models, the behavior of the 68 integrated Paris laws should be the baseline. Thus the baseline scatter is 0.02769 log_{10} units.

Comparing the simulation's standard deviation of 0.19778 log_{10} units with the integrated Paris law baseline shows the simulation to have overestimated the scatter by 0.19778 / 0.02769 or about *7.1X*. This is awful. Such a simulation would be worse than useless since it would likely compel a costly redesign. Put in perspective, the probability of failure before about 207 000 cycles is 0.1%, determined from the mean and standard deviation of the 68 specimens' (*log*-transformed) lives. The MC simulation puts this failure rate at about 33%, an overestimation of failure rate of over *300X*.

This absurd simulation result has been observed by every engineer who has performed similar MC simulations, since it doesn't require any statistics to detect an answer that is wrong by a factor approaching an order of magnitude in standard deviation. Sadly the most common palliatives proposed as remedies do not perform much better.

---

[2] The analyses were carried out using $log_{10}(cycles)$, and again using untransformed cycles. The reported log_{10} result can not, of course, be determined simply by taking the *log* of the mean and standard deviation of the untransformed results. All calculations are summarized in Tables 1 and 2 and Figure 5.

What Went Wrong?

The model parameters, *C* and *n*, are assumed to be normally distributed. Is this a good assumption in this case?

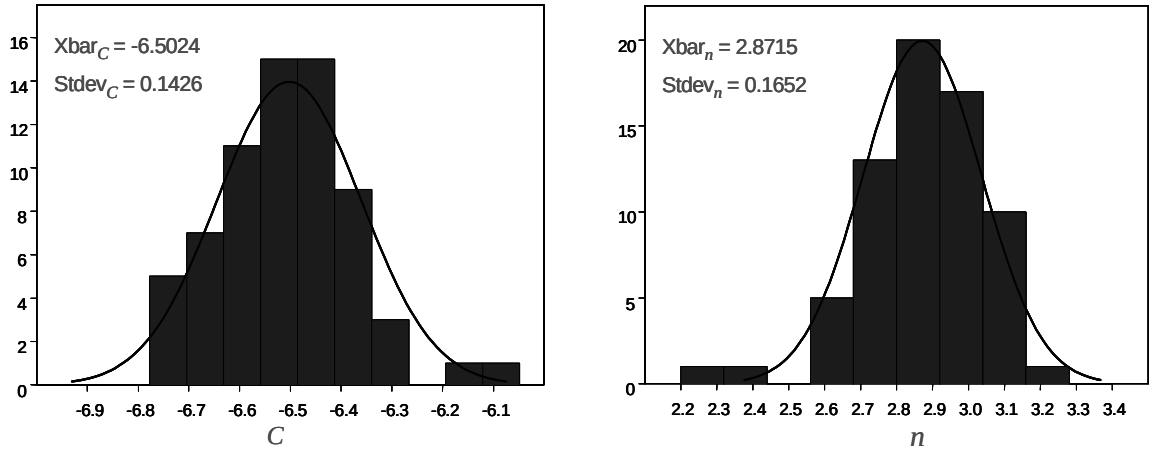

Figure 2 - *Histograms of Paris Model Parameters C and n*

Figure 2 presents histograms of both model parameters. While somewhat approximate, the normal density is not an altogether improper model; surely these departures from the normal could not have caused the *7X* inflation of the standard deviation. A closer look at the figures provides a clue. There are two observations that are *high* for parameter *C*, and two that are *low* for parameter *n*. Perhaps these should be considered as pairs, rather than as independent observations. Figure 3, a schematic plot of crack growth rate *vs.* stress intensity on a *log-log* grid, shows why *C* and *n* behave in tandem: when the slope, *n*, is shallow the intercept, *C*, must be larger for the resulting line to go through the data. Similarly, a steeper slope requires a smaller intercept.

Possible Remedies (all of them wrong):

Assuming *C* and *n* to be independent, when they obviously are not (the most common error in Monte Carlo modeling), results in unacceptable error in simulated lifetime scatter. Possible remedies that have been suggested are:

1. *n* assumed fixed, *C* is normal
2. *C* assumed fixed, *n* is normal
3. *C* assumed a linear function of *n*.

Fixing either *n* or *C* seems at first blush like a reasonable solution, and it does reduce the over-prediction of scatter from *7.1X* to *5.1X* (*n* fixed) or *5.4X* (*C* fixed). While this is an obvious improvement, the error remains wildly unacceptable. Sadly, it is at this stage when the standard deviation of *C* or *n* is arbitrarily "adjusted," *i.e.* fudged until a believable result is achieved.

Figure 4 also shows why assuming either *C* or *n* as fixed is not reasonable. The horizontal line is at $n = 2.87$, the average of 68 Paris slopes. This is a reasonable value only when $-6.58 < C < -6.45$. When *C* is outside this range, as it will be often, the resulting simulated combination is very, *very* improbable. In fact observations in either the first or third quadrants (large *n* with large *C*, or small *n* with small *C*) are exceedingly unlikely in reality but occur about half the time in uncorrelated simulation.

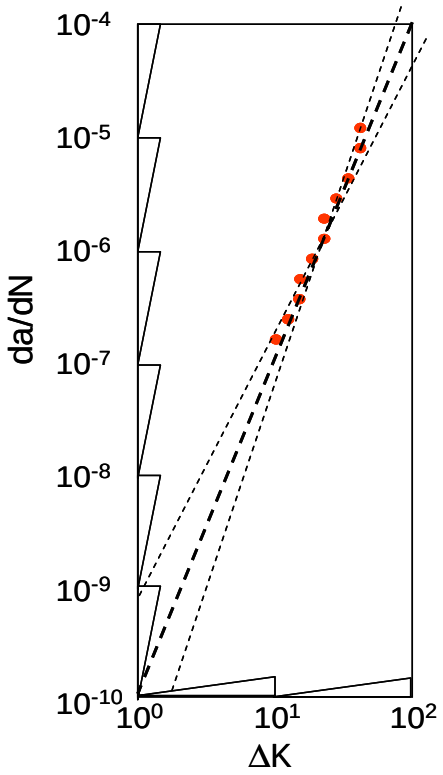

Figure 3 - *Schematic showing why Paris Parameters must be correlated. Note that in this schematic the intercept is C = log_{10}(da/dN)= -10, at log_{10}(DK)=0.*

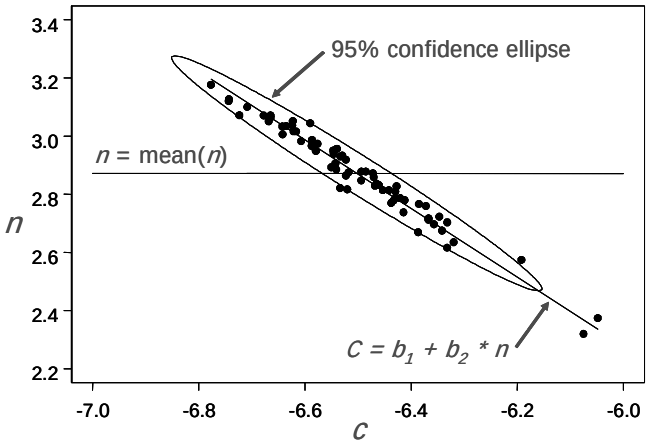

Figure 4- *Paris Parameters C and n are obviously correlated (r=0.982).*

Another option for remedy suggests itself since the two parameters are obviously so closely related: let one be a function of the other. A linear fit of $C=b_1 + b_2 n$, with n being sampled from a normal density, does indeed improve things. But this time the resulting error ratio is 0.51, *i.e.:* the scatter has been *over*-corrected, and now is *underestimated* by almost half. Clearly this nonconservative result is also unacceptable.

To understand why such an appealing suggestion should have such an undesirable result, look again at Figure 4 which also shows the 95% confidence ellipse for the C and n pairs. Assuming that one is a linear function of the other, in effect collapses this ellipse into a line, thus *underestimating* the overall variability. (The confidence ellipse also suggests that two of the tests may be different from the others, as was noted earlier.)

The Right Way:

We have considered four very common oversights in Monte Carlo modeling. So, how do you do it correctly?

Parameters estimates for C and n are jointly distributed. (Notice that this is not optional. It is how regression model parameters naturally behave. You can't choose the ratio of a circle's circumference to its diameter to be an integer because it might be more convenient. The fact is that π is inconveniently transcendental. Similarly, regression parameter estimates are asymptotically multivariate normal, and correlated, so any realistic simulation must sample from their correlated joint density.) Modeling them as bivariate normal in a MC simulation produces a standard deviation of 0.02802 in *log10* integrated lifetime for 1000 samples, which is very close to the standard deviation of the integrated individual Paris fits, 0.02769. The ratio of standard deviations is 1.012. **Modeling the joint behavior correctly reduces the greater than 700% error in the estimate of the standard deviation to about 1%.**

Notice, too, that replacing a constant n (the horizontal line in Figure 4) with a (conditional) probability density has the paradoxical effect of *decreasing* the resulting variability in calculated lifetime, since it corrects Mistake #2 (see Tables 1 and 2). This refutes the common misconception that replacing a constant with a probability density in a Monte Carlo simulation always results in increased scatter in the output. All these results are summarized in Tables 1 and 2 and in Figure 5.

Table 1- *MC Results Assuming Cycles are LogNormally Distributed*

			Correct	Mistake #1	Mistake #2	Mistake #3	Mistake #4
	Actual N	Eqn 3 N	C, n joint	C, n indept.	n fixed	C fixed	$C=b_0 + b_1*n$
mean	5.40916	5.39773	5.39909	5.41404	5.39414	5.39911	5.42217
stdev	0.03015	0.02769	0.02802	0.19778	0.14084	0.14872	0.01426
mean ratio		1.000	1.000	1.003	0.999	1.000	1.005
stdev ratio		1.000	1.012	7.143	5.087	5.371	0.515

Table 2 - *MC Results Assuming Cycles are Normally Distributed*

			Correct	Mistake #1	Mistake #2	Mistake #3	Mistake #4
	Actual *N*	Eqn 3 *N*	*C*, *n* joint	*C*, *n* indept.	*n* fixed	*C* fixed	$C=b_0 + b_1*n$
mean	257,165	250,389	251,182	287,841	261,278	265,859	264,487
stdev	18,447	16,332	16,240	140,261	88,189	94,453	8,703
mean ratio		1.000	1.003	1.150	1.043	1.062	1.056
stdev ratio		1.000	0.994	8.588	5.400	5.783	0.533

These results are from random samples of 1000. Other random samples would differ slightly. The superiority of the lognormal model over the normal model for these data is evident from the behavior of the ratios of the means to the baseline mean, with the skewed data causing the symmetrical normal's means to be overestimated slightly (penultimate row). While the choice of distribution model does influence the numerical results, it does not change the conclusion that common errors in Monte Carlo modeling can overestimate the variability by factors of seven or more, or underestimate it by half, and thus aren't even consistently conservative or anticonservative.

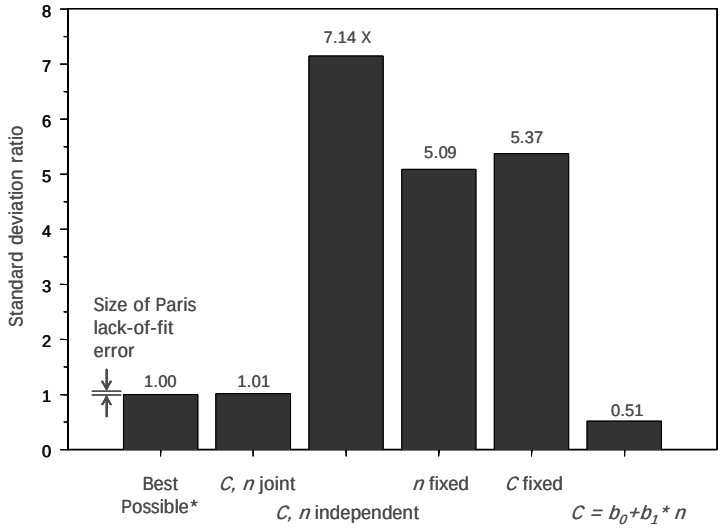

Figure 5- *Assuming C and n independent overestimates scatter by 7.1 X*

How to Sample from a Joint Probability Density:

As a consequence of the Central Limit Theorem in statistics (see the Appendix), regression model parameters are asymptotically multivariate normal. Thus while the assumption of Gaussian behavior isn't always appropriate for physical parameters, it *is* often justified for regression parameters. The following algorithm can be used to sample from a bivariate normal density.

Let z_1, z_2 be *iid* (independent and identically (from the same probability density) distributed) *N(0,1)*, and let $x_1 = \mu_1 + s_1 z_1$ and

$$(x_1, x_2)' \sim BN\left[(\mu_1, \mu_2)', s_1, s_2\ \rho_{1,2}\right] \qquad \text{then}$$

$$x_2 = \mu_2 + s_2\left(\rho_{1,2}\, z_1 + \sqrt{1-\rho_{1,2}^2}\; z_2\right) \tag{4}$$

where the symbol "~" is read "is distributed as," $N(\mu, \sigma^2)$ is a normal density with mean μ and variance σ^2, and $BN[(\mu_1, \mu_2)', \sigma_1^2, \sigma_2^2, \rho]$ represents a bivariate normal density with means μ_1, μ_2, variances σ_1^2, σ_2^2, and correlation, ρ. Equation 4 can be generalized to higher dimension regression models, which will of course require the parameter covariance matrix as the extension of s_1, s_2 and $\rho_{1,2}$ here.

Sampling from higher dimension, non-normal, joint densities:

Generalizations of equation 4 are *not* often justified in modeling the joint behavior of physical parameters since such situations are rarely multivariate normal. Rare too is statistical independence. It is lamentable then that many Monte Carlo users hope to avoid these difficulties by assuming them away. (If all the variables are assumed to be independent, then their marginal densities can be used.) Convenience, however, is scant justification, and consensus is a poor measure of veracity. Mother Nature will do as she will whether our simulations portend effectively or not.

All is not hopeless, however, and great progress has been made by taking advantage of conditional independence, and modeling the joint density as a network connected by statements of conditional probability [5]. A practical example is presented in [6].

Another underappreciated difficulty with direct-sampling Monte Carlo is what is referred to in the Bayesian literature as the "curse of dimensionality" [*cf*: 7]. This is the requirement that the number of sampled points must increase exponentially with the number of random variables to maintain a given level of precision. This places a practical limit on direct-sampling Monte Carlo.

A "new" method, Markov Chain Monte Carlo, (fifty years old but only recently rediscovered [8]) isn't encumbered by this impediment. Direct-sampling methods must sample directly from the entire probability space to obtain a sample from the joint probability density of interest. In contrast, Markov Chain Monte Carlo methods can sample directly from the desired joint probability density itself. Because they do not have to sample everywhere in the probability space, and only sample where the variables most probably reside, MCMC methods are not fettered by the problem of large dimensions. MCMC has revolutionized Bayesian statistics during the past decade, yet remains almost unknown to the engineering community.

Putting Things in Perspective:

If engineering Monte Carlo analysis is vulnerable to such enormous errors why do so many MC studies produce reasonable results? Here, flouting Murphy's Law, serendipity provides an explanation. First, many physical phenomena are indeed statistically independent, and thus do not conflict with that implicit assumption. *All*[3] regression model parameters *are* correlated, however, so any MC simulation based on them is

[3] Under some circumstances, for example when the data are centered at $\overline{X}, \overline{Y}$ some of the model covariances are zero.

vulnerable to the errors illustrated here. (For at least 75 years it has been well known in the applied statistics community that regression model parameters are correlated [*cf.:* 9], yet that fact is almost universally unknown to us engineers.)

The effects of ignoring model parameter correlations are sometimes mitigated by a second piece of good luck: the effects of scatter in model parameters may be overwhelmed by other sources of variability in the system being simulated. For example, the erroneously large standard deviation caused by treating Paris model parameters, C and n, as being independent is about 0.2 log_{10} units, in this instance. If there were another, independent source of variability of say 0.5 log_{10} units (admittedly a large error), the resulting effect of having ignored correlation would be about $(0.2^2 + 0.5^2)^{½} = 0.54$ log_{10} units, an increase in the total error of less than 10%.

Summary:

There is more to Monte Carlo simulation than replacing constants with probability densities. We have explored four common Monte Carlo modeling oversights and demonstrated their unacceptable consequences, using the 68 specimen Virkler-Hillberry data as an example. These errors and their consequences are not confined to the example data, nor to only Paris crack growth rate modeling, but must be considered in *any* Monte Carlo analysis that relies on regression models (and most do).

We have further demonstrated that correctly modeling the regression parameters as multivariate normal nearly eliminates the MC model error in this example.

Monte Carlo simulation is a powerful engineering analysis tool. Used properly it can provide insights that are otherwise unattainable. Lamentably, many practitioners are not aware of the statistical assumptions they are making, and that violating any one of them could eviscerate their analysis.

Acknowledgements:

I wish to thank Professor Ben Hillberry of Purdue University for graciously making the data, as well as specimen geometry and testing details, available for this study. I also wish to thank my longtime friend and colleague Dr. Al Berens of the University of Dayton Research Institute for suggesting the data to me.

APPENDIX - Review of Probability Relationships:

joint probability: $f(x, y \mid \theta)$ where f is the probability of x and y together as a pair, given the distribution parameters, θ.

multivariate distribution: A joint probability density of two or more variables. It is often summarized by a vector of parameters, θ. For example, the MVnormal is summarized (sufficiently) by a mean vector and covariance matrix.

marginal probability: : $f(x \mid \boldsymbol{\theta})$ where f is the probability density of x, for all possible values of y, given the distribution parameters, $\boldsymbol{\theta}$. The marginal probability is determined from the joint distribution of x and y by integrating over all values of y, thus integrating out the variable y. In applications of Bayes's Theorem, y is often a matrix of possible parameter values.

conditional probability: $f(x \mid y; \boldsymbol{\theta})$ where f is the probability of x by itself, given specific value of variable y, and the distribution parameters, $\boldsymbol{\theta}$. If x and y represent events A and B, then $P(A|B) = n_{AB}/n_B$,where n_{AB} is the number of times both A and B occur, and n_B is the number of times B occurs. $P(A|B) = P(AB)/P(B)$, since $P(AB) = n_{AB}/N$ and $P(B) = n_B/N$ so that $P(A|B) = \dfrac{n_{AB}/N}{n_B/N} = n_{AB}/n_B$ Note that in general the conditional probability of A given B is **not** the same as B given A. The probability of **both** A and B together is $P(AB)$, and $P(A|B) \times P(B) = P(AB) = P(B|A) \times P(A)$, if both $P(A)$ and $P(B)$ are non-zero. This leads to a statement of **Bayes's Theorem**: $P(B|A) = P(A|B) \times P(B)/P(A)$. Conditional probability is also the basis for statistical dependence and independence.

Joint, marginal and conditional densities are summarized in Figure A-1.

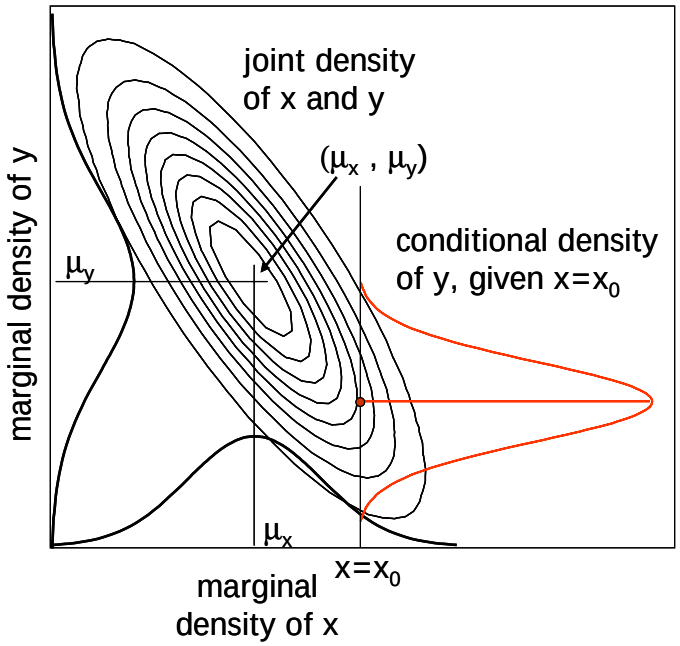

Figure A-1 - *Schematic Showing Joint, Conditional and Marginal Densities*

The **Central Limit Theorem** justifies using a multivariate normal density to model the collective behavior of regression model parameters. The CLT states that the distribution of an average tends to be normal, even when the distribution from which the average is computed is decidedly *non*-normal. Furthermore, this normal distribution will have the same mean as the parent distribution, and variance equal to the variance of the parent divided by the sample size. This result does *not* depend on the original distribution of x, only that the mean and variance are finite. And "large" n may be on the

order of a dozen observations. Formally the CLT says that if $x_1, x_2, \ldots x_n$ are a sequence of independent identically distributed (*iid*) random variables, with finite mean μ_x and variance σ_x^2 then z_n *converges in distribution* to *N(0, 1)* as *n* becomes large, and

$$z_n = (\bar{x}_n - E(\bar{x}_n) / \sqrt{\text{var}(\bar{x}_n)} = (\bar{x}_n - \mu_x) / (\sigma_x / \sqrt{n})$$

where *E( )* is the expectation (averaging) operator. For a discrete density, *f(x)*, the expectation operator is $E(x)=\Sigma xf(x)$, and for continuous density, *f(x)*, $E(x) = \int xf(x)dx$

References:

[1] Virkler, D.A., B.M. Hillberry, and P.K. Goel (1979), "The Statistical Nature of Fatigue Crack Propagation," AFFDL-TR-78-43, Air Force Flight Dynamics Laboratory, April, 1978

[2] Ostergaard, D.F, and B. M. Hillberry, "Characterization of the Variability in Fatigue Crack Propagation Data," *Probabilistic Fracture Mechanics and Fatigue Method: Applications for Structural Design and Maintenance, ASTM STP 798*, J.M. Bloom and J.C. Ekvall, Eds., American Society for Testing and Materials, 1983, pp. 97-115

[3] Ildiko Frank and Jerome Freidman "A Statistical View of some Chemometrics Regression Tools," *Technometrics*, V35, N2, 1993, p 110

[4] Annis, C.G., Jr., R. M. Wallace, and D. L. Sims, "An Interpolative Model for Elevated Temperature Fatigue Crack Propagation," AFML-TR-76-176, Part I, November 1976, presented at 1977 Society for Experimental Stress Analysis Spring Meeting, Dallas, Texas, May 1977.

[5] Jordan, Michael I, (editor) *Learning in Graphical Models*, MIT Press, 1999

[6] Annis, Charles, "Bayesian Network Analysis of the 1-D Oscillator Problem - A Directed Acyclic Graph Paradigm for Probabilistic High Cycle Fatigue Risk Assessment," supported under Independent Contractor Agreement 01-S441-48-01-C4, Universal Technology Corporation Prime Contract F33615-98-C-2807, September, 2001

[7] Carlin, Bradley P. and Thomas A. Louis, *Bayes and Empirical Bayes Methods for Data Analysis*, Chapman and Hall/CRC, 1996

[8] Annis, Charles, "Modeling High Cycle Fatigue with Markov Chain Monte Carlo: A New Look at an Old Idea," AIAA 2002-13800, presented at 43rd AIAA/ASME/ASCE/AHS Structures and Dynamics Conference, Denver, CO, 22-25 April, 2002

[9] Fisher, Ronald A., *Statistical Methods for Research Workers.* (First published in 1925; 14th edition was ready for publication in 1962, when Fisher died, and was published in 1990, by the Oxford University Press, along with *Experimental Design* and *Scientific Inference*.)