LA-UR-16-20344
Approved for public release; distribution is unlimited.

Title: Pu239 Cross-Section Variations Based on Experimental Uncertainties and Covariances

Author(s): Sigeti, David Edward
Williams, Brian J.
Parsons, D. Kent

Intended for: Report

Issued: 2016-10-18 (rev.1)

PU239 CROSS-SECTION VARIATIONS BASED ON EXPERIMENTAL UNCERTAINTIES AND COVARIANCES

DAVID E. SIGETI

LANL, XCP-8

BRIAN J. WILLIAMS

LANL, CCS-6

D. KENT PARSONS

LANL, XCP-5

Abstract. Algorithms and software have been developed for producing variations in plutonium 239 neutron cross sections based on experimental uncertainties and covariances. The varied cross-section sets may be produced as random samples from the multivariate normal distribution defined by an experimental mean vector and covariance matrix, or they may be produced as Latin-Hypercube/Orthogonal-Array samples (based on the same means and covariances) for use in parametrized studies. The variations obey two classes of constraints that are obligatory for cross-section sets and which put related constraints on the mean vector and covariance matrix that detemine the sampling. Because the experimental means and covariances do not obey some of these constraints to sufficient precision, imposing the constraints requires modifying the experimental mean vector and covariance matrix. Modification is done with an algorithm based on linear algebra that minimizes changes to the means and covariances while insuring that the operations that impose the different constraints do not conflict with each other.

E-mail addresses: sigeti@lanl.gov, brianw@lanl.gov, dkp@lanl.gov.
Date: May 6, 2016.

Contents

1. Introduction

This report documents algorithms developed by the authors for producing variations in plutonium 239 neutron cross sections based on experimental uncertainties and covariances. Two kinds of sampled variations may be produced:

(1) Random samples from a multivariate normal distribution for cross sections defined by an experimental mean vector and covariance matrix; or
(2) Latin-Hypercube/Orthogonal-Array (LH/OA) samples for use in parametrized studies of the effects of cross-section and other uncertainties on the outputs of physics codes for which the cross sections and other varied quantities are inputs.

The task of performing either kind of sampling is complicated by the need for the samples to obey certain constraints, a nonnegativity constraint (Section 4.1) and a set of (linear) summation constraints (Section 4.2). These constraints on the sampled cross sections in turn put

constraints on the mean vector and covariance matrix that define the two sampling processes, constraints that the experimental means and covariances only imperfectly obey. We have developed a methodology for "repairing" the mean vector and covariance matrix that minimizes changes to the means and covariances while insuring that the operations that impose the different constraints do not conflict with each other. This methodology is explained in Section 4 on constraints.

Random sampling from a multivariate normal distribution in which the covariance matrix obeys summation constraints cannot be handled by the standard method for multivariate normal sampling, and the task is further complicated by the need for the samples to strictly obey the nonnegativity constraint. We explain how to do such sampling in Section 5.

LH/OA sampling has a couple of additional complications, in addition to the need to respect the constaints that it shares with normal sampling. In the first place, handling correlated variations in LH/OA sampling is not straightforward. The method for sampling from a correlated and constrained normal distribution that we describe in Section 5 provides a basis for doing LH/OA sampling that both obeys the constraints and respects the correlations between cross sections, as we explain in Section 6.

In addition, LH/OA sampling requires us to drastically reduce the degrees of freedom in the cross-section variations in order to characterize them with a small number of parameters, as is required for combining cross-section variations with other uncertainties in integrated studies. Clearly, we need to do so in a way that captures the variations in cross sections that maximize the variations in the outputs of the physics codes to which the cross sections are inputs. Here, as in the problem of correlated and constrained LH/OA sampling, the method for doing constrained and correlated normal sampling provides a basis for solving the problem, as we explain in Section 6.

In Section 7 we explain how to do a modified type of LH/OA sampling that uses *flux weights* to increase the ability to capture the cross section variations that maximize the variations in physics code outputs with a minimum of parameters. This kind of sampling captures 98% of the variance in the weighted cross sections with just three parameters.

We begin in Section 2 immediately below with a brief discussion of the source and structure of our experimental mean vector and covariance matrix. In Section 3, we then present some elementary formulas and facts about multivariate normal distributions and covariance matrices that we will use in the rest of the report. In Section 4, we begin

the main technical topics of the paper. We conclude with appendices presenting some mathematical results used in the body of the paper.

2. Experimental Cross Section Mean Vector and Covariance Matrix

2.1. **Sources.** The experimental mean vector, μ_e, and covariance matrix, Σ_e, used in this work were produced by one of the authors (Parsons) using the LANL nuclear cross-section processing code NJOY [3] based on evaluated nuclear data from the ENDF/B VII.1 library [1].

2.2. **Cross-Section Structure.** A complete set of Pu239 neutron cross sections is specified for the total cross section and for six reactions (elastic scattering, inelastic scattering, n2n, n3n, fission, and absorption)[1]. Within each of these seven blocks, cross sections are specified for each of 30 energy groups, which are the same across all blocks. There are thus 210 individual cross sections.

The experimental mean vector for the cross sections naturally has the same structure as just described for a cross-section set.

The experimental covariance matrix for the cross sections necessarily operates on a vector space with this structure and is thus a 210x210 matrix. The diagonal elements of this matrix give the variances (squares of the standard deviations) of the individual cross sections, and the off-diagonal elements provide information on the correlations between cross sections.

All of the experimental means and covariances are given in single precision.

3. The Multivariate Normal Distribution and Covariance Matrices

In this section, we present some elementary facts that we will use in the rest of the report about multivariate normal distributions and covariance matrices.

3.1. **The Multivariate Normal Distribution for Cross Sections.** For the purposes of producing cross-section variations consistent with a mean vector, μ, and covariance matrix, Σ, we treat these quantities

[1]The total cross section is specified separately because it is measured separately and very accurately.

as defining a multivariate normal distribution for a cross-section set, s,

$$\begin{aligned} p(s) &\equiv \mathrm{N}(s \,|\, \mu, \Sigma) \\ &\equiv (2\pi)^{-k/2} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2} (s-\mu)^T \Sigma^{-1} (s-\mu) \right) \end{aligned} \tag{1}$$

where k is the number of cross-sections.

3.2. **Covariance Matrices: Structure and Constraints.** In this section, we give a brief description of the structure and certain properties of covariance matrices that we will use extensively in the remainder of this report.

The covariance matrix for a multivariate random variable encodes information about both the uncertainties of the individual components of the random vector and about the correlations between the components. The diagonal elements of the covariance matrix specify the uncertainties, as they are the variances (the squares of the standard deviations) of the corresponding components. The off-diagonal elements, on the other hand, contain the information about the correlations. Specifically, for a multivariate random variable, s, with a covariance matrix, Σ,

$$\Sigma_{ij} \equiv \mathrm{E}\big(\left[s_i - \mathrm{E}(s_i) \right] \left[s_j - \mathrm{E}(s_j) \right] \big), \tag{2}$$

where E is the expectation operator.

This structure, together with the requirement that the multivariate normal distribution defined in Equation 1 be a proper, normalizable probability distribution, puts a number of constraints on the covariance matrix.

(1) Because they are the variances of random variables, it is clear that the diagonal elements must be nonnegative;
(2) It is clear from the definition in Equation 2 of the off-diagonal elements of the covariance matrix that the matrix must be symmetric, as interchanging i and j does not change the value of Σ_{ij}.
(3) In addition, for the multivariate normal distribution defined by the mean and covariance to be normalizable, it is necessary that the covariance matrix be nonnegative-definite.

We might expect from the presence of the inverse of Σ in Equation 1 that the covariance matrix would need to be invertible and thus not just nonnegative-definite but actually positive-definite. Note that this is not the case. A non-trivial null space for the covariance matrix corresponds to a direction in which the normal distribution does not vary, indicating that the probability is confined to an affine subspace in the k-dimensional space of s. In addition, as we will show in Section 5,

the algorithm for sampling from a multivariate normal does not require the inverse of the covariance matrix and can be modified to handle the case where the matrix has a non-trivial null space.

Some properties of symmetric nonnegative-definite matrices and thus of covariance matrices that we will use in the following are,

(1) A symmetric matrix has a spectral factorization,

$$\Sigma = UDU^T, \tag{3}$$

where U is an orthogonal matrix the columns of which are the normalized eigenvectors of Σ and D is a diagonal matrix of corresponding eigenvalues, all of which are real.

(2) A symmetric matrix is nonnegative-definite if and only if all of its eigenvalues are nonnegative.

(3) We may therefore re-write the spectral factorization of a nonnegative-definite matrix in the form,

$$\Sigma = AA^T, \tag{4}$$

where,

$$A = U\sqrt{D}, \tag{5}$$

and the square root may be taken because the eigenvalues are nonnegative.

(4) For a nonnegative-definite matrix A, the off-diagonal elements satisfy

$$|A_{ij}| \leq \sqrt{A_{ii}A_{jj}}. \tag{6}$$

An additional property of covariance matrics that we will use is that, if x is a multidimensional random variable with covariance matrix Σ,

$$\mathrm{cov}(x) = \Sigma, \tag{7a}$$

then a random variable, y defined by multiplying x by a matrix, A,

$$y \equiv Ax, \tag{7b}$$

has covariance matrix $A\Sigma A^T$.

$$\mathrm{cov}(y) = A\Sigma A^T. \tag{7c}$$

4. Constraints on Cross Sections

The task of sampling from the multivariate normal distribution for cross sections that is defined by the experimental mean vector and covariance matrix is complicated by the requirement that cross-section sets produced by the sampling must obey a number of constraints. In particular,

(1) No cross section may be negative. This requirement is absolute, as the physics codes that use the cross sections cannot, in general, handle negative cross sections.
(2) For each energy group, the cross section for that group in the total block must equal the sum of the cross sections for that group in the individual reaction blocks. With finite-precision arithmetic this requirement can only be satisfied approximately so this is all that the physics codes require.

These two basic constraints in turn put a number of constraints on the mean vector and covariance matrix, not all of which are obeyed to sufficient precision by the experimental mean and covariance, which then requires that we "repair" the mean and covariance before we can sample from them. In the next two sub-sections we discuss these constraints on the mean vector and covariance matrix and how the mean and covariance may be modified to better satisfy the constraints.

In addition to the constraints on the mean and covariance arising directly from the nonnegativity and summation constraints on cross-section sets, the constraint that the covariance matrix be nonnegative-definite will require an additional "repair" of the matrix. This repair comes naturally after the imposition of the summation constraints, so we discuss it in the final sub-section of this section.

4.1. **The Nonnegativity Constraint.** Strictly speaking, sampling from a normal distribution cannot obey a nonnegativity constraint, as the support of the normal distribution is the entire real line. However, if every component of the mean vector is positive and the corresponding standard deviation from the covariance matrix is small compared to the mean, samples with negative components will be rare. In this situation and for many applications, it makes sense simply to disgard such cases. This is the strategy that we adopt in our work because, given our experimental mean vector and covariance matrix, for both normal and LH/OA sampling, including the latter with flux weights, the fraction of cross-section sets with negative cross sections is never larger than about 0.06

In addition to being able to handle the case where standard deviations are small with respect to means by simply disgarding the resulting rare cross-section sets with negative cross sections, it is also possible to accomodate zero components in the mean vector if the corresponding standard deviations are also zero, although some care is required in handling such components, as we discuss in the remainder of this subsection.

In any case, the requirement that no cross section be negative obviously requires that no element of the mean vector be negative which, thankfully, is the case with our experimental mean vector.

Our experimental mean vector does have some zero cross sections[2]. Clearly, a zero for a component of the mean vector must be matched by a zero for the corresponding variance on the diagonal of the covariance matrix because otherwise sampling will produce a negative cross section half the time. Fortunately, the experimental covariance matrix has a zero variance for every zero component of the experimental mean vector.

A somewhat less obvious constraint that a zero value in the i-th component of the mean vector puts on the covariance matrix is that the entire i-th row and column of the matrix must be zero. This follows from Equation 6 and the requirement that the variance of such a component be zero. Again, our experimental covariance matrix obeys this constraint.

Since the nonnegativity constraint must be strictly obeyed, it is essential to insure, when we are drawing samples from the multivariate normal distribution in Equation 1, that roundoff error from finite-precision arithmetic in the sampling computation not produce small nonzero values in components of samples that should be strictly zero, as there is no way to guarantee that such small nonzero values will not be negative. Since the sampling algorithms described in Sections 5 and 6 involve the computation of the spectral decomposition of the covariance matrix and the construction of samples as linear combinations of the eigenvectors of the matrix, a complex computation that operates globally over the entire covariance matrix, pretty much the only way to insure that roundoff error does not produce small nonzero values in the components of the samples that should be zero is to remove the zero components of the mean vector and the corresponding zero rows and columns of the covariance matrix prior to performing the sampling computation, and then to restore the zero components in the samples after they are produced.

In addition, as we discuss in the next subsection, the experimental mean vector and covariance matrix do not obey the summation constraints to sufficient precision and it is therefore necessary to modify them slightly in order to improve the compliance with those constraints. As with the computation of samples, the need to prevent roundoff error

[2]These zeros occur in blocks for reactions that have an energy threshold—the inelastic scattering, n2n, and n3n blocks—so they are always in the lower energy groups in these blocks.

in this computation from corrupting the components of the mean vector and the rows and columns of the covariance matrix that should remain zero requires us to remove these before performing any modifications on the mean vector or the covariance matrix, and to restore the removed zero components, rows, and columns only after the modifications have been made.

4.2. **The Summation Constraints.** The summation constraint on cross-section sets described above may be represented in matrix form by defining a constraint matrix,

$$C^T = \left[\mathrm{I}_g \quad \overbrace{-\mathrm{I}_g \quad -\mathrm{I}_g \quad \cdots \quad -\mathrm{I}_g}^{b-1 \text{ blocks}} \right] \tag{8}$$

where g is the number of energy groups, I_g is the $g \times g$ identity matrix, b is the number of blocks (the number of reactions plus one for the total block), and it is assumed that the total block comes first. With this notation, the summation constraints on a cross-section vector, s, may be written as,

$$C^T s = 0. \tag{9}$$

If want our sampled cross-section sets to obey these constraints, it is clear that the mean vector, μ, for our multivariate normal distribution (Equation 1) must also obey this constraint,

$$C^T \mu = 0. \tag{10}$$

The question then arises of what we are to do if the experimental mean vector, μ_e, does not obey the constraint in Equation 10 to sufficient precision. It appears that we should somehow repair the mean vector in that case. To see how to do this, begin by noting that the rows of C^T (the columns of C) are vectors in the space of cross sections and that they therefore span a subspace, $\mathscr{C}$. Equation 10 tells us that μ must be orthogonal to each of these vectors. Geometrically, this means that μ must lie in the orthogonal complement of $\mathscr{C}$, $\mathscr{C}_\perp$. The natural way to repair μ_e if it does not, in fact, lie in $\mathscr{C}_\perp$, is to compute the projection of μ_e onto $\mathscr{C}_\perp$ and use the projected vector as our repaired μ.

To compute this projection of μ_e, note that the columns of C are orthogonal, as, for any index i and any two columns of C, c and d, it is never the case that both c_i and d_i are nonzero[3]. If we normalize the columns of C to produce a modified matrix $\widehat{C}$, we then have an

[3]This lack of shared nonzero components arises because there is a one-to-one correspondance between vectors and groups and each vector has nonzero values only for cross sections for its group.

orthonormal basis for $\mathscr{C}$. It is easy to show that, given that $\widehat{C}$ is a matrix of column vectors that constitute an orthonormal basis for $\mathscr{C}$, the matrix that projects an arbitrary vector onto $\mathscr{C}$ is given by,

$$P_{\mathscr{C}} = \widehat{C}\widehat{C}^T. \tag{11}$$

Then, the matrix that projects onto $\mathscr{C}_\perp$ is,

$$P_{\mathscr{C}_\perp} = \mathrm{I} - \widehat{C}\widehat{C}^T, \tag{12}$$

where I is the identity matrix on the space of cross sections. Therefore, we may repair μ_e simply by multipying it by $P_{\mathscr{C}_\perp}$,

$$\mu \equiv P_{\mathscr{C}_\perp}\mu_e \equiv \left(\mathrm{I} - \widehat{C}\widehat{C}^T\right)\mu_e. \tag{13}$$

In addition to imposing the constraint in Equation 10 on the mean vector, it is clear that the requirement that samples from our distribution obey the constraint in Equation 9 must impose constraints on the covariance matrix as well. If we look at the second expression for the multivariate normal distribution in Equation 1, it is clear that we want $(s-\mu)^T\Sigma^{-1}(s-\mu)$ to be very large (infinite, really) whenever s is not in $\mathscr{C}_\perp$. Since our repaired μ is in $\mathscr{C}_\perp$ and $\mathscr{C}_\perp$ is a linear subspace, s not being in $\mathscr{C}_\perp$ is equivalent to $s-\mu$ not being in $\mathscr{C}_\perp$. So, we want $u^T\Sigma^{-1}u$ to be infinite for any cross-section vector u that does not lie in $\mathscr{C}_\perp$.

This will be the case if and only if Σ operating on any vector in $\mathscr{C}$ gives zero,

$$\Sigma\, C = 0. \tag{14}$$

This is the constraint on the covariance matrix induced by the constraint on samples given in Equation 9.

If Σ_e does not satisfy this constraint to sufficient precision, we may fix it by replacing it with the covariance matrix for vectors that have been projected onto $\mathscr{C}_\perp$,

$$\begin{aligned}\Sigma &\equiv P_{\mathscr{C}_\perp}^T \Sigma_e P_{\mathscr{C}_\perp} \\ &= P_{\mathscr{C}_\perp} \Sigma_e P_{\mathscr{C}_\perp} \\ &= \left(\mathrm{I} - \widehat{C}\widehat{C}^T\right)\Sigma_e\left(\mathrm{I} - \widehat{C}\widehat{C}^T\right),\end{aligned} \tag{15}$$

where we have used the fact that $P_{\mathscr{C}_\perp}$, being a projection matrix, is symmetric (as is also obvious from the definition of $P_{\mathscr{C}_\perp}$ in terms of C). It should be immediately clear that Σ so defined satisfies the constraint in Equation 14.

The actual proceedure for repairing μ_e and Σ_e is more complex than described above, because of the presence of the zero components in μ_e and the corresponding zero rows and columns in Σ_e that must be

exactly preserved in the repaired versions. The solution is to remove the zero components, rows, and columns, perform the repair on reduced versions of μ_e and Σ_e, and then restore the zero components, rows, and columns to the reduced, repaired versions. Note that this requires removing the same components from the columns of C as are removed from μ_e and renormalizing the columns in the new, reduced space. Note also that the reduction does not change the orthogonality of the columns of C, because the orthogonality is not due to any cancelation of quantities in a dot product but rather to the fact that there no two columns have nonzero values for the same component so every term in the relevant dot product is identically zero.

The proceedure for fixing μ_e and Σ_e is therefore as follows:

(1) Construct the constraint matrix C according to Equation 8.
(2) Remove the zero components of μ_e and the corresponding rows and columns of Σ_e to produce reduced versions, μ'_e and Σ'_e.
(3) Remove the corresponding rows of C to produce a reduced version, C'.
(4) Normalize the columns of C' to produce a normalized, reduced version, $\widehat{C}'$.
(5) Compute the projection matrix onto the subspace orthogonal to the subspace spanned by $\widehat{C}'$, $P_{\mathscr{C}'_\perp}$, using Equation 12, but substituting $\widehat{C}'$ for $\widehat{C}$.
(6) Apply $P_{\mathscr{C}'_\perp}$ to μ'_e and Σ'_e as in Equations 13 and 15 to produce repaired, reduced versions, μ'_r and Σ'_r.
(7) Restore the missing zero components to μ'_r and the missing zero rows and columns to Σ'_r to produce our final, full, repaired versions, μ and Σ.

4.3. **The Nonnegative Definiteness Constraint.** In addition to the nonegativity and summation constraints, the covariance matrix is required to be nonnegative-definite. As discussed in Section 3.2 above, this nonnegative-definiteness is equivalent to all eigenvalues being nonnegative. Given that *some* of the eigenvalues of the covariance matrix must be zero, it is clear that this constraint can only be satisfied approximately, because round-off error from finite-precision arithematic is certain to make at least some of the zero eigenvalues have small negative values. This problem is likely to be worse than it needs to be because the experimental covariance matrix is only given in single precision. Since doing either sampling from the normal distribution in Equation 1 or sampling based on the normal distribution as part

of a Latin-Hypercube/Orthogonal-Array design involves complex computations with the covariance matrix, it seems best to insure that the matrix obeys the constraint on its eigenvalues to within double precision before doing any sampling.

This constraint on the eigenvalues may be enforced simply by taking the repaired Σ from the previous sub-section, where we imposed the summation constraints, performing the spectral factorization, setting all negative eigenvalues to zero, and inverting the factorization.

Note that there is no danger that such an operation will undo the summation constraints that have previously been imposed, as these constraints show up in the spectral factorization as a set of eigenvectors with (near) zero eigenvalues. Leaving these eigenvectors unmodified (which the operation does with all eigenvectors) and setting the correponding eigenvalues that have small negative values to exactly zero clearly will do nothing to violate the summation constraints.

In fact, it is important that the constraint on the eigenvalues be imposed only *after* the summation constraints are imposed precisely so that we may be certain that the eigenvectors that correspond to the summation constraints are exactly as they should be (to within double precision, of course).

In practice, along with setting all negative eigenvalues to zero, we identify the eigenvectors that correspond to the summation constraints, and set all of their eigenvalues (whether negative or not) to zero before inverting the spectral factorization. We identify the appropriate eigenvectors in the first place by noting that their eigenvalues should have the smallest absolute magnitudes (as we have just imposed the summation constraints) and then confirm this identification by showing that the subspace spanned by the identified eigenvectors is the same as the subspace of the summation constraints. We make the latter identification by using a standard method for comparing subspaces by finding the largest principal angle between them (see [2, Section 6.4.3]).

Of course, these operations, like pretty much any complex operations on the covariance matrix, should only be carried after removing the rows and columns of the matrix that should be exactly zero, with the removed rows and columns restored once the operation is completed. In practice, this can be done after the imposition of the summation constraints described in the preceeding section, or as a part of it after the reduced matrix is multiplied by the projection matrices for the constraints but before the zero rows and columns are restored.

With this final repair, we have a covariance matrix that satisfies all constraints as well as possible. Henceforth in this report, we will

assume that all computations using the mean vector and covariance matrix will be carried out using the fully repaired versions μ and Σ.

5. Sampling from a Normal Distribution with a Covariance Matrix with Zero Eigenvalues

If the covariance matrix, Σ, of a multivariate normal distribution has zero eigenvalues, then the conventional algorithm for sampling from a multivariate normal distribution cannot be used, as the algorithm relies on taking the Cholesky factorization of Σ, which cannot be done for a matrix that is not strictly positive-definite. There is, fortunately, an alternative algorithm that can be used in such cases. The alternative algorithm uses the spectral factorization in place of the Cholesky factorization.

Both algorithms depend on the following property of multivariate normal distributions, which is a simple consequence of the formula for the covariance of the linear transformation of a multidimensional random variable given in Equations 7: If x is a multidimensional random variable with a standard normal distribution,

$$x \sim \mathrm{N}(0, \mathrm{I}), \tag{16}$$

where I is the identity matrix, then a random variable defined by an affine transformation of x,

$$y \equiv \mu + Ax, \tag{17}$$

has a transformed normal distribution,

$$y \sim \mathrm{N}\left(\mu, AA^T\right). \tag{18}$$

Thus, given an arbitrary multivariate normal distribution with mean μ and covariance matrix Σ, $\mathrm{N}(\mu, \Sigma)$, we can sample from the distribution if we can write Σ as a product of a matrix and its transpose, $\Sigma = AA^T$. With this factorization of Σ, a sample may be obtained by taking a (multivariate) sample, x, from a standard multivariate normal of the appropriate size, multiplying the sample vector by the matrix A, and adding μ as in Equation 17. Note that a sample from a standard multivariate normal may be produced by any algorithm for producing independent samples from a standard univariate normal, as the identity covariance matrix in Equation 16 specifies that the individual components of the multivariate sample x are independent. More independent samples from the multivariate normal are obtained by simply repeating the process (assuming that the underlying random number generator continues to produce independent samples from the standard univariate normal).

We pointed out in Section 3.2 that the spectral factorization of any covariance matrix, Σ, (Equation 3) may be written in the form $\Sigma = AA^T$ (Equation 4), where $A = U\sqrt{D}$, U is the matrix the columns of which are the eigenvectors of Σ, and D is the diagonal matrix of eigenvalues. Thus, the spectral factorization provides precisely the kind of factorization of Σ that is needed for sampling from $\mathrm{N}(\mu, \Sigma)$, provided we set

$$A \equiv U\sqrt{D} \tag{19}$$

in Equation 17 to compute,

$$y = \mu + U\sqrt{D}x. \tag{20}$$

Numerically, the algorithm is a bit more complex, because of the effects of round-off error in the computation of the spectral factorization of Σ. Even if Σ is strictly nonnegative-definite, with some zero but no negative eigenvalues, round-off error from finite-precision arithmetic will cause some eigenvalues in the computed spectral factorization to be slightly negative. The obvious solution is to set all of the negative eigenvalues in D to zero before taking the square root. If this prescription is followed, then the algorithm as described above properly samples from the multivariate normal distribution defined by μ and Σ, whether or not Σ has zero eigenvalues.

The presence of zero components of μ complicates the sampling algorithm further, as it is essential that the samples exactly reproduce the zero components. This is accomplished as in the discussion of imposing summation constraints in Section 4.2 by producing reduced versions of μ and Σ with the zero components of μ and the corresponding rows and columns of Σ removed, performing the sampling algorithm described above using the reduced μ and Σ, and restoring the zero components, rows, and columns in the resulting samples. This works because the eigenvectors of the reduced Σ become eigenvectors of the full Σ once we restore the zero components, which follows from the fact that the removed components were zero.

Note that computing $y = \mu + U\sqrt{D}x$, where D is diagonal and nonnegative, is equivalent to computing

$$y = \mu + \sum_i x_i \sqrt{\lambda_i} u_i \tag{21}$$

where the x_i are the components of x, the λ_i are the elements of the diagonal matrix D (the eigenvalues of Σ) and the u_i are the columns of U (the corresponding normalized eigenvectors of Σ). We may thus view a sample from the multivariate normal $\mathrm{N}(\mu, \Sigma)$ as the sum of μ and a linear combination of the eigenvectors of Σ, with each eigenvector

weighted by the square root of its eigenvalue and multiplied by a random sample from a standard univariate normal. This way of looking at the sampling in Equation 20 will prove useful in Section 6, where we explain how to produce cross-section variations for parametric studies based on Latin-Hypercube/Orthogonal-Array sampling.

It is also important to note that the eigenvectors of Σ for which the corresponding eigenvalues are non-zero will be orthogonal to any vector, v, for which $\Sigma\, v = 0 = 0 \cdot v$, because such a v is clearly an eigenvector of Σ with eigenvalue 0 and eigenvectors of a symmetric matrix with distinct eigenvalues are orthogonal. In particular, if v represents a constraint on Σ as in Equation 14, then, provided μ obeys the related constraint in Equation 10, a sample, s, formed from the sum of μ and a linear combinations of the eigenvectors of Σ as in Equations 20 and 21 will obey the constraint in Equation 9, as the eigenvectors with eigenvalue 0 contribute nothing to the sample by virtue of the weighting by $\sqrt{\lambda_i}$.

Finally, we note that the eigenvectors of Σ contribute to the variance of the cross-section variations produced by sampling from $\mathrm{N}(\mu, \Sigma)$ in proportion to their eigenvalues. This observation will be important in Section 6 where we seek to reduce the cross-section variations to functions of a small number of parameters.

6. Latin-Hypercube/Orthogonal-Array Sampling

In addition to being able to sample from the multivariate normal distribution defined by the experimentally-determined mean vector μ and covariance matrix Σ (both properly repaired of course), we also want to be able to produce varied cross-section sets appropriate for parametrized studies of the effects of various uncertainties on outputs of physics codes. This requires being able to compute cross-section variations as a function of a small number of parameters, certainly no more than, say, six. Given that the number of cross sections is 210 (albeit with 30 constraints, a number of cross sections with zero mean and variance, and substantial correlations), this is a formidable demand for dimensional reduction.

Values of these intended parameters for a study, together with values of other parameters representing other kinds of uncertainties, are to be produced as part of a Latin-Hypercube/Orthogonal-Array (LH/OA) design. The inputs for such a design are a minimum and a maximum value (that is, a range), for each parameter. These ranges collectively define a hypercube in the multidimensional space of the parameters.

Given these inputs, the LH/OA design algorithm produces a collection of parameter sets that are optimally distributed within this hypercube.

For example, if one has several parameters, each of which may be varied independently of the others and for each of which the uncertainty may be characterized by a normal distribution, one will typically take 2-σ or 3-σ ranges for the parameters as the inputs to the LH/OA algorithm.

However, it is not clear how to apply the LH/OA algorithm when it does not make sense to vary certain subsets of the parameters independently, as when their variations should be correlated or when they must obey certain multivariate constraints. The solution is to write the parameters in question as (usually affine) functions of a new set of the parameters the values of which *can* be varied independently.

In incorporating cross-section variations into our studies, then, we seek a set of vectors in the space of cross sections that meet the following requirements:

(1) There is a natural way to write a cross-section set as a linear combination of the vectors (plus a mean vector);
(2) Such a cross-section set is guaranteed to satisfy the constraints;
(3) The coefficients of the vectors in the linear combination that computes a cross-section set can reasonably be varied independently;
(4) There is a reasonable way to choose these coefficients so as to at least approximate the correlated uncertainties that characterize our knowledge of the cross sections (as represented by the multivariate normal distribution in Equation 1).
(5) It can be shown that a small number of these vectors are sufficient to capture the variations in the cross sections that significantly affect the outputs of our physics codes.

If we can find such a set of vectors, then the coefficients in the linear combination that produces a cross-section set become our new parameters that can be varied independently.

The eigenvectors of the covariance matrix, weighted by the square roots of their eigenvalues, that are used to do random sampling from the multivariate normal distribution in Equation 21 above are good candidates for such a set of vectors because:

(1) Equation 21 gives a general cross-section set as a linear combination of the eigenvectors;
(2) A cross-section set computed using Equation 21 is guaranteed to satisfy the constraints;

(3) It makes good sense to vary the coefficients of the eigenvectors (the x_i in the linear combination of eigenvectors in Equation 21) independently in an LH/OA design as the eigenvectors are statistically independent in $\mathrm{N}(\mu, \Sigma)$.
(4) Choosing the coefficients x_i as independent samples from a standard univariate normal gives us our samples from the multivariate normal distribution in Equation 1 through Equation 21, so the eigenvectors are directly related to our knowledge of the cross sections embodied in the multivariate normal.
(5) The eigenvalue corresponding to an eigenvector quantifies the contribution of that eigenvector to the overall variance in the cross sections, suggesting that we may be able to reduce the dimensionality of our representation of cross-section variations by using only the eigenvectors with the largest eigenvalues.

In order to use the eigenvectors in a LH/OA design, then, we order the eigenvectors[4] in descending order of their eigenvalues and choose a number of eigenvectors to retain, m. When building our LH/OA design we specify m cross-section variation parameters. The values of these parameters in each parameter set in the design will just be substituted for the first m x_i in Equation 21 to produce the actual cross-section samples. If we want to explore a 2-σ range, we specify the ranges of the x_i to run from -2 to 2 and *mutatis mutandis* for 3-σ ranges or any other desired range.

Note that there is no guarantee that using only the eigenvectors with the largest eigenvalues will actually meet our last requirement for the vectors on the coefficients of which we will build our LH/OA design because the vectors that produce the maximum variance in the cross-sections are not necessarily the vectors that will produce the largest variations in the outputs of the physics codes that use the cross sections. Nevertheless, the eigenvectors are the obvious set with which to start and any more sophisticated choice for the vectors to use in the design should probably be expressed as linear combinations of the eigenvectors with non-zero eigenvalues, as any such linear combinations will be guaranteed to obey the constraints and the relationship of the eigenvectors to sampling from the distribution for the cross sections, $\mathrm{N}(\mu, \Sigma)$, allows us to map their coefficients onto the LH/OA model.

[4]Of course, we must compute the eigenvectors as described in Section 4.2, working on reduced versions of μ and Σ with the zero components of μ and the corresponding rows and columns of Σ removed, in order to satisfy the requirement to preserve the zero components in the samples exactly.

7. Latin-Hypercube/Orthogonal-Array Sampling with Per-Group Flux Weights

In this section, we describe how we use flux weights for the energy groups in our Latin-Hypercube/Orthogonal-Array sampling in order to increase the probability that the low-dimensional parametric representation of cross-section variations discussed in Section 6 captures the variations that have the most effect on the outputs from the physics codes.

Flux weights are a kind of importance weighting for each energy group. They are specified by group (rather than, say, for each reaction) because they are computed based on the neutron populations in each group in a typical application. The 30 group flux weights that we use in this work were derived from the fission spectrum and the DT fusion peak of the traditional Los Alamos "TD" weight function. The $1/E$ and Maxwellian flux shapes at lower incident neutron energies were not included[5].

Given a set of weights for the energy groups as a vector w of length g, where g is the number of groups, we construct a diagonal matrix, W, of size $gb \times gb$, where b is the number of blocks (including the total block), by repeating w down the diagonal of W, once for each block. Then, we may compute a weighted cross-section set, t, from an ordinary set, s, as,

$$t = Ws. \tag{22}$$

We then compute the covariance matrix for the weighted cross sections, using Equations 7,

$$\Upsilon \equiv \mathrm{cov}(t) \equiv W\Sigma W^T = W\Sigma W, \tag{23}$$

where the last expression comes from the obvious symmetry of W.

We may draw samples for weighted cross sections by using Υ just as we would draw samples for unweighted cross sections by using Σ, using in this case the spectral decomposition of Υ rather than the spectral decomposition of Σ as described in Section 5 for normal sampling and Section 6 for LH/OA sampling. Doing so does not actually buy us anything if we use all of the eigenvectors of Υ, as the resulting samples need to be transformed back to unweighted cross-section sets in order to be used by the physics codes and we will just have samples with the same statistical properties as before (assuming that there are no zero

[5]The "TD" weight function is also sometimes referred to as the "CLAW" weight function. It is numerically specified as a 49 point log-log interpolation table and may be found as the `IWT=9` weight function of the `GROUPR` subroutine of the LANL neutron cross-section processing code NJOY [3].

weights, see below). However, if we want to do LH/OA sampling using a limited number of eigenvectors in order to have a parametric representation with a small number of parameters, as described in Section 6, using a small number of eigenvectors of Υ with the largest eigenvalues rather than a similar set of eigenvectors of Σ is likely to give us samples that do a better job of capturing the variations in cross sections that have the greatest effect on the outputs of the physics codes.

The situation is actually a bit more complex than this because there *are* some zero weights, for the first six energy groups (starting from the lowest energy group). In this case, we need to take care in inverting the weighting transformation to go from the sampled weighted cross sections to the unweighted cross sections used by the physics codes. Essentially, we simply impose the condition that there are no variations in the first six energy groups and all samples will simply have the mean value for these groups (in all blocks). We do this by removing the rows and columns of the zero-weighted groups from Υ (as well as the zero rows and columns that Υ inherits from Σ) before performing the spectral decomposition. We then proceed with sampling as usual using the eigenvectors and eigenvalues of Υ, transform back to unweighted samples, and then restore the zero components to the samples.

More precisely, our proceedure is:

(1) Compute the diagonal matrix of weights, W.
(2) Using the repaired covariance matrix from Equation 15, compute the weighted covariance matrix Υ from Equation 23.
(3) Remove the zero rows and columns of Υ (both the ones that are zero due to zero weights and the ones that were zero in the original Σ).
(4) Compute a reduced version of W by removing the same rows and columns that were removed from Υ.
(5) Compute the spectral decomposition of the resulting reduced Υ.
(6) Use the eigenvectors and eigenvalues of Υ to compute sample weighted cross-section sets according to Equation 21, where here the u_i and λ_i refer to the eigenvectors and eigenvalues of Υ, not Σ, and we ignore the mean vector μ for the moment.
(7) Transform the samples back to unweighted samples by multiplying them by the inverse of the reduced version of W, which exists because the zero elements on the diagonal of W were eliminated in doing the reduction.
(8) Restore the zero components to the samples.
(9) Add μ to each sample.

In this proceedure, the x_i in the sampling algorithm are chosen as independent samples from a standard univariate normal if we are doing normal sampling, or according to a LH/OA design with ranges from, say, -2 to 2 if we are doing LH/OA sampling over a 2-σ range.

When we perform LH/OA sampling using flux weights, we capture 98% of the variance in the weighted cross sections using just the first three eigenvectors of the weighted covariance matrix.

8. Appendices

8.1. **Quantifying Violation of the Summation Constraints.** In Section 4.2 we discussed how to impose the summation constraints on the mean vector, μ, and the covariance matrix, Σ, if the given experimental mean and covariance do not satisfy the constraints to sufficient precision, but we did not explain how to determine whether or not a violation of the constraints large enough to require "repair" of the mean or covariance is present, or even how to quantify such a violation. In this section, we explain how to quantify such violations of the constraints.

To begin with, note that the definition of the constraints in Equation 9 for an arbitrary cross-section vector, s, leads naturally to the definition,

$$\Delta_s \equiv C^T s, \tag{24}$$

where Δ_s is a vector of length g, where g is the number of energy groups, that quantifies, for each energy group, the degree to which the cross sections for the individual reactions fail to sum to the value in the total block[6].

In practice, it is probably more meaningful to give the discrepancy between the total of the reactions and the value in the total block as a fraction of the value in the total block, so we define a vector, δ_s, of length g by,

$$(\delta_s)_i \equiv (\Delta_s)_i \,/\, (s_t)_i\,, \quad i = \{1, \ldots, g\}, \tag{25}$$

where s_t is just the total block of s.

Quantifying violations of the summation constraints in the mean vector is now straightforward—we simply use the absolute quantity, Δ_μ,

$$\Delta_\mu \equiv C^T \mu, \tag{26}$$

of the relative quantity, δ_μ,

$$(\delta_\mu)_i \equiv (\Delta_\mu)_i \,/\, (\mu_t)_i\,, \quad i = \{1, \ldots, g\}, \tag{27}$$

where μ_t is just the total block of the mean vector.

[6]Note that the definition of Δ_s involves C and not $\widehat{C}$.

Quantifying the degree to which Σ violates the summation constraints is more complex because, unlike μ, Σ is not itself a cross-section vector. We proceed by considering the cross-section vectors that will produced by sampling from a multivariate normal distibution with mean 0 and covariance Σ, $\mathrm{N}(0, \Sigma)$ (where we are ignoring μ because we want to look just at the contribution of Σ to any violation of the constraints). If we define s as a multivariate random variable with this distribution, then we may take Equation 24 as defining a new random variable, Δ_s, by multiplying s by the matrix C^T. Then, according to Equation 7, the covariance of Δ_s is,

$$\mathrm{cov}(\Delta_s) = C^T \Sigma \left(C^T\right)^T = C^T \Sigma\, C, \tag{28}$$

(and, in fact, the distribution of Δ_s is $\mathrm{N}\big(0, C^T \Sigma\, C\big)$).

If we then consider the covariance matrix $C^T \Sigma\, C$ for Δ_s, we see that its diagonal elements of are precisely the variances of the elements of Δ_s. The square roots of these elements are the standard deviations of the elements of Δ_s, and we may reasonably take these as measures of the degree that Σ produces violations of the summation constraints in sample cross sections, Δ_Σ,

$$\Delta_\Sigma \equiv \sqrt{\mathrm{diag}\big(C^T \Sigma\, C\big)}. \tag{29}$$

Finally, as with violations of the constriants by the mean vector, it probably makes sense to look at the violations coming from Σ relative to the mean values in the total block, δ_Σ,

$$(\delta_\Sigma)_i \equiv (\Delta_\Sigma)_i \,/\, (\mu_t)_i\,, \quad i = \{1, \ldots, g\}. \tag{30}$$

Alternatively, we can take the components of Δ_Σ relative to the standard deviations of the groups in the total block, $\sqrt{\mathrm{diag}(\Sigma_b)}$, where Σ_b is the sub-matrix of Σ corresponding to the total block,

$$(d_\Sigma)_i \equiv (\Delta_\Sigma)_i \,/ \sqrt{\mathrm{diag}(\Sigma_b)_i}, \quad i = \{1, \ldots, g\}. \tag{31}$$

References

1. M. B. Chadwick et al., *ENDF/B-VII.1 Nuclear data for science and technology: Cross sections, covariances, fission product yields and decay data*, Nucl. Data Sheets **112** (2011), 2887.
2. Gene H. Golub and Charles F. Van Loan, *Matrix computations*, fourth ed., Johns Hopkins University Press, 2012.
3. R.E. MacFarlane and A.C. Kahler, *Methods for processing ENDF/B-VII with NJOY*, Nucl. Data Sheets **111** (2010), no. 12, 2739–2890.