XI.C. MISCELLANEOUS COMMENTS AND SUGGESTIONS

Bayes' method

Users should bear in mind that there is a fundamental difference between SAMMY and most other data analysis programs with which they might be familiar: SAMMY was designed to use Bayes' equations rather than least-squares equations to update parameter values. The approach taken to data analysis can be quite different from that taken with least squares. Some of these differences are indicated here.

SAMMY can be used sequentially* to obtain results valid over extremely large energy regions or over many different data sets. One need not do a series of independent analyses followed by much labor to provide consistency over all the data; consistency is a by-product of Bayes' equations. Parameters relevant to all data sets being analyzed should be flagged for fitting.

Parameters irrelevant to the data currently being analyzed (e.g., energies or widths for far-away resonances) may be varied (flagged for fitting) without ill effect. These parameter values will not be changed significantly, nor will their uncertainties. (This feature, in particular, is quite different from many implementations of the least-squares method, in which significant computation time is often consumed looking for the "best" value of parameters about which the current data have no information.)

During the initial stages of an analysis, when starting values are poorly known, ** SAMMY may be used much like least-squares programs are used, analyzing the same data with different starting values, varying different parameters, etc., until reasonable parameter values are found for the data being studied. The user should bear in mind that those values are accurate only under the assumption that all the fixed parameters remain fixed; hence, it is necessary to later do a global fitting of all parameters.

Occasionally a user may find it convenient to include a data set more than once during a sequential fit. However, in this case, the parameter covariance matrix (PCM) produced by the sequential analyses must not be considered to be accurate: This is equivalent to including the same data twice in a simultaneous fit, that is, to reducing the uncertainty on each data point by the square root of 2. Sometimes this can be used to advantage: For example, to obtain extremely accurate parameter values, one might run sequentially through all data sets 100 times. Upon completion, the final PCM must be multiplied by 100 (the uncertainties multiplied by 10) to give the actual PCM.

---

* A "sequential fit" is one in which the output PARameter and COVariance file from one run (analyzing one energy region or one of the data sets) are used as input to the next run (using another energy range or another data set), analyzing each and every data point once and only once.

** When starting values are far from true values, the assumption of linearity (implicit in both Bayes' equations and the least-squares equations) is not even approximately true. Hence the blind use of even the iterative form of either set of equations will often not produce a good fit of theory to data.

It is possible to use an output PARameter file as input to a future run without using the COVariance file. To do this, the user may either (i) edit the PARameter file, deleting the final line which states "COVariance matrix is in binary form" and using instead an initial (prior) large uncertainty on each parameter, (ii) include in the input file a card that says "`IGNORE INPUT BINARY covariance matrix file`", which will cause the default prior uncertainties to be used, or (iii) answer "`IGNORE`" when SAMMY requests the name of the binary file.

Prior (initial) uncertainties for each varied parameter are required as input to SAMMY (although they will be assigned default values if not stated explicitly). In the case where it may be assumed that we have no prior knowledge about the value of a certain parameter (i.e., usually), the value for the uncertainty should ideally be infinite. (Recall that Bayes' equations, in the limit of infinite prior uncertainty, reduce to the familiar least-squares equations.) Computers, however, require finite numbers. Moreover, it has not been possible to program SAMMY to operate via least squares for some parameters and via Bayes' equations for others. Compromise is therefore necessary. A convenient and practical compromise is that formulated in SAMMY's default values for prior uncertainties (see Section V1.B). As far as Bayes' equations are concerned, these default uncertainties are effectively infinite, yet they are usually not so large as to cause singularity problems in inverting the matrix.* A major exception is the default uncertainty on resonance energy, which may be much too large, especially for large s-wave resonances.

The ratio of the output uncertainty for a given parameter to its initial (prior) uncertainty is a measure of the effect the data have on that parameter. An unchanged or slightly changed uncertainty implies that the data have little or no effect on the parameter. A drastic reduction in the uncertainty (i.e., a ratio which is, say, less than 0.05 or 0.10) implies that this is an important parameter for these data, and, further, that the prior value was effectively infinite. An increase (a ratio > 1.0) perhaps indicates that numerical problems exist.

Occasionally there are significant changes in values of parameters which should be unaffected in the region being analyzed. This might be the case, for example, for the width of a large s-wave resonance. If, however, the uncertainty on that parameter remains large and the correlations between it and other parameters are also large, possibly the value of that parameter will resume its expected magnitude once data directly affecting it are analyzed. Values of correlated parameters will be altered to ensure that the fit to the data analyzed earlier is not degraded. On the other hand, it may happen that SAMMY is unable to recover from major changes in parameter values. In this case, the analyst should carefully consider whether some data are inconsistent with the others and make appropriate changes. If all data are deemed to be good, it will likely be necessary to alter the sequential order in which data are analyzed or to decrease prior uncertainties on parameters. Alternatively, one could do simultaneous runs, as described in Section IV.E.1.

* Experience has shown that the choice of prior uncertainties can, in fact, make some difference in the final results, both for parameter values and for parameter uncertainties. This has often been stated as a deficiency in Bayes' method, since the output depends upon the input. The analyst should keep in mind, however, that output from least-squares analyses also depends upon the input, but with least squares the user has no choices regarding that particular piece of input – the prior uncertainties are always infinite.

Spin group definitions

It is legitimate to specify a spin group in the INPut file without giving any resonance for that spin group in the PARameter file. Potential scattering (via hard-sphere phase shifts) will be calculated for that spin group. It is not permitted, however, to request a spin group or channel in the PARameter file without specifying that group in the INPut file.∗

The ENDF so-called "Reich Moore" format (LRF=3), however, does not allow spin groups to be defined without at least one resonance in the spin group. (See the next paragraph for the exception to this rule.) Evaluators who require an "empty" spin group are encouraged to use the new, less restrictive "R Matrix Limited" format (LRF=7).

The original ENDF LRF=3 format made the implicit but unwarranted assumption that all resonances having a given l and J would have the same s. Because this assumption is not always valid, the format was modified in 1999 to specify which channel spin would be used; see Section VI.F.2 for details. The format also requires that "the other" channel spin be added if it is absent, in order to obtain the correct hard-sphere-phase-shift cross section.

A critic has suggested that SAMMY should *require* that all analyses always use all possible spin groups and channels (up to a user-determined value of l, presumably); the critic is recalling an early ENDF evaluation for which "the other channel spin" was omitted during the analysis, causing great consternation when processor codes reconstructed a cross section which differed from the original. Nevertheless, the SAMMY author has chosen not to hamstring the SAMMY user in this fashion, relying instead on the expertise of today's evaluators to avoid repetition of a single 25-year-old mistake.

There are legitimate reasons for giving the analyst the ability to define spin groups as s/he chooses: (1) Only the analyst understands all aspects of the evaluation, so only he or she should be responsible for including all relevant information. (2) Forcing the inclusion of spin groups or channels with negligible contribution will add nothing of value to the calculation, and only increase the CPU time used for calculating zero. (3) It is often useful for the analyst to consider the effect of an individual spin group without including the others.

To be sure that your final evaluation includes the "complete" set of spin groups and channels (and to justify the SAMMY author's faith in the thoroughness of today's evaluators), users are strongly encouraged to make use of program SAMQUA for assistance in determining spin groups.

∗ When the command "`QUANTUM NUMBERS ARE in parameter file`" is given and the spin-group definitions are at the beginning of the PARameter file rather than in the INPut file, the same arguments hold. A channel or spin group cannot be used in the resonance-parameter list if it is not first defined in the spin-group list.

Negative reduced-width amplitudes

In Section II.B.1, in the paragraph following Eq. (II B1.7), it is stated that reduced-width amplitudes can be either positive or negative. However, the reduced-width amplitude itself is not printed in either SAMMY or ENDF files. Instead, the partial width (related to the square of the reduced-width amplitude) is printed; if the value in the file is given as negative, the sign is assumed to belong to the amplitude rather than to the partial width. (This potentially confusing convention is standard practice, is unambiguous, and soon becomes second-nature to experienced R-matrix practitioners.)

In light of the previous paragraph, one might ask why so few negative signs are observed in SAMMY or ENDF files. There are several reasons:

- The "natural" thing to do is to use positive signs until and unless the need for a negative sign becomes apparent. Because it is the easy thing to do, we assign all the amplitudes as positive unless there is evidence to the contrary.

- The sign is irrelevant for widely spaced resonances of the same spin groups (and for resonances of different spin groups). Even for neighboring resonances of the same spin group, only the relative sign is sometimes observable.

- Interferences between neutron widths are often effectively invisible under experimental conditions, lost in the noise and the Doppler and the resolution broadening, etc.

- Negative signs are most often observed for fission widths, when two or more fission widths are specified for each resonance. These signs appear naturally as needed during the course of an analysis – but again, they only appear when needed, and they are not needed very often.

- Rather than starting with all positive amplitudes, a more realistic approach would be to assign the signs randomly for the starting parameters. This is rarely done, most often with fission widths.

- A basic assumption underlying the Reich-Moore approximation is that the signs of the gamma-width amplitudes are randomly distributed. This randomness plus the sheer number of capture channels allow us to treat the gamma width as an average, non-interfering, positive quantity.

Transmission vs. total cross section

Analysts will want to make use of transmission data rather than total cross section data whenever possible. Total cross section data are typically derived from transmission data, which cannot reliably be translated to cross section data for all resonances for any given sample thickness. Ideally, transmission data from several different-thickness samples should be included in the evaluation, in order to determine the peaks for resonances of widely differing widths.

ENDF files

The official manual for ENDF formats is [ENDF-102]. A subset of that manual relevant to the resolved-resonance region, with additional comments of interest to SAMMY users, is given in an ENDF-for-SAMMY-users guide [NL07b].

Anyone creating or using ENDF files for the resolved resonance region should be aware of idiosyncrasies in the ENDF formats.

- The ENDF multilevel Breit-Wigner format (LRF = 2) is "multi" only for elastic scattering. Other cross sections are single level.

- The so-called Reich-Moore format (LRF = 3) permits only a limited subset of options available in Reich-Moore. Use the new "R-Matrix Limited format" (LRF = 7) for situations that do not fit the LRF = 3 format. That is, use LRF = 7 for more than one entrance channel, for more than two fission channels, for other reaction channels (inelastic, proton, alpha, etc.), for combinations of orbital angular momenta and/or channel spins, and for other generalizations.

- The so-called Reich-Moore format (LRF = 3) requires that, for a given l and J, any non-identified (i.e., missing) channel spin s must be added to the configuration as a new spin group. With this format, it is not possible to explicitly define a spin group or channel for which there are no resonances. Hence, missing channels are implicitly assumed to be present so that the hard-sphere phase shift contribution will be added to the cross sections.

- In contrast, the R-Matrix Limited format (LRF = 7) requires that each and every spin group and channel be explicitly defined. It is the evaluator's responsibility to be sure that all non-negligible spin groups and channels are included.

- Except for the new LRF = 7 format, no ENDF formats have provision for defining parity explicitly. Hence, for example, the ENDF file for ^{235}U suggests that the ground state is $7/2^+$ when it is in fact $7/2^-$.

- The ENDF formats, like the SAMMY PARameter file format, use $G = A\Gamma = A\left(2P\gamma^2\right)$, where A is the sign of γ, for the input quantity.

- Unlike SAMMY, ENDF resonance parameter covariance matrices (RPCM) are written in terms of the input quantity G rather than of the more physical quantity Γ. (As of this writing, this topic is not addressed in the ENDF-102 manual, though it is expected to be included in future releases of that document. The convention that ENDF RPCM be written in terms of G is used in SAMMY to create ENDF File 32, is used in the two processor codes ERRORJ [GC04] and AMPX [MD02], and has been recommended to CSEWG. In contrast, the RPCM printed in the SAMMY.LPT file has always been expressed in terms of Γ rather than G.)