How to handle calibration uncertainties in high-energy Astrophysics

Vinay L. Kashyap[a], Hyunsook Lee[a], Aneta Siemiginowska[a], Jonathan McDowell[a], Arnold Rots[a], Jeremy Drake[a], Pete Ratzlaff[a], Andreas Zezas[a], Rima Izem[b], Alanna Connors[c], David van Dyk[d], and Taeyoung Park[e]

[a]Harvard-Smithsonian Center for Astrophysics, Cambridge, MA, USA;
[b]Harvard University, Cambridge, MA, USA;
[c]Eureka Scientific, Arlington, MA, USA;
[d]Univ. of California, Irvine, CA, USA;
[e]Univ. of Pittsburgh, Pittsburgh, PA, USA

ABSTRACT

Unlike statistical errors, whose importance has been well established in astronomical applications, uncertainties in instrument calibration are generally ignored. Despite wide recognition that uncertainties in calibration can cause large systematic errors, robust and principled methods to account for them have not been developed, and consequently there is no mechanism by which they can be incorporated into standard astronomical data analysis. Here we present a framework where they can be encoded such that they can be brought within the scope of analysis. We describe this framework, which is based on a modified MCMC algorithm, and propose a format standard derived from experience with effective area measurements of the ACIS-S detector on *Chandra* that can be applied to any instrument or method of codifying systematic errors. Calibration uncertainties can then be propagated into model parameter estimates to produce error bars that include systematic error information.

Keywords: Calibration, Systematic Error, Astrophysics, X-ray, Chandra, ACIS-S, Effective Area, Format

1. INTRODUCTION

The importance of accounting for statistical errors has been well established in astronomical analysis: a measurement is of little value without an estimate of its credible range. Numerous robust strategies have been developed to compute model uncertainties due to imperfect data. A major component of this analysis is good knowledge of the instrument characteristics, described by the instrument calibration data. Indeed, without the transformation from measurement signals to physically interesting units afforded by the instrument calibration, the observational results cannot be understood in a meaningful way. However, even though it is well known that the measurements of the instrument's properties (e.g., quantum efficiency of a CCD detector, point spread function of a telescope, etc.) have associated measurement uncertainties, the calibration of instruments is taken on faith, with only nominal estimates used in data analysis, even when it is recognized that these uncertainties can cause large systematic errors in the inferred model parameters. Calibration uncertainty is usually ignored entirely, or in some cases, it is assumed that the calibration error is uniform across an energy band or an image area. Such a process can often cause an erroneous interpretation of the data.

The reason that calibration uncertainties have been often ignored is that a robust and principled method to account for them had not existed. Every user or builder of instruments will strongly agree that incorporating

Send correspondence to VLK; E-mail: vkashyap@cfa.harvard.edu, Telephone: 1 617 495 7173

This paper was published in *Observatory Operations: Strategies, Processes, and Systems*, Edited by Roger J. Brissenden & David R. Silva, Proceedings of SPIE Vol. 7016, pp. 70160P-70160P-8 (2008) and is made available as an electronic reprint with permission of SPIE.

calibration uncertainties is important. But work is hindered by having no common language or standard procedure for doing this. In the past, individual groups have tried instrument-specific methods. These have ranged from bootstrapping,[1] to raising and lowering response "wings" by hand,[2,3] and in one case, even analytical marginalization over a particular kind of instrumental uncertainty.[4] However, in general, especially in important cross-instrument comparisons, all but the crudest methods (say, multiplying each instrument's total effective area by a fitted "uncertainty factor"[5]) are very difficult to handle.

In an earlier paper,[6] we described a "brute force" Monte Carlo method of including calibration uncertainties in typical parameter estimation studies. In that work, many realizations of an instrument response were generated to sample the range in the uncertainties of its constituent calibration components. These realizations were used in repeated model fitting exercises to determine their influence on retrieved model parameters. Here, we extend that work and develop a more general framework where calibration uncertainties can be codified in such a manner that they can be brought within the scope of regular data analysis. In §2, we describe this framework, which is based on a modified Markov-Chain Monte Carlo model fitting algorithm. The primary hindrance to universal use of this framework however lies in the manner in which calibration uncertainties are codified in practice. In §3, we describe how the knowledge of instrument uncertainties can be distilled into a convenient and generalizable form. In §4, we distill this description into a format that can be generalized to different instruments, different calibration products, and even different ways of codifying the uncertainty.

2. INCORPORATING SYSTEMATIC ERROR IN ANALYSIS

In general, the response of a detector to incident photons arriving at time t can be written as

$$M(E', \mathbf{p}', t) = \int dE\, d\mathbf{p}\; S(E, \mathbf{p}, t; \theta)\; \mathcal{R}(E, E', \mathbf{p}'; t)\; \mathcal{P}(\mathbf{p}, \mathbf{p}', E; t)\; \mathcal{A}(E, \mathbf{p}'; \mathbf{p}, t) \tag{1}$$

where $\mathbf{p}'$ and E' are the measured photon location energy (or the detector channel as the case may be), while $\mathbf{p}$ and E are the true photon energy and sky location; the source physical model $S(E, \mathbf{p}, t; \theta)$ describes the energy spectrum, morphology (point, extended, diffuse, etc.), and variability with parameters θ; and $M(E', \mathbf{p}', t)$ are the expected counts in detector channel space. Calibration is carried out using well known instances of $S(E, \mathbf{p}, t; \theta)$ to determine the quantities

$$\begin{aligned} \mathcal{R}(E, E', \mathbf{p}'; t) &\equiv \text{Energy Redistribution} \\ \mathcal{P}(\mathbf{p}, \mathbf{p}', E; t) &\equiv \text{Point Spread Function} \\ \mathcal{A}(E, \mathbf{p}'; \mathbf{p}, t) &\equiv \text{Effective Area} \end{aligned} \tag{2}$$

It is important to note that all of the quantities in Equation 1 have uncertainties associated with them.

2.1 Types of Calibration Uncertainty

The term "calibration uncertainty" covers three distinct types of errors:

1. Statistical errors: Calibration products are derived by comparing data obtained in strictly controlled conditions with predictions from well-defined sources, either in the lab, or some special well-understood astrophysical sources. That is, the calibration products $\mathcal{R}(.)$, $\mathcal{P}(.)$, $\mathcal{A}(.)$ are estimated via Equation 1 using a well-defined source model $S(.)$. Parameterized models are fit to these data to derive best-fit parameters that are then used to define the relevant calibration products, e.g., the gain or quantum efficiency of a detector, the transmission efficiency of a grating, etc. The errors on the best-fit values of these parameterized models, which carry the information on how accurately the calibration is known, are however discarded during subsequent definitions of the calibration products. If recorded, they can be used to generate Monte Carlo simulations that include these errors. Alternatively, the source model may be modified by multiplying it with a generic function (such as a spline or a polynomial model) whose parameters are also fit; the generic function is increased in complexity until the quality of the fit is deemed adequate, and the errors on the parameters of the generic model is taken to characterize the magnitude of the calibration error[7]

2. Correlated uncertainties or systematic errors: The calibration at one point for one instrument may depend on the measurements at a different point or for a different instrument. Thus, changing the calibration values at one place will have a ripple effect that changes the likely calibration everywhere. For example, the *Chandra*/LETGS+HRC-S effective area is calibrated by using the power-law source PKS 2155-304. Because of the contributions from higher orders to the spectrum, which cannot be disentangled, the index of the power-law used is predicated on an analysis of the same source with data obtained contemporaneously with the HETGS+ACIS-S. Thus, changes in the HETGS+ACIS-S effective area will affect the longer-wavelength LETGS+HRC-S effective area. Similarly, the uncertainties in each subassembly component, such as the telescope mirror effective area, the grating transmission coefficients, the transmittance of the UV/Ion-shield, etc., will have a direct and wide-ranging effect on the measured response of the detector to incoming photons. The result is a set of possible effective area curves which are all equally applicable, but where using one curve rather than the other may cause significant systematic differences in derived model parameters. A detailed accounting of the uncertainties present in *Chandra*/ACIS-S effective areas is given by Drake et al.,[6] who also describe a natural way to track the variations: they generate numerous instances (N_{ARF}) of possible realizations of the ACIS-S effective areas, which together form an ensemble

$$\{\mathcal{A}_i(E_j), i = 1..N_{ARF}, j = 1..N_{bin}\} \tag{3}$$

where each simulation i of the effective area curve is written as a function of the energy E_j.

3. Subjective expertise: The most difficult aspect of calibration uncertainties, one that cannot be easily quantified, is the knowledge encompassed by the experience of the scientist. Often, the existence of systematic uncertainty is suspected, but is not quantified until suitable data are acquired or cross-instrument comparisons are made.[8] Proper calibration of complex instruments, such as those on modern space-based telescopes, is in some sense an art, and it is beyond the capacity or the charge of most astronomers who use the data at one remove from the calibration effort. Different mechanisms have been proposed to quantify this type of uncertainty, ranging from adopting ad hoc distributions such as truncated Gaussians[6] or uniform deviations over a specified range. But as long as it can be characterized even loosely, statistical theory does provide a mechanism by which such information can be included in the analysis. In a Bayesian framework, knowledge of the instrument uncertainties are described by a prior probability. This is multiplied by the probability of the data, given the instrument characteristics (and astrophysical model parameters) to form a posterior probability – or, in practice, a joint likelihood of both. See, e.g., the manner in which known atomic data constraints are included in the reconstruction of a coronal emission measure.[9] Thus, once quantified, in any manner seen fit by calibration scientists, such information can be incorporated trivially in the MCMC based method described below. Furthermore, the MCMC technique can also be used to firm up the loose estimates of uncertainty by subjecting the *a priori* distribution to rigorous inference. Given high quality data, this process allows them to be updated.

2.2 Markov-Chain Monte Carlo

Previously, the best available strategy for incorporating calibration uncertainties had been to square-add the measurement errors to estimates of the calibration errors and to then fit the source model to the data. This is not an optimal strategy for a number of reasons: it makes unwarranted assumptions that both the measurement error and the calibration errors are Gaussian; that the calibration errors are uncorrelated; and that the uncertainty on the calibration products can be uniquely translated to an uncertainty in each bin in data space.

A better solution was proposed by Drake et al.,[6] who carried out model fits separately for each instance of a simulated effective area and then estimated the effect of the systematic error via the variance in the maximum likelihood estimates in the model parameters. This solution produces reasonable estimates of the parameter variances, and can be relatively easily implemented even within existing analysis software packages such as XSPEC and Sherpa. However, it too has certain limitations, ranging from the theoretical to the practical. Theoretically, it does not have a robust Bayesian foundation and the heuristic can be considered ad hoc. The relative magnitudes of the statistical and systematic components of the total variance is as yet unsettled. In practice, it also requires verification whenever different models are considered or even different parts of the parameter space are explored. Furthermore, the large number of fits required imposes a heavy computational

cost. Most importantly, this strategy requires that either instrument specific code be available in order to construct new simulations of the calibration product, or a comprehensive database of possible realizations of the product be made available *a priori*. The former imposes a heavy burden on both the software maintainers and the calibration groups. The latter is unrealistic due to excessive storage costs. Our focus here is to describe a solution that works around all of these problems.

Our preferred solution is to adopt a principled computational strategy that is robust, reliable, and fast. We therefore adopt a Markov-Chain Monte Carlo (MCMC) algorithm[10] which, with simple modifications, can handle any level of complexity present in both the astrophysical models and in the calibration uncertainty. Simply put, a Markov chain is an ordered sequence of parameter values such that any particular value in the sequence depends on the history of the sequence only through its immediate predecessor. Markov chains can be used to construct powerful Monte Carlo sampling schemes that are able to explore interesting regions in high-dimensional parameter spaces and, for instance, determine best-fit values of model parameters. In this context, it is used as a fitting engine, similar to, but far more robust than, Levenberg-Marquardt, Powell, Simplex, and other minimization algorithms. One of its main advantages is that it is highly flexible and can be applied to a wide variety of problems. A single run of the MCMC algorithm is sufficient to fully describe the variations in the model parameters that arise due to both statistical and systematic errors, which therefore leads to reduced computational costs.

Our strategy is simply stated: calibration uncertainties can be directly incorporated within the general MCMC algorithm in a modular fashion, by selecting and using a different realization of the calibration values at each iteration of the MCMC sampler. How these realizations are obtained will vary on a case by case basis.

A Monte Carlo sample is obtained by sampling the model parameters θ given the data D and calibration information Z,

$$\theta^{(k)} \sim p(\theta^{(k)}|D, Z),$$

where k is the iteration number, $\theta^{(k)}$ are the values of the parameters at iteration k, $p(.)$ is a probability distribution, and the notation "|" refers to conditioning, i.e., $p(A|B)$ is to be read as "the probability of A *given* that B is true." When calibration uncertainty is included, we can no longer condition on Z as a known value of the calibration values. Instead we add a new step that updates Z according to the calibration uncertainties. In particular, $\theta^{(k)}$ is updated using the same iterative algorithm as above, with an additional step at each iteration that updates Z. Suppose at iteration k, $Z^{(k)}$ is the realization of calibration values. Then the new algorithm consists of the following two steps:

$$\begin{array}{lll} \theta^{(k)} & \text{are sampled from} & p(\theta^{(k)}|Z^{(k-1)}, D), \text{ and} \\ Z^{(k)} & \text{are sampled from} & p(Z^{(k)}|\theta^{(k)}, D)\,. \end{array} \tag{4}$$

We expect $Z^{(k)}$ to be nearly independent of both $\theta^{(k)}$ and D, so that the prior distribution $p(Z)$ is the dominant contributor to the chain of $Z^{(k)}$. We emphasize the modular structure of this algorithm. It effectively separates the complex problem of model fitting in the presence of calibration uncertainties into two simpler problems: fitting a model with known calibration and the quantification of calibration uncertainties.

The main issue before us is how to specify $p(Z)$ so that valid calibration values are used at each iteration. A simple strategy is to specify a library of possible calibration values that cover the total range of possibilities (this may impose a heavy storage cost; see below for alternatives). Some may be more likely than others, and weights that indicate the relative likelihood are required. With this library in hand, one of the possible curves would be sampled at each iteration of the MCMC sampler. Bayes' theorem would be used to update the prior weights to the posterior weights that are contingent on D and θ. (Again we expect the difference between the prior and posterior weights to be small.) Below, we describe a specific instance of a practical strategy that aims at summarizing a continuum of potential calibration values in a mathematically concise form. Note also that this strategy is not limited to handling uncertainties in instrument response alone; it can also be used to include uncertainty in atomic data, such as the emissivities of emission lines, in the analysis of high-resolution spectra.[9]

3. PRACTICAL CONSIDERATIONS

As pointed out above (§2.2), one of the primary issues that needs to be addressed is how to store the knowledge of the uncertainty in the calibration products in such a way that neither software costs nor storage costs are prohibitive. We first note that the best source for establishing the magnitude of the calibration uncertainties are the instrument calibration teams, who are indeed knowledgeable about and are best suited to judge the quality of the calibration. In principle, the calibration teams of each instrument, or each telescope, could write and make public appropriate code that would simulate a new effective area, response matrix, or point spread function (Equation 2) and which could be incorporated within Sherpa and XSPEC. However, it is extremely unrealistic to expect the instrument teams to write software that is compatible with multiple analysis systems, and to maintain it such that it remains compatible with future evolutions of these systems. The alternative, for the calibration teams to generate thousands of independent simulations of the calibration products once and for all, and for the analysis programs to sample from this database, is also unrealistic because of the extremely large storage space that is required to handle every possible observational case.

The solution to the above conundrum lies in the realization that the majority of the simulations differ very little from each other. In other words, even though the size of the database is large, the information contained within it can be considered sparse. Thus, if we reduce the dimensionality of the database from thousands of identical stored products to only a few that capture most of the variation present in the calibration, then a database of calibration uncertainty products becomes feasible. Below, we describe one implementation of such a dimensionality reduction, applied to *Chandra*/ACIS-S effective areas, using Principal Components analysis.

3.1 PC Decomposition

We describe the process by which we reduce the dimensionality of calibration uncertainties by using the specific example of *Chandra*/ACIS-S effective areas. It is possible to generate[6] simulations of the total effective area of this instrument by explicitly including uncertainties in each of the subsystems (UV/Ion shield transmittance, CCD Quantum Efficiency, and the telescope mirror reflectivity) that comprise it. Each realization represents a version of the combined instrument effective area that differs from the reference effective area by the uncertainties in the calibrations of subassembly instrument components. The ensemble of simulated effective areas can be represented as $\{\mathcal{A}_{ij}\}$ (see Equation 3). The dimensionality of this ensemble can be drastically reduced by carrying out a Principal Components analysis.[6]

Principal Component Analysis (PCA), also called Karhunen-Loeve expansion, or Proper Orthogonal Decomposition, and related to Factor Analysis, is a standard exploratory statistical tool for analyzing variations in multivariate data. PCA finds the linear, orthogonal, directions of variation of greatest variability in the data. It is often used to reduce the dimensionality of the data since the number of these linear directions is usually smaller than the original dimension of the data. PCA has a simple linear algebra derivation – principal component directions are the orthogonal eigenvectors of the variance-covariance matrix of the data. The standardized eigenvalues associated with the eigenvectors give the percentage of the variation explained by each direction. When the multivariate data are a set of curves, an extension of PCA is used where the variance covariance function is decomposed in an orthonormal space of eigenfunctions.

Once a set of effective areas $\{\mathcal{A}_{ij}\}$ are generated such that they encompass the range of uncertainty present in our knowledge of the effective area, the differences with respect to a reference effective area, $\mathcal{A}_0$, that is usually provided to the users,

$$\delta\mathcal{A}_{ij} = \mathcal{A}_{ij} - \mathcal{A}_{0j} \tag{5}$$

are calculated. These carry the full uncertainty information. The $\{\delta\mathcal{A}_{ij}\}$ are then decomposed via a principal components analysis to generate eigenvalues $\{e_i\}$ and eigenvectors $\{\nu_i(E_j)\}$. The fraction of the variance in $\{\delta\mathcal{A}_i\}$ accounted for by the k^{th} component,

$$f_k = \frac{e_k^2}{\sum_{i=1}^{N_{ARF}} e_i^2}\,, \tag{6}$$

and they are ordered such that $e_1 \geq e_2 \geq ... \geq e_{N_{ARF}}$.

The PCA step reduces dimensionality because a small subset of the eigenvectors can account for almost all the variation in $\delta\mathcal{A}_{ij}$. These components, together with the mean sample deviation

$$\overline{\delta\mathcal{A}_j} = (1/N_{ARF}) \sum_i \delta\mathcal{A}_{ij} \,, \tag{7}$$

which is the average deviation from the nominal effective area, describes most of the variation in the original sample. A new Monte Carlo realization of the effective area, $\mathcal{A}'_j$, which shares all the properties of the original sample, can then be obtained as

$$\mathcal{A}'_j \;=\; \mathcal{A}_{0j} + \overline{\delta\mathcal{A}_j} + \sum_{k=1}^{N_{comp}} \; r_k \; e_k \; \nu_{jk} \;\; + r_{(N_{comp}+1)} \; \xi_j \,, \tag{8}$$

where $r_k \sim \mathcal{N}(0,1)$ are Gaussian deviates, and N_{comp} is formally equal to the number of effective areas in the ensemble N_{ARF}, but can be reduced as needed to discard components that are ignorable, and the discarded components can be accounted for as a single residual component,

$$\xi_j = \sum_{k=N_{comp}+1}^{N_{ARF}} \; e_k \; \nu_{jk} \,. \tag{9}$$

The average deviation, $\overline{\delta\mathcal{A}_j}$ can be thought of as the bias term, and is expected to be close to 0. In the above example, we have $N_{ARF} = 1000$ (i.e., 1000 separate curves, each representing a possible realization of the *Chandra*/ACIS-S effective area), but this set can be reduced to $N_{comp} \approx 20$ orthogonal sets of principal components which account for $> 99\%$ of the variation.

3.2 Codifying the Uncertainty

The reconstruction of the simulated effective areas based on PC decomposition, as described in Equation 8, can be rephrased to be in the more general form

$$\text{Simulated} = \text{Nominal} + \text{Bias} + \text{randomized components} + \text{residuals} \,. \tag{10}$$

This formulation removes the necessity of depending solely on PCA methods, and allows us to use a variety of methods to generate the simulated curves. For example, it can even include such loosely stated measures of uncertainty as "the effective area is uncertain by X% at wavelength Y".

Furthermore, note that this form is not limited to describing effective areas alone, but can also be used to encompass the calibration uncertainty in response matrices and point spread functions. The precise method by which the randomized components are generated will vary widely, but in all foreseeable cases they can be described as in Equation 10, with a bias term and a randomized component added to the nominal product, and with an optional residual component. The ensemble of calibration products simulated in this way form an informative prior $p(Z)$ (see Equation 4). Some potential methods of describing the randomized component are listed below:

– When a large library of simulated calibration curves is available, the randomized component is simply the full set of curves from the simulation. During the MCMC iterations, a random index is chosen and the curve corresponding to that index is picked for that iteration. This process preserves the weights of the initial simulation. Note that in this case the residual component is identically zero.

– The description of a 2-dimensional calibration product such as the spectral response matrix cannot be done with a single PC decomposition. However, it is possible to define a sequence of hierarchical PCs, where first the "within variances" are characterized (i.e., how the response varies inside a single response matrix) and then the "between variances" are determined (i.e., PCA is carried out on the Principal Components of the "within variance"). This reduces the required storage by many orders of magnitude.

– If the calibration uncertainties are characterized via a multiplicative polynomial term in the source model, the randomized component in Equation 10 can be obtained via a Monte Carlo bootstrap of the parameters of the

polynomial, starting from their best-fit values and the estimated errors. These simulated curves can then be used to modify the nominal effective areas inside each iteration. Thus, the bias and residual terms are identically zero, and only the polynomial parameter best-fit values and errors need to be stored.
– Often, it turns out that the calibration is identified to be systematically offset over a small passband, and that this had not been previously identified. In such cases, users can specify their own estimate of calibration uncertainty as a randomized offset over the relevant range. This is essentially equivalent to using a correction, as above, with a polynomial of order 1. The stored quantities here are the average offset, the bounds over which the offset can range, and a pointer specifying whether to generate uniform or Gaussian deviates over that range.

4. GENERALIZED FILE FORMAT

Here we describe our proposed file specification, with some key header values that are based on the Principal Components decomposition of ACIS-S effective areas. We are working on ways to make files written in this format be accessible transparently within Sherpa.

In general, each calibration product, whether an effective area, response matrix, or a point spread function, is stored in a FITS file with at least the following 3 blocks:

PRIMARY – with an empty data section

BIAS – containing the nominal calibration product as well as the bias term
(i) Specifically in our example with the effective areas, the BIAS block is reminiscent of the SPECRESP block present in the HEASARC ARF standard (CAL/GEN 92-002).
(ii) In addition to the columns represending the energy grid (ENERG_LO and ENERG_HI), we also have a column SPECRESP containing the nominal effective area $\mathcal{A}_{0j}$, and a column BIAS that contains $\overline{\delta\mathcal{A}_j}$.
(iii) The BIAS column is optional, so that if it is missing, the bias is assumed to be identically zero.
(iv) Additional columns, such as the equivalent wavelength grid (BIN_LO and BIN_HI) may also be present.

UNCERTAINTY – containing the parameters or arrays necessary to generate a randomized component, as well as the residual term
(i) For the PCA decomposition of Equation 8, this contains 4 columns, COMPONENT (containing the component number k), FVARIANCE (containing the fraction of the variance accounted for by that component, f_k), EIGENVAL (the eigenvalue e_k), and EIGENVEC (the eigenvector ν_{jk}, stored in a single cell as a variable length array).
(ii) There are $N_{comp}+1$ rows in this extension. The last row contains the residuals ξ_j in the EIGENVEC column, and is further identified by a component number $k=-1$ and an eigenvalue $e_{-1}=0$.

It is easy to generalize the same file structure to account for RMFs and PSFs. The files require a number of keywords in addition to those that are part of the HEASARC specification. In the BIAS extension, we must have the header keywords

EMETHOD Specifies how the randomized component is generated. In our example, it is set to PCA1D. Other possible choices are SIMS (stored simulations), PCPC (within and between PCA), POLY1D (multiplicative polynomial term), etc.

HDUCLASS We recommend setting this to "CXC", in order to distinguish it from the HEASARC standard.

HDUCLAS1 Case dependent; for ARFs, set to RESPONSE

HDUCLAS2 Case dependent; for ARFs, set to SPECRESP

HDUCLAS3 Always set to BIAS, to indicate that this extension contains the BIAS term.

HDUNAME Always set to BIAS.

ERREXT Set to the name of the extension that contains the uncertainty values.

and in the UNCERTAINTY extension, we must have

EXTNAME Specifies the extension name, and is required to be the same as ERREXT from the BIAS extension

EMETHOD Identical to that in the BIAS extension, it is repeated here for purposes of clarity.

HDUCLASS Same as that set in the BIAS block

HDUCLAS1 Set to CALERR

HDUCLAS2 Case dependent; for ARFs, set to PCACOMP

HDUCLAS3 Unused, but currently set to EMETHOD.

5. SUMMARY

Our goal is to obtain realistic error bars on astrophysical source model parameters that include both statistical and systematic errors. For this purpose, we have developed a general and comprehensive strategy to describe and store calibration uncertainty and then to incorporate them within data analysis via a Markov-Chain Monte Carlo scheme. This scheme treats the possible variations in calibration as an informative prior while estimating the posterior probability distributions of the source model parameters. Thus, the effects of calibration uncertainty is automatically described in the result of a single fit.

In general, any analysis scheme that includes systematic errors must do so via a large set of simulated calibration products. There are two strategies that can be used to produce such simulations – a software based one or a storage based one. Building and maintaining independent software code that can carry this out for numerous detectors on various telescopes for observations at different epochs is difficult, and further requiring that these diverse codes be accessible in standard analysis packages such as Sherpa or XSPEC is unrealistic. We therefore propose a storage method that compresses the uncertainty information and allows it to be stored along with the nominal calibration products. The compression is defined such that it preserves the statistical properties of the calibration uncertainty, such as the magnitude of the possible variations and the correlations that exist therein.

We propose a specific file format to contain the calibration uncertainty information, built in the context of describing the uncertainties in *Chandra*/ACIS-S on-axis effective area. Work is under way to have this file format be accessible in Sherpa.

ACKNOWLEDGMENTS

We thank Herman Marshall, Randall Smith, and Nancy Brickhouse for useful discussions. This research was supported by NASA/AISRP grant NNG06GF17G and NASA contract NAS8-39073 to the Chandra X-Ray Center.

REFERENCES

[1] Simpson & Mayer-Hasselwander, 1986, A&A, 162, 340
[2] Forrest, D.J., 1988, BAAS, v20, p740
[3] Forrest, D.J., Vestrand, W.T., & McConnell, M., 1997, Technical Report, New Hampshire Univ. Durham, NH
[4] Bridle, S.L., et al., 2001, MNRAS, 335, 1193
[5] Hanlon, L.O., et al., 1995, Ap&SS, 231, 157
[6] Drake, J.J., Ratzlaff, P., Kashyap, V., Edgar, R., Izem, R., Jerius, D., Siemiginowska, A., & Vikhlinin, A., 2006, Proc. SPIE, v6270, p627019
[7] Marshall, H., 2006, at IACHEC, Lake Arrowhead, CA
[8] David, L., Nevalainen, J., Bonamente, M., Jerius, D., Zhao, P., Edgar, R., Gaetz, T., Marshall, H., Grant, K., & Bautz, M., 2007, *Chandra Calibration Workshop*, #2007.23
[9] Kang, H., van Dyk, D.A., Kashyap, V.L., Connors, A., 2005, *X-Ray Diagnostics of Astrophysical Plasmas: Theory, Experiment, and Observation*, AIP Conf. Proc. Volume 774, pp. 373
[10] van Dyk, D., Connors, A., Kashyap, V.L., & Siemiginowska, A. 2001, ApJ 548, 224