An Application of Semiparametric Bayesian Isotonic Regression to the Study of Radiation Effects in Spaceborne Microelectronics

Marian Farah

MRC Biostatistics Unit, University of Cambridge, UK

Athanasios Kottas

University of California, Santa Cruz, USA

Robin D. Morris

Palo Alto, California, USA

Summary. This work is concerned with the vulnerability of spaceborne microelectronics to single event upset, a change of state caused by high-energy charged particles in the solar wind or the cosmic ray environment striking a sensitive node. To measure the susceptibility of a semiconductor device to single event upsets, testing is conducted by exposing it to high-energy heavy ions or protons produced in a particle accelerator. The number of upsets is characterized by the interaction cross-section, an increasing function of linear energy transfer (LET). The prediction of the on-orbit upset rate is made by combining the device geometry and cross-section vs. LET curve with a model for the orbit-specific radiation environment. We develop a semiparametric isotonic regression method for the upset count responses, based on a Dirichlet process prior for the cross-section curve. The proposed methodology allows the data to drive the shape of the cross-section vs. LET relationship, resulting in more robust predictive inference for the on-orbit upset rate than conventional models based on Weibull or lognormal parametric forms for the cross-section curve. We illustrate the modelling approach with data from two particle accelerator experiments.

Keywords: Cross-section vs. LET curve; Dirichlet process prior; Markov chain Monte Carlo; On-orbit upset rate

Address for correspondence: Athanasios Kottas, Department of Applied Mathematics and Statistics, School of Engineering, University of California, Santa Cruz, CA 95064, U.S.A. E-mail: thanos@ams.ucsc.edu

1. Introduction

Single Event Upsets (SEUs) are non-destructive errors caused by the interaction of a charged particle with a circuit, resulting in a change of logic state (Dodd and Massengill, 2003). For example, SEUs appear as bitflips in a storage element, and as transient voltage pulses in arithmetic logic units. Characterizing the vulnerability of a spaceborne electronic device to SEU is critical for the success of space missions. The fundamental quantity of interest is the upset rate of the device for the space radiation environment specified by a particular orbit.

It is known that the upset rate depends on the linear energy transfer (LET) of the incident particles, the cross-section of interaction, and the fluence (Petersen et al., 1992). LET is a measure of the energy transferred to a device per unit length when an ionizing particle passes through it, and the common unit for LET is $MeV \cdot cm^2/mg$; fluence, given in $particles/cm^2$, is the number of particles that intersect a unit area; cross-section is a measure of the device upset response to ionizing radiation – it is calculated as the number of errors divided by fluence (Xapsos et al., 1993; Petersen, 1996).

To measure the susceptibility of an electronic device to SEUs, the device is subjected to heavy ion tests in a particle accelerator. The results of such tests include the fluence and the number of upsets recorded at each combination of ion and energy (and thus LET).

The prediction of on-orbit upset rates from test data typically proceeds in two stages. The first stage of the analysis is to fit the cross-section vs. LET curve to the experimental data. In the second stage of the analysis, the estimated cross-section vs. LET curve is combined with the rectangular parallepiped approximation to the geometry of the sensitive volume and a model of the space radiation environment to produce predictions of on-orbit upset rates (Tylka et al., 1997). A number of codes have been developed to compute on-orbit upset rates; in this work, we use CREME96 (Cosmic Ray Effects on Micro-Electronics), a widely-used code for modelling radiation environments in near-Earth orbits to evaluate radiation effects in spacecraft (Tylka et al., 1997). We note that some recent work uses explicit Monte Carlo simulation of the space environment and of the interaction of the charged particles with the device (Warren et al., 2005; Weller et al., 2009). However, this approach requires detailed knowledge of the device geometry, which is often not available. We do not pursue this approach in this paper.

Our focus is on flexible predictive inference for the cross-section vs. LET curve, which is taken to be an increasing function of LET, ℓ, with a plateau related to the physical cross-section of the device. Upsets are caused when the charge deposited in the device's sensitive volume exceeds a critical amount. Charge deposition is proportional to energy deposition, so the cross-section vs. LET curve is monotonically increasing. In particular, the cross-section vs. LET curve can be generically written as $\sigma_0 G(\ell)$, where σ_0 is the limiting cross-section, and $G(\cdot)$ is a cumulative distribution function (cdf) on $\mathbb{R}^+$.

Standard practice is to assume $G(\cdot)$ is the cdf of a Weibull distribution (e.g., Petersen et al., 1992; Pickel, 1996; Swift et al., 2008). The Weibull parameters are, typically, estimated using (weighted) least squares, thus avoiding probabilistic modelling of the response distribution for the upset counts. The use of the Weibull parametric form for the cross-section vs. LET curve is conventional, with little physical justification to support it. Other mod-

els (e.g., lognormal) can be used (Petersen, 1996), and, as we demonstrate in this paper, the resulting model uncertainty may have a substantial impact on the accuracy and/or uncertainty of the on-orbit upset rate predictions.

We address uncertainty in the cross-section vs. LET curve by treating the entire function $G(\cdot)$ as the key unknown parameter of the model. Specifically, we work with a Dirichlet process prior (Ferguson, 1973; Antoniak, 1974) for $G(\cdot)$, which, along with a parametric prior for σ_0, defines the prior probability model for the mean of the response distribution, which is taken to be Poisson. Hence, from a methodological point of view, we develop a Bayesian semiparametric isotonic regression model for count responses.

We apply the method to data obtained from two particle accelerator experiments. One of the devices tested was a commercial microprocessor, and the other was a field-programmable gate array (FPGA), a semiconductor device that can be configured by the designer after manufacturing. To illustrate the utility of the semiparametric approach relative to existing techniques, we compare it with the standard Weibull model, as well as with a lognormal model. The Dirichlet process based semiparametric model allows the data to drive the shape of the cross-section vs. LET curve, and can thus result in predictive inference that is less sensitive to modelling choices than parametric models. This can further result in more reliable inference for the upset rate distribution. Quantifying the uncertainty in the predicted on-orbit upset rates can give valuable information regarding upset mitigation needed when the part is used in space, and also when sufficient testing has been performed.

The outline of the paper is as follows. In Section 2, we provide a description of the data sets from the two devices. Section 3 develops the semiparametric model for the cross-section vs. LET curve, including methods for prior specification, posterior inference, and model comparison with customary parametric approaches. Section 4 reports on the data analysis results for the two devices. In Section 5, we extend the model to account for uncertainty in the fluence values recorded, and Section 6 concludes with a discussion. The appendices include technical details on implementation of inference under the semiparametric model.

2. The Data

The data used to illustrate the proposed methodology are obtained from two particle accelerator experiments. The first dataset corresponds to tests performed on a Freescale MC7447AT PowerPC (subsequently "Device A"). It records the errors in the data cache bits. In these tests the number of distinct LET values is small (only 4 LET values are used), but there are repeated measurements at each LET value. The second dataset corresponds to tests conducted for a rare failure mode on a Xilinx XQR4VLX200 FPGA (subsequently "Device B"). The failure mode considered is termed "POR SEFI" (Power-On Reset Single Event Functional Interrupt). It is triggered by an upset to certain elements of the FPGA's control logic and is a global error, as opposed to the single-bit errors considered for Device A. In this dataset, the number of upsets reported at any given LET is small. Due to the sparsity of the data, we expect significant statistical uncertainty in the estimated cross-section vs. LET response.

The data, D, for each device are obtained in the form $D = \{(c_i, f_i, \ell_i) : i = 1, \ldots, N\}$,

where c_i is the upset count, f_i is the fluence, and ℓ_i is the LET value. The two data sets are given in Tables 1 and 2. Figure 1 shows a plot of the cross-section vs. LET from the experimental data for both devices. For device A, the data are collected for the whole device, but the fit is performed on a per-bit scale (σ_0 is multiplied by the number of bits when performing the fit). Thus, the observed cross-section per bit values are computed as $c_i/(kf_i)$, where k is the number of storage bits in the device. For device B the failure mode is per-device, so the observed cross-section values are computed as c_i/f_i. Device B has unique measurements recorded at 7 distinct LET values, while device A has repeated measurements at 4 distinct LET values.

3. Methods

The semiparametric modelling approach is presented in Section 3.1. In Section 3.2, we outline the method for posterior inference. In Section 3.3, we give the specification of the prior distributions used, and model comparison is addressed in Section 3.4. Technical details are provided in the appendices.

3.1. Bayesian semiparametric model

We model the upset counts, c_i, assuming a Poisson distribution, Poisson$(c_i; \mu_i)$, with mean which, as a function of LET, ℓ_i, is given by

$$\mu_i = f_i \sigma_0 G(\ell_i). \tag{1}$$

Hence, the mean of the Poisson response distribution is specified by the cross-section vs. LET curve, $\sigma_0 G(\cdot)$, adjusted by the observed fluences, f_i. In Section 5, we develop a model that allows for measurement error in the fluence values, which yields as a special case the Poisson likelihood specification in (1).

As discussed in the Introduction, our key modelling objective is to avoid potentially restrictive parametric assumptions regarding the form of cdf $G(\cdot)$, which specifies the shape of the cross-section vs. LET curve. We thus employ a Dirichlet process (DP) prior for $G(\cdot)$ with centering (base) distribution G_0 and precision parameter α. Here, $G_0(\ell) \equiv G_0(\ell; \boldsymbol{\psi})$ is a parametric cdf on $\mathbb{R}^+$ with parameters $\boldsymbol{\psi}$. Moreover, α controls the variability of DP realizations around G_0; as α gets larger, cdf realizations from the DP prior get *closer* to the centering cdf $G_0(\ell)$. We write DP(α, G_0) to denote the DP prior for $G(\cdot)$.

The semiparametric approach to modelling $G(\cdot)$ is more flexible than standard parametric techniques, since it obviates the need to specify a particular functional form for the cross-section vs. LET curve, which may not be supported by the data. It thus enables quantification of uncertainty due to the functional form of the random cdf $G(\cdot)$, which is essential for accurate estimation of the upset rate distribution. At the same time, an appealing feature of the DP prior model is that it can be centered around familiar parametric models. In our context, the commonly used Weibull parametric form is a natural choice for G_0, resulting in

$$G_0(\ell; w, s) = 1 - \exp\left(-(\ell/w)^s\right). \tag{2}$$

Alternatively, the lognormal cdf can be used, giving

$$G_0(\ell; m, w) = \Phi((\ln(\ell) - m)/w), \tag{3}$$

where $\Phi(\cdot)$ is the standard normal cdf. In the former case, we have $\boldsymbol{\psi} = (w, s)$ whereas, in the latter, $\boldsymbol{\psi} = (m, w)$.

Denote by $\theta_i = G(\ell_i)$, $i = 1, ..., N$, the parameters defining the cdf at the observed LET values, $\ell_1 < \ell_2 < ... < \ell_N$. Then, the DP$(\alpha, G_0)$ prior for $G(\cdot)$ implies an ordered Dirichlet prior distribution for $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$. Specifically,

$$p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi}) = \frac{\Gamma(\alpha)}{\prod_{i=1}^{N+1} \Gamma(d_i)} \theta_1^{d_1 - 1} (\theta_2 - \theta_1)^{d_2 - 1} \ldots (\theta_N - \theta_{N-1})^{d_N - 1} (1 - \theta_N)^{d_{N+1} - 1}, \tag{4}$$

where $d_1 \equiv d_1(\alpha, \boldsymbol{\psi}) = \alpha G_0(\ell_1; \boldsymbol{\psi})$, $d_i \equiv d_i(\alpha, \boldsymbol{\psi}) = \alpha(G_0(\ell_i; \boldsymbol{\psi}) - G_0(\ell_{i-1}; \boldsymbol{\psi}))$, for $i = 2, \ldots, N$, and $d_{N+1} \equiv d_{N+1}(\alpha, \boldsymbol{\psi}) = \alpha(1 - G_0(\ell_N; \boldsymbol{\psi}))$. The prior model for $G(\cdot)$ is completed with hyperpriors for α and for $\boldsymbol{\psi}$. Prior specification for the DP hyperparameters as well as for σ_0 is discussed in Section 3.3.

Regarding the first stage distribution of the model, in the case of device B, the data set includes a single count c_i for each distinct LET value ℓ_i, $i = 1, \ldots, N$. Hence, based on the conditionally independent Poisson responses with means in (1), the likelihood is given by

$$\mathcal{L}(\boldsymbol{\theta}, \sigma_0; D) \propto \sigma_0^{\sum_{i=1}^{N} c_i} \exp\left(-\sigma_0 \sum\nolimits_{i=1}^{N} f_i \, \theta_i\right) \prod_{i=1}^{N} \theta_i^{c_i}. \tag{5}$$

In the data set corresponding to device A, we have repeated measurements, with fluences and counts (f_{ij}, c_{ij}), $i = 1, ..., n_j$, corresponding to distinct LET value ℓ_j, $j = 1, ..., p$ (with $N = \sum_{j=1}^{p} n_j$). In this case, the first stage distribution can be written as

$$\mathcal{L}(\boldsymbol{\theta}, \sigma_0; D) \propto \sigma_0^{\sum_{j=1}^{p} \sum_{i=1}^{n_j} c_{ij}} \exp\left(-\sigma_0 \sum\nolimits_{j=1}^{p} \theta_j \sum\nolimits_{i=1}^{n_j} f_{ij}\right) \prod_{j=1}^{p} \prod_{i=1}^{n_j} \theta_j^{c_{ij}}. \tag{6}$$

Note that the number of DP induced parameters $\theta_j = G(\ell_j)$, $j = 1, ..., p$, is smaller than N in this case.

We note that the model developed here can be applied to more general isotonic regression problems, where the DP prior will provide flexibility for the shape of the isotonic regression function. For instance, the Poisson response distribution could be replaced with a negative Binomial distribution, and, in principle, more general mixture distributions for count responses can be considered. A different Bayesian isotonic regression method for count responses was developed in Dunson (2005), where the focus was on a prior that assigns positive probability to flat regions for the regression function. The use of the DP prior for isotonic regression functions with continuous responses was explored in Lavine and Mockus (1995). The work on bioassay modelling with DP priors for the dose-response curve (e.g., Gelfand and Kuo, 1991; Mukhopadhyay, 2000; Kottas et al., 2002) is also related to the approach developed here.

3.2. Posterior predictive inference

The joint posterior distribution, $p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D)$, is obtained by combining the likelihood in either (5) or (6), for the data from devices B or A, respectively, with the prior for $\boldsymbol{\theta}$ in (4) and the priors for σ_0, α, and $\boldsymbol{\psi}$. We use a hybrid Markov chain Monte Carlo (MCMC) algorithm to sample from this distribution. The MCMC method comprises Metropolis-Hastings steps for α and $\boldsymbol{\psi}$ and slice sampling steps for the components of $\boldsymbol{\theta}$; see Appendix A for details.

The MCMC algorithm provides samples of the posterior distribution for the cross-section vs. LET curve at the observed LET values. Full inference for the curve requires the posterior distribution of $G(\ell_k)$ over a sufficiently fine grid of LET values ℓ_k. In particular, extrapolating the random cdf $G(\cdot)$ beyond the largest and smallest observed LET values is necessary for prediction of the upset rate distribution. Such inference can also reveal important differences in the prediction of upset rates between customary parametric models and the proposed semiparametric alternative. Appendix B provides the details on predictive inference for $G(\cdot)$.

Once the cross-section vs. LET curve is available, it can be used to predict the SEU rate for the device when used in a particular orbit (Pickel, 1996). A heavy-ion incident on a device will cause the deposition of charge within the device. This charge is generated by ionization as the ion loses energy. The charge is collected in certain areas of the device, due to the electric fields present. If the charge collected on a sensitive node is sufficient, it will cause an upset.

The sensitive volume is typically modeled as a rectangular parallelepiped (RPP). The $x-$ and $y-$ dimensions are given by the square-root of the limiting cross-section (σ_0), and the $z-$ dimension is termed the thickness. Here, z is not the device's physical thickness, but rather the depth of the region in which deposited charge is collected. The charge deposited by an ion traversing the RPP is proportional to the ion's LET multiplied by the path length within the sensitive volume. The flux of ions in space is omnidirectional, so the distribution of charge deposited depends on the flux of heavy ions and the distribution of chord lengths within the sensitive volume.

Assuming a step threshold for producing an upset results in an upset rate as a function of energy $R(E)$. The cross-section vs. LET determined from heavy ion tests is then used to weight $R(E)$, giving the rate prediction as $R = \int R(E) f(E) \mathrm{d}E$, where $f(E)$ is the normalized density function corresponding to the cross-section vs. LET curve. Clearly, the accurate determination of $G(\cdot)$, and hence $f(\cdot)$, is vital for accurate rate prediction. The computation of predicted rates as outlined here is performed by the CREME96 software, available through the web interface at `https://creme-mc.isde.vanderbilt.edu/`.

Results on posterior inference for on-orbit upset rates are presented in Section 5.2 under the practically important extension of the model to incorporate uncertainty in the recorded fluence values.

3.3. Prior Specification

To complete the model developed in Section 3.1, we must specify priors on: α, the precision parameter of the DP; σ_0, the limiting cross-section; and $\boldsymbol{\psi}$, the parameters of the DP

centering distribution G_0.

For the precision parameter we considered two different priors, a gamma prior centered around 25 with moderate variance (see Figure 2, right panel), and a uniform prior on $(0, 100)$. (To avoid numerical instabilities in the MCMC algorithm implementation, the gamma prior is truncated at 100.) The gamma prior favors moderate to large deviations from G_0 corresponding to moderate to small α values, respectively; note that α values in $(50, 100)$ yield prior realizations for G that get increasingly closer to G_0.

The prior information for σ_0 is different for the two devices. Device A is manufactured using a radiation-hard-by-design methodology (Lacoe et al., 2000), and the test LET values are known to cover the range that includes most or all of the rising portion of the cross-section vs. LET curve. When the data covers the plateau, or comes close to the plateau, the limiting cross-section, σ_0, is typically well-determined by the data. Hence, for device A, we placed a diffuse exponential prior on σ_0. Device B is a radiation-tolerant device, and so it is unknown a-priori whether the range of LET values used in the tests reaches the plateau – the data may only cover the left tail of the distribution associated with the cross-section vs. LET curve. This results in weak identifiability between σ_0 and the scale parameter of G_0, which can be resolved by using an informative prior for σ_0 chosen to be uniform up to a maximum value $\sigma_{\max}$. An absolute upper-bound for σ_0 is the physical size of the device. This can be reduced, in consultation with a device expert, by considering the proportion of the area of the device that is occupied by active circuit elements, and the fraction of the active circuit area that can be considered to be sensitive to ion strikes. For device B we based the choice of $\sigma_{\max}$ on the values given in Swift et al. (2008). We note that a device expert could have been consulted to give an informative upper bound for σ_0 for device A also, but, as explained above, this was not necessary.

Finally, we used independent exponential priors for the components of parameter vector $\boldsymbol{\psi}$, where $\boldsymbol{\psi} = (w, s)$ when G_0 is Weibull, and $\boldsymbol{\psi} = (m, w)$ when G_0 is lognormal (see expressions (2) and (3), respectively). The parameters of the exponential priors were determined by assuming a fairly noninformative range of LET values (twice the observed range), which we use as an estimate of the difference between the 0.975 and 0.025 percentiles of G_0, with the midrange value used as a rough estimate of the median of G_0. This results in two equations with two unknowns, which are the elements of $\boldsymbol{\psi}$. The solutions are taken to be the means of the exponential priors for w and s, or for m and w. Note that this approach requires only a plausible range of LET values, which, in general, can be obtained from a device expert prior to conducting the experiment.

3.4. Model comparison

For each data set we study two versions of the DP-based semiparametric model (using the Weibull and lognormal centering distributions), and for each model we have two different priors (uniform and gamma) for the DP precision parameter.

We also consider the parametric Weibull and lognormal fits to the data, using the same priors for σ_0 and $\boldsymbol{\psi}$ specified earlier. The posterior distribution, $p(\sigma_0, \boldsymbol{\psi} \mid D)$ ($\boldsymbol{\psi} = (w, s)$ or (m, w)) under these parametric models can be sampled using straightforward Metropolis-Hastings algorithms. Predictive inference under the two parametric models is similarly

straightforward.

In Section 4, we perform empirical comparison of the two commonly used parametric models with the more general semiparametric model. To supplement graphical comparison with more formal model comparison, we use a cross-validation posterior predictive criterion based on the log-pseudo marginal likelihood (LPML) statistic (Geisser and Eddy, 1979).

In general, let $\mathcal{M}_1$ and $\mathcal{M}_2$ be two models which are to be compared based on available data $D = \{y_j : j = 1, ..., N\}$, and denote by $p_r(\cdot \mid D)$ the posterior predictive density under model $\mathcal{M}_r$, $r = 1, 2$. The conditional predictive ordinate (CPO) for observation y_j under model $\mathcal{M}_r$ is given by $\text{CPO}_{r,j} = p_r(y_j \mid D_{(j)})$, where $D_{(j)}$ denotes the $(N-1)$-dimensional data vector resulting from D with y_j removed. The ratio $\text{CPO}_{1,j}/\text{CPO}_{2,j}$ describes how well model $\mathcal{M}_1$ supports the observation y_j relative to model $\mathcal{M}_2$, based on the remaining data $D_{(j)}$. The LPML statistic under model $\mathcal{M}_r$ is given by $N^{-1}\sum_{j=1}^{N}\log\{p_r(y_j \mid D_{(j)})\}$, and it provides an aggregate "leave-one-out" cross-validation measure of model predictive utility.

In our context, the LPML statistic offers a natural choice for model comparison, since good predictive performance of a postulated model for the cross-section vs. LET curve is essential for reliable inference of the upset rate distribution. In particular, given the availability of repeated measurements at each LET value for the data from device A, we also consider "block" cross-validation. Under this approach, the LPML statistic is computed from $\text{CPO}^*_{ij} = p(c_{ij} \mid (f_{ij}, \ell_j), D^*_{(j)})$, where $D^*_{(j)}$ is the reduced data vector resulting from removing the entire block of observations $\{c_{ij} : i = 1, ..., n_j\}$ for each LET value ℓ_j. Arguably, the "block" cross-validation strategy provides a more general test for the models under consideration, since it emphasizes predictive performance at new (unobserved) values over the LET space.

Given the small sample sizes for our data, computing for the LPML statistic is feasible using directly the definition of the CPOs. Alternatively, CPO computing is possible using the MCMC output from the fit of the model to the full data vector. The relevant technical details for the semiparametric model can be found in Appendix C. The latter approach is more efficient, but requires careful implementation to avoid numerical instabilities.

The predictive performance of the parametric and semiparametric models is assessed in Section 4 using the estimated CPO values and the corresponding LPML statistics.

4. Results

In Sections 4.1 and 4.2 we present inference results for cross-section vs. LET responses for the two devices under the semiparametric model, discuss the sensitivity of the responses to our modelling assumptions, and study how the semiparametric model compares with the parametric models.

4.1. Device A

Figure 2 shows the DP fit of the cross-section vs. LET curve under the Weibull centering distribution for both the gamma and uniform priors for α, together with the prior and posterior densities for α. The posterior mean cross-section is very similar for the two

priors for α, however, the 95% intervals differ – substantially when extrapolating above the highest observed LET. We also note that the 95% intervals become wider between the observations. There is a small amount of learning for α, with the posterior distribution under the uniform prior favoring larger values (corresponding to less divergence from the centering distribution). This is consistent with the extrapolation results. The cross-section results were found to be insensitive to the priors on σ_0 and $\boldsymbol{\psi}$ $(= (w, s))$.

Figure 3 (left panel) shows the parametric Weibull and lognormal fits. It can be seen clearly that the choice of parametric function can have a large impact on the estimated cross-section vs. LET curve. We note also that the parametric models underestimate the cross-section uncertainty between the observed LETs. The semiparametric model is expected to be much less sensitive to the choice of centering distribution. Figure 4 (left panel) shows the estimated cross-section vs. LET for the semiparametric model based on the lognormal centering distribution. While the posterior mean estimate for the cross-section vs. LET curve is rising slowly for large LET values, the difference between the two semiparametric models is much smaller than for the parametric models. Learning for α is again relatively weak, although in this case, under both priors for α, the posterior densities (figure 4, right panel) favor smaller values. This suggests that the data provide less support to the lognormal parametric model than the Weibull model.

These observations are confirmed by the LPML values given in Table 3. Under each of the "leave-one-out" or the "block" cross-validation approaches (discussed in Section 3.4), the LPML values for all the semiparametric models are very similar, demonstrating the insensitivity with respect to the centering distribution. Under both cross-validation settings, the semiparametric modelling approach outperforms the parametric models. In particular, its predictive power is highlighted by the "block" cross-validation LPML results, where the lognormal parametric model fares significantly worse, as would be expected considering the results in the left panel of Figure 3.

4.2. Device B

Figure 5 shows the semiparametric fits based on the Weibull and lognormal centering distribution for the two prior specifications for α. The inference for the cross-section curve is very similar for the two α priors, and the two centering distributions; the only noticeable difference is the somewhat larger posterior uncertainty in the extrapolation region under the Weibull choice for G_0. In general, inference is dominated by the sparsity of the data – the small number of counts at each LET value results in wide posterior probability intervals. As with the device A data, the posterior densities for α (not shown) depicted a relatively small amount of learning for the DP precision parameter.

Figure 3 (right panel) shows the Weibull and lognormal parametric estimates of the cross-section vs. LET curve. There is again a larger difference between the two parametric estimates than the semiparametric models centered on the same parametric forms; however, that difference is less emphatic than in the results for the device A data. The robustness of the DP-based semiparametric predictive inference is confirmed by the LPML values in Table 4. We note, however, that for this data set the parametric models yield slightly larger LPML values, indicating that adequate prediction can be obtained from a simpler model,

provided the parametric form for the cross-section curve is appropriate for the data.

5. Modelling Extension for Random Fluences

The model considered so far has assumed that the fluences, f_i, are measured without error. However, the dosimetry in heavy-ion accelerator facilities is not perfect. The traditional assumption by practitioners was that up to 10% error in the dosimetry was to be expected (e.g., Ladbury et al., 2004). A number of recent papers have shown, however, that this may in fact be over-conservative. For instance, Morris and Foster (2010) analyse data sets which show between 10% and 25% error in the fluence, and Petersen (2010) finds fluence errors between 7.3% and 93%. On the other hand, Morris and Foster (2011) analyse data for which the fluence is especially well controlled, and find an error of only 3%. It is useful therefore to expand the statistical model of Section 3 to allow for random fluences. Section 5.1 presents such an extension, and in Section 5.2, we discuss the corresponding results for the data from the two devices.

5.1. The modelling approach

We focus on the setting that involves a single count c_i for each distinct LET value ℓ_i (as for device B), and denote by f_i^*, $i = 1, \ldots, N$, the corresponding unobserved (random) fluences, which now become part of the parameter vector for the full model. The approach is similar for the data set from device A, which includes repeated counts for each LET value and thus additional parameters f_{ij}^*, $i = 1, ..., n_j$, $j = 1, ..., p$, for the random fluences.

The extension of the modelling framework of Section 3.1 requires a joint likelihood specification for the (f_i, c_i), $i = 1, ..., N$. We retain the Poisson distribution for the upset counts, c_i, and add an inverse gamma distribution for the observed fluences, f_i, with parameters for both distributions that depend on the f_i^*. Specifically,

$$(f_i, c_i) \mid f_i^*, \sigma_0, G(\cdot), \phi \overset{ind.}{\sim} \text{inv-gamma}(f_i; \phi, \phi f_i^*)\, \text{Poisson}(c_i; f_i^* \sigma_0 G(\ell_i)), \quad i = 1, ..., N \quad (7)$$

where inv-gamma(a, b) denotes the inverse gamma distribution with mean $b/(a-1)$ (provided $a > 1$). The Poisson specification for the upset counts is similar to the one in Section 3.1, where now the means in (1) are specified through the random fluences. The choice of the inverse gamma distribution for the observed fluences, along with gamma priors for the random fluences, facilitates MCMC posterior simulation. Moreover, under the particular parametrization for the distribution of the f_i, given f_i^* and ϕ, we obtain the fixed fluence model of Section 3.1 as a special limiting case. Note that $\text{E}(f_i \mid f_i^*, \phi) = f_i^*/(1 - \phi^{-1})$ and $\text{Var}(f_i \mid f_i^*, \phi) = \phi^2 (f_i^*)^2 / \{(\phi - 1)^2(\phi - 2)\}$, and thus, as $\phi \to \infty$, the distribution for f_i reduces to a point mass at $f_i = f_i^*$, for each $i = 1, ..., N$.

The part of the model that involves the cross-section vs. LET curve remains unchanged, that is, we keep the same priors as before for σ_0, $G(\cdot)$, and the DP parameters. Although the first-stage inverse gamma distribution for fluence measurement uncertainty does not depend on LET, we assign (independent) gamma(a_i^*, b_i^*) priors to the f_i^* with parameters specific to the LET value ℓ_i. Based on the discussion above regarding the range of values

of fluence error observed in typical data sets, the prior for the fluence should be sufficiently diffuse to support both moderate and large dosimetry errors. For the results reported in Section 5.2, we specified the priors for the f_i^* using the observed fluences, but, in general, eliciting from the device expert a rough LET-specific range of plausible fluence values would suffice for the analysis. Finally, parameter ϕ controls the dispersion of the inverse gamma distribution for the observed fluences. For a generally applicable prior choice for ϕ, it is useful to note that the joint model for (f_i, c_i) and f_i^*, implied by (7) and the gamma(a_i^*, b_i^*) prior for f_i^*, can also be written proportional to the Poisson$(c_i; f_i^* \sigma_0 G(\ell_i))$ model with a gamma$(\phi + a_i^*, b_i^* + \phi f_i^{-1})$ distribution for f_i^*. If we further consider the limiting version of the gamma(a_i^*, b_i^*) prior (with $a_i^* \to 0^+$ and $b_i^* \to 0^+$), we obtain a version of the standard hierarchical Poisson model, that is, the negative Binomial response distribution for the c_i when the f_i^* are integrated out. This connection with hierarchical Poisson generalized linear models facilitates the prior choice for the precision parameter ϕ; in particular, for the results of Section 5.2, we used the prior $p(\phi) = (1+\phi)^{-2}$, $\phi > 0$ (as in, e.g., Albert, 1988).

Under the above formulation for the DP-based model with random fluences, the joint posterior distribution for σ_0, $(\boldsymbol{\theta}, \alpha, \boldsymbol{\psi})$, ϕ, and $\{f_i^* : i = 1, ..., N\}$ is proportional to

$$\sigma_0^{\sum_{i=1}^N c_i} \prod_{i=1}^N \left\{ (f_i^* \theta_i)^{c_i} \exp\left(-\sigma_0 f_i^* \theta_i\right) \frac{(\phi f_i^*)^\phi}{\Gamma(\phi)} f_i^{-(\phi+1)} \exp\left(-\frac{\phi f_i^*}{f_i}\right) \right\}$$
$$\times \left(\prod_{i=1}^N \text{gamma}(f_i^*; a_i^*, b_i^*) \right) \times p(\sigma_0) p(\phi) p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi}) p(\alpha) p(\boldsymbol{\psi}),$$

where $p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})$ is the DP induced prior for the $\theta_i = G(\ell_i)$, $i = 1, ..., N$, given in (4), and $p(\phi)$, $p(\sigma_0)$, $p(\alpha)$, $p(\boldsymbol{\psi})$ denote the prior densities for the respective parameters.

The approach to MCMC posterior simulation follows closely the corresponding MCMC method for the model with fixed fluences (given in Appendix A). Updating α and $\boldsymbol{\psi}$ proceeds exactly as before, and updating σ_0 and the θ_i, given the f_i^*, involves replacing f_i with f_i^*. The additional sampling that is required is for the f_i^*, $i = 1, ..., N$, and for ϕ. The latter parameter is sampled with a Metropolis-Hastings step; the former parameters can be sampled directly, since the posterior full conditional of each f_i^* is given by a gamma distribution with shape parameter $a_i^* + c_i + \phi$, and rate parameter $b_i^* + \sigma_0 \theta_i + \phi f_i^{-1}$. Finally, posterior predictive inference for the cross-section vs. LET curve is again based on interpolation (and extrapolation) for the random cdf $G(\cdot)$, which proceeds exactly as detailed in Appendix B.

5.2. Results

For the data illustrations, we consider only the gamma prior for α (predictive inference for the cross-section function was similar under the uniform prior). Regarding the gamma priors for the random fluence parameters, we used the observed fluence values to specify their effective range. In particular, for device A, for each LET value ℓ_j, $j = 1, ..., 4$, the parameters of the prior for the f_{ij}^*, $i = 1, ..., n_j$, were chosen to yield a gamma distribution with most of its probability mass in $(0.25 \min\{f_{ij}\}, 2 \max\{f_{ij}\})$. For device B, the gamma prior for f_i^*, corresponding to LET value ℓ_i, had most of its mass in $(0.25 f_i, 2.5 f_i)$, $i = 1, ..., 7$. Again, we

also report results under the parametric Weibull and lognormal models which were extended to incorporate random fluences in the same fashion as with the DP semiparametric model.

Considering first device A, Figure 6 (left panel) includes the inference results for the cross-section vs. LET curve under the semiparametric modelling approach. The model is relatively invariant to the choice of the centering distribution, the main difference being in the uncertainty bands between the second and third LET values and in the extrapolation region. Posterior uncertainty is overall greater than the fixed fluence case, and there is also a smaller amount of reduction in the uncertainty of the cross-section at the LET values at which data was recorded (contrast with the left panel of Figures 2 and 4). Posterior uncertainty has increased more significantly under the parametric models (Figure 6, right panel), for which the main observation is that, as in the fixed fluence case, the parametric lognormal estimate is substantially different from the parametric Weibull estimate. The posterior densities for the random fluences were, in general, similar under the two different centering distributions for the DP-based model. For all f_{ij}^*, the posterior density was significantly concentrated relative to the corresponding prior density, thus suggesting a fair amount of learning for the random fluence parameters. The majority of the observed f_{ij} were near the mode of the posterior densities for the corresponding f_{ij}^*, but a few were well into the tails. The left panel of Figure 8 shows, under the DP model with the Weibull centering distribution, prior and posterior densities for a representative subset of the f_{ij}^*; the order of the observed fluences within each LET value follows from Table 1.

Turning to results for device B, due to the sparsity of the data, the difference between the semiparametric and parametric estimates is smaller than for the device A data. This is illustrated in Figure 7, which compares the estimated cross-section vs. LET for the semiparametric model with the Weibull or lognormal centering distributions with the corresponding parametric models. Under the Weibull centering distribution for the DP-based model, Figure 8 (right panel) plots prior and posterior densities for the random fluences, all of which include the observed fluences. Compared to the results for device A, there was a smaller amount of prior to posterior learning for the f_i^*, although for all seven LET values the posterior density was less dispersed than the prior density.

Regarding the precision parameter ϕ, inference under the DP semiparametric model was very robust to the choice of the centering distribution. Under the Weibull choice for G_0, posterior mean and 90% interval estimates were given by 49.29 and $(24.15, 81.80)$ for device A, and by 25.78 and $(3.36, 76.43)$ for device B. Note that under the $(1+\phi)^{-2}$ prior for ϕ, essentially all the prior probability mass is placed in $(0, 25)$. Hence, there was substantial prior to posterior learning for ϕ, which suggests that, even with sparse data, the magnitude of the uncertainty in the fluences can be reasonably well determined.

Finally, we report results on predictive inference for on-orbit upset rates, following the approach discussed in Section 3.2. We consider a nominal low-Earth orbit (450km altitude, 51.6 deg inclination, solar minimum, 100 mils aluminum shielding). In this orbit protons will also contribute to the upset rate; here, we consider only the contribution due to heavy-ion radiation. Under each of the models entertained in the earlier analyses, and for each of the two devices, Figure 9 plots the posterior densities for on-orbit upset rates. For device A, the semiparametric models give similar on-orbit upset distributions, with the posterior density

based on the Weibull centering distribution having thicker right tail. The distributions from the parametric models are both centered at larger values of on-orbit upset rate. However, for device B, as expected from the inference results for the cross-section curve, there is closer agreement in the estimated on-orbit upset rates between all the models.

6. Discussion

We have presented a Bayesian approach to modelling the cross-section vs. LET response of microelectronic parts that are candidates for use in space. The methodology was based on a semiparametric isotonic regression model for count responses, using a Dirichlet process prior for the monotonically increasing regression function. We have expanded the modelling approach in a hierarchical fashion to include uncertainty in the fluence measurements, and showed the effect of this on the predicted on-orbit upset rate distributions.

Prediction of the on-orbit upset rate for the microelectronic device is the key engineering problem that motivated this work. To improve upon standard practice in the parts testing literature, we have focused on flexible statistical modelling for the cross-section vs. LET responses. Physical modelling aspects have been limited to the monotonicity restriction for the cross-section curve as well as device-specific information to guide prior choice for the limiting cross-section parameter. As a consequence, a possible criticism for the approach revolves around the posterior realizations for the cross-section vs. LET curve, which may be physically implausible. In particular, the estimated responses for device A show inflexions at the observed LET values (left panel of Figure 2 and 4) although their extent is reduced under the more general model with random fluence (Figure 6, left panel). This aspect of the posterior inference for the cross-section curve is an artifact of the DP prior, exacerbated by the repeated measurements at each LET value. We note however that widening posterior uncertainty bands between the observations is a general characteristic of nonparametric priors and not specific to the DP prior. An improved model for the cross-section curve, more strongly based on the physics of upsets, may help alleviate this issue.

The feature of the proposed semiparametric model for the cross-section curve is that it is less sensitive to modelling assumptions than parametric methods. In particular, when one of the commonly used parametric models is sufficient for the data, the nonparametric prior will result in similar inferences, typically, at the expense of (somewhat) larger uncertainty (compare the two panels in Figure 7). However, in general, the semiparametric model yields more robust predictive inference, especially, for extrapolation (contrast the two panels in Figure 6), including prediction for the scientifically relevant on-orbit upset rate.

Acknowledgements

This research is part of the Ph.D. dissertation of M. Farah, completed at University of California, Santa Cruz, and was performed while R.D. Morris was employed by the Universities Space Research Association, Mountain View, USA. The authors thank Gary Swift (Xilinx) and Steve Guertin (JPL) for providing the test data, Chuck Foster for useful discussions, and Allan Tylka (NRL) for providing the CREME96 executable. They also wish to thank

an Associate Editor for helpful suggestions on the formulation of the model in Section 5, and two reviewers for several useful comments.

Appendix A: Posterior simulation method for the semiparametric model

Here, we provide details on MCMC posterior inference for the DP-based semiparametric model of Section 3.1. We focus on the scenario where there is a single count c_i for each distinct LET value ℓ_i, $i = 1, \ldots, N$ (as in the data for device B). The extension to the repeated measurements case (device A) is straightforward.

Therefore, based on the likelihood in (5) and the prior in (4), and under, say, the Weibull centering distribution in (2), the joint posterior distribution can be expressed as

$$\begin{aligned} p(\sigma_0, \boldsymbol{\theta}, \alpha, w, s \mid D) \propto \sigma_0^{\sum_{i=1}^N c_i} \exp(-\sigma_0 \sum\nolimits_{i=1}^N f_i\, \theta_i) \prod_{i=1}^N \theta_i^{c_i} \\ \times \frac{\Gamma(\alpha)}{\prod_{i=1}^{N+1} \Gamma(d_i)} \theta_1^{d_1-1} (\theta_2 - \theta_1)^{d_2-1} \ldots (\theta_N - \theta_{N-1})^{d_N-1} (1-\theta_N)^{d_{N+1}-1} \\ \times p(\sigma_0) p(\alpha) \exp(-b_w w) \exp(-b_s s), \end{aligned}$$

where $p(\sigma_0)$ and $p(\alpha)$ are the prior densities for σ_0 and α, respectively (the prior choices for σ_0, α and $\boldsymbol{\psi} = (w, s)$ were discussed in Section 3.3). Moreover, $\theta_i = G(\ell_i)$, $i = 1, ..., N$, and the d_i, $i = 1, ..., N+1$, are defined in Section 3.1; note that the d_i depend on DP prior hyperparameters α and $\boldsymbol{\psi}$.

The posterior full conditional for σ_0 is given by a gamma distribution under an exponential prior for σ_0, and by a truncated gamma distribution over $(0, \sigma_{\max})$ under the uniform prior on $(0, \sigma_{\max})$ for σ_0.

The DP prior hyperparameters enter the expression of the posterior distribution in a complex fashion that does not allow direct sampling. We update $\boldsymbol{\psi} = (w, s)$ with a random walk Metropolis-Hastings step based on a bivariate normal proposal on the logarithmic scale. We use an analogous Metropolis-Hastings step for the DP precision parameter α. Our implementation resulted in acceptance rates between 25% and 30% for $\boldsymbol{\psi}$ and α, respectively, with relatively low autocorrelation.

Updating the parameters in $\boldsymbol{\theta}$ is the most challenging aspect of the posterior simulation approach. Even with carefully chosen and tuned proposal distributions, a Metropolis-Hastings step for the entire vector $\boldsymbol{\theta}$ proved to be impractical. We instead use slice sampling (e.g., Damien, Wakefield and Walker, 1999; Neal, 2003) for each component of $\boldsymbol{\theta}$. This approach yields readily implemented *automatic* updates for each θ_i as described below. The resulting posterior samples are heavily correlated, but this can be controlled with appropriate levels of thinning of the MCMC output.

The posterior full conditional for θ_1 is given by

$$p(\theta_1 \mid \theta_2, \sigma_0, \alpha, \boldsymbol{\psi}, D) \propto \theta_1^{c_1+d_1-1} \exp(-\sigma_0 f_1 \theta_1)(\theta_2 - \theta_1)^{d_2-1}$$

and, therefore, we can introduce latent variables u_1 and u_2, each defined on $\mathbb{R}^+$, such that

$$p(\theta_1, u_1, u_2 \mid \theta_2, \sigma_0, \alpha, \boldsymbol{\psi}, D) \propto \exp(-\sigma_0 f_1 \theta_1) I\left\{u_1 < \theta_1^{c_1+d_1-1}\right\} I\left\{u_2 < (\theta_2 - \theta_1)^{d_2-1}\right\}.$$

Now, the MCMC algorithm is extended to draw from the full conditionals for u_1, u_2, and θ_1, arising from $p(\theta_1, u_1, u_2 \mid \theta_2, \sigma_0, \alpha, \boldsymbol{\psi}, D)$. The full conditional distributions for u_1 and u_2 are uniform over $(0, \theta_1^{c_1+d_1-1})$ and $(0, (\theta_2 - \theta_1)^{d_2-1})$, respectively. The full conditional for θ_1 is given by an $\text{Exp}(\sigma_0 f_1)$ distribution truncated such that $u_1^{(c_1+d_1-1)^{-1}} < \theta_1 < \theta_2 - u_2^{(d_2-1)^{-1}}$, if $d_2 > 1$, or $\theta_1 > \max\{u_1^{(c_1+d_1-1)^{-1}}, \theta_2 - u_2^{(d_2-1)^{-1}}\}$, if $0 < d_2 < 1$. The truncated exponential conditional for θ_1 can be efficiently sampled using the inverse cdf method.

Next, for each $i = 2, ..., N-1$, the posterior full conditional for θ_i,

$$p(\theta_i \mid \theta_{i-1}, \theta_{i+1}, \sigma_0, \alpha, \boldsymbol{\psi}, D) \propto \theta_i^{c_i} \exp(-\sigma_0 f_i \theta_i)(\theta_i - \theta_{i-1})^{d_i-1}(\theta_{i+1} - \theta_i)^{d_{i+1}-1}.$$

We now introduce $\mathbb{R}^+$-valued latent variables v_1, v_2, and v_3, such that

$$p(\theta_i, v_1, v_2, v_3 \mid \theta_{i-1}, \theta_{i+1}, \sigma_0, \alpha, \boldsymbol{\psi}, D) \propto \exp(-\sigma_0 f_i \theta_i) I\left\{v_1 < \theta_i^{c_i}\right\} \times I\left\{v_2 < (\theta_i - \theta_{i-1})^{d_i-1}\right\} I\left\{v_3 < (\theta_{i+1} - \theta_i)^{d_{i+1}-1}\right\},$$

which leads to $\text{Unif}(0, \theta_i^{c_i})$, $\text{Unif}(0, (\theta_i - \theta_{i-1})^{d_i-1})$, and $\text{Unif}(0, (\theta_{i+1} - \theta_i)^{d_{i+1}-1})$ full conditionals for v_1, v_2, and v_3, respectively. Moreover, the full conditional for θ_i is truncated $\text{Exp}(\sigma_0 f_i)$, where:

- $\max\left(v_1^{c_i^{-1}}, \theta_{i-1} + v_2^{(d_i-1)^{-1}}\right) < \theta_i < \theta_{i+1} - v_3^{(d_{i+1}-1)^{-1}}$, if $d_i > 1$ and $d_{i+1} > 1$
- $\max\left(v_1^{c_i^{-1}}, \theta_{i-1} + v_2^{(d_i-1)^{-1}}, \theta_{i+1} - v_3^{(d_{i+1}-1)^{-1}}\right) < \theta_i$, if $d_i > 1$ and $d_{i+1} < 1$
- $v_1^{c_i^{-1}} < \theta_i < \min\left(\theta_{i-1} + v_2^{(d_i-1)^{-1}}, \theta_{i+1} - v_3^{(d_{i+1}-1)^{-1}}\right)$, if $d_i < 1$ and $d_{i+1} > 1$
- $\max\left(v_1^{c_i^{-1}}, \theta_{i+1} - v_3^{(d_{i+1}-1)^{-1}}\right) < \theta_i < \theta_{i-1} + v_2^{(d_i-1)^{-1}}$, if $d_i < 1$ and $d_{i+1} < 1$.

Finally, the posterior full conditional for θ_N,

$$p(\theta_N \mid \theta_{N-1}, \sigma_0, \alpha, \boldsymbol{\psi}, D) \propto \theta_N^{c_N} \exp(-\sigma_0 f_N \theta_N)(\theta_N - \theta_{N-1})^{d_N-1}(1 - \theta_N)^{d_{N+1}-1}$$

is augmented with $\mathbb{R}^+$-valued latent variables z_1, z_2, and z_3, to

$$p(\theta_N, z_1, z_2, z_3 \mid \theta_{N-1}, \sigma_0, \alpha, \boldsymbol{\psi}, D) \propto \exp(-\sigma_0 f_N \theta_N) I\left\{z_1 < \theta_N^{c_N}\right\} \times I\left\{z_2 < (\theta_N - \theta_{N-1})^{d_N-1}\right\} I\left\{z_3 < (1 - \theta_N)^{d_{N+1}-1}\right\}.$$

Here, the full conditionals for z_1, z_2, and z_3 are $\text{Unif}(0, \theta_N^{c_N})$, $\text{Unif}(0, (\theta_N - \theta_{N-1})^{d_N-1})$, and $\text{Unif}(0, (1 - \theta_N)^{d_{N+1}-1})$, respectively, and the full conditional for θ_N is a truncated $\text{Exp}(\sigma_0 f_N)$ distribution, where

- $\max\left(z_1^{c_N^{-1}}, \theta_{N-1} + z_2^{(d_N-1)^{-1}}\right) < \theta_N < 1 - z_3^{(d_{N+1}-1)^{-1}}$, if $d_N > 1$ and $d_{N+1} > 1$

- $\max\left(z_1^{c_N^{-1}}, \theta_{N-1} + z_2^{(d_N-1)^{-1}}, 1 - z_3^{(d_{N+1}-1)^{-1}}\right) < \theta_N$, if $d_N > 1$ and $d_{N+1} < 1$
- $z_1^{c_N^{-1}} < \theta_N < \min\left(\theta_{N-1} + z_2^{(d_N-1)^{-1}}, 1 - z_3^{(d_{N+1}-1)^{-1}}\right)$, if $d_N < 1$ and $d_{N+1} > 1$
- $\max\left(z_1^{c_N^{-1}}, 1 - z_3^{(d_{N+1}-1)^{-1}}\right) < \theta_N < \theta_{N-1} + z_2^{(d_N-1)^{-1}}$, if $d_N < 1$ and $d_{N+1} < 1$.

Appendix B: Predictive inference for the cross-section vs. LET curve

This section develops the approach to posterior predictive inference for the cross-section curve at values within or outside the range of observed LET values.

Consider first prediction at new LET values, $\tilde{\boldsymbol{\ell}} = (\tilde{\ell}_1, \ldots, \tilde{\ell}_M)$, such that $\ell_i < \tilde{\ell}_1 < \cdots < \tilde{\ell}_M < \ell_{i+1}$, for any $i = 1, \ldots, N-1$. The full model, including the corresponding new $\tilde{\theta}_m = G(\tilde{\ell}_m)$, $m = 1, ..., M$, can be written as

$$\begin{aligned} p(\sigma_0, \tilde{\boldsymbol{\theta}}, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D) &\propto p(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi}) p(\sigma_0) p(\alpha) p(\boldsymbol{\psi}) \prod_{i=1}^{N} \text{Poisson}(c_i; f_i \sigma_0 \theta_i) \\ &= p(\tilde{\boldsymbol{\theta}} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi}) p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D), \end{aligned}$$

where $\tilde{\boldsymbol{\theta}} = (\tilde{\theta}_1, \ldots, \tilde{\theta}_M)$, and $p(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})$ is the joint prior for $(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta})$ induced by the DP prior for $G(\cdot)$, with density that extends the form for $p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})$ in (4). Specifically,

$$\begin{aligned} p(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi}) = \frac{\Gamma(\alpha)}{\Gamma(d'_{i+1}) \prod_{\substack{n=1 \\ n \neq i+1}}^{N+1} \Gamma(d_n) \prod_{m=1}^{M} \Gamma(\tilde{d}_m)} \; & \theta_1^{d_1 - 1} (\theta_2 - \theta_1)^{d_2 - 1} \ldots (\theta_i - \theta_{i-1})^{d_i - 1} \\ & \times (\tilde{\theta}_1 - \theta_i)^{\tilde{d}_1 - 1} (\tilde{\theta}_2 - \tilde{\theta}_1)^{\tilde{d}_2 - 1} \ldots (\tilde{\theta}_M - \tilde{\theta}_{M-1})^{\tilde{d}_M - 1} (\theta_{i+1} - \tilde{\theta}_M)^{d'_{i+1} - 1} \\ & \times (\theta_{i+2} - \theta_{i+1})^{d_{i+2} - 1} \ldots (\theta_N - \theta_{N-1})^{d_N - 1} (1 - \theta_N)^{d_{N+1} - 1}, \end{aligned} \tag{8}$$

where $d'_{i+1} = \alpha(G_0(\ell_{i+1}; \boldsymbol{\psi}) - G_0(\tilde{\ell}_M; \boldsymbol{\psi}))$, $\tilde{d}_1 = \alpha(G_0(\tilde{\ell}_1; \boldsymbol{\psi}) - G_0(\ell_i; \boldsymbol{\psi}))$, and for $m = 2, \ldots, M$, $\tilde{d}_m = \alpha(G_0(\tilde{\ell}_m; \boldsymbol{\psi}) - G_0(\tilde{\ell}_{m-1}; \boldsymbol{\psi}))$. Moreover, the d_i, $i = 1, ..., N+1$, are as defined in Section 3.1.

Hence, $p(\tilde{\boldsymbol{\theta}} \mid D) = \int p(\tilde{\boldsymbol{\theta}} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi}) p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D) \, \mathrm{d}\sigma_0 \mathrm{d}\boldsymbol{\theta} \mathrm{d}\alpha \mathrm{d}\boldsymbol{\psi}$, and therefore, using Monte Carlo integration based on the posterior draws from $p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D)$, we can sample $p(\tilde{\boldsymbol{\theta}} \mid D)$ by additional sampling from $p(\tilde{\boldsymbol{\theta}} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi})$.

Using the expressions for $p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})$ and $p(\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})$ in (4) and (8), respectively, and the fact that $d_{i+1} = d'_{i+1} + \sum_{m=1}^{M} \tilde{d}_m$, the density for $p(\tilde{\boldsymbol{\theta}} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi})$ can be obtained as

$$\frac{\Gamma(d_{i+1})(\theta_{i+1} - \theta_i)^{-M}}{\Gamma(d'_{i+1}) \prod_{m=1}^{M} \Gamma(\tilde{d}_m)} \left(\frac{\tilde{\theta}_1 - \theta_i}{\theta_{i+1} - \theta_i}\right)^{\tilde{d}_1 - 1} \cdots \left(\frac{\tilde{\theta}_M - \tilde{\theta}_{M-1}}{\theta_{i+1} - \theta_i}\right)^{\tilde{d}_M - 1} \left(\frac{\theta_{i+1} - \tilde{\theta}_M}{\theta_{i+1} - \theta_i}\right)^{d'_{i+1} - 1}.$$

Hence, $p(\tilde{\boldsymbol{\theta}} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi})$ can be readily sampled, since the density above corresponds to random vector $(\theta_i + (\theta_{i+1} - \theta_i)\omega_1, \ldots, \theta_i + (\theta_{i+1} - \theta_i)\omega_M)$, where $(\omega_1, \ldots, \omega_M)$ follows an ordered Dirichlet distribution with parameters $(\tilde{d}_1, \ldots, \tilde{d}_M, d'_{i+1})$.

We next turn to extrapolation for the cross-section vs. LET curve beyond the largest observed LET value. Consider prediction at new LET values $\boldsymbol{\ell'} = (\ell'_1, \ldots, \ell'_M)$, such that $\ell_N < \ell'_1 < \cdots < \ell'_M$. The model augmented with $\boldsymbol{\ell'}$ and the corresponding $\boldsymbol{\theta'} = (\theta'_1, ..., \theta'_M)$, where $\theta'_m = G(\ell'_m)$, $m = 1, ..., M$, becomes $p(\sigma_0, \boldsymbol{\theta'}, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D) = p(\boldsymbol{\theta'} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi}) p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D)$. In this case, $p(\boldsymbol{\theta'} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi})$ can be recognized as the density for random vector $(\theta_N + (1-\theta_N)\omega_1, \ldots, \theta_N + (1-\theta_N)\omega_M)$, where $(\omega_1, \ldots, \omega_M)$ has an ordered Dirichlet distribution with parameters $(d'_1, \ldots, d'_M, d'_{M+1})$. Here, $d'_1 = \alpha\,(G_0(\ell'_1; \boldsymbol{\psi}) - G_0(\ell_N; \boldsymbol{\psi}))$, $d'_m = \alpha\left(G_0(\ell'_m; \boldsymbol{\psi}) - G_0(\ell'_{m-1}; \boldsymbol{\psi})\right)$, $m = 2, \ldots, M$, and $d'_{M+1} = \alpha\,(1 - G_0(\ell'_M; \boldsymbol{\psi}))$, such that $d_{N+1} = \sum_{m=1}^{M+1} d'_m$.

Finally, extrapolation in the region below the smallest observed LET value is obtained in a similar fashion. Consider prediction at new LET values $\boldsymbol{\ell''} = (\ell''_1, \ldots, \ell''_M)$, such that $\ell''_1 < \cdots < \ell''_M < \ell_1$. Let $\boldsymbol{\theta''} = (\theta''_1, ..., \theta''_M)$, where $\theta''_m = G(\ell''_m)$, for $m = 1, ..., M$, and $d''_1 = \alpha G_0(\ell''_1; \boldsymbol{\psi})$, $d''_m = \alpha\left(G_0(\ell''_m; \boldsymbol{\psi}) - G_0(\ell''_{m-1}; \boldsymbol{\psi})\right)$, for $m = 2, \ldots, M$, and $d''_{M+1} = \alpha\,(G_0(\ell_1; \boldsymbol{\psi}) - G_0(\ell''_M; \boldsymbol{\psi}))$, such that $d_1 = \sum_{m=1}^{M+1} d''_m$. Then, the required conditional density, $p(\boldsymbol{\theta''} \mid \boldsymbol{\theta}, \alpha, \boldsymbol{\psi})$, corresponds to the distribution of random vector $(\theta_1\omega_1, \ldots, \theta_1\omega_M)$, where $(\omega_1, \ldots, \omega_M)$ has an ordered Dirichlet distribution with parameters $(d''_1, \ldots, d''_M, d''_{M+1})$.

Appendix C: CPO estimation

Here, we develop an efficient approach to estimating the CPOs and the LPML statistic, discussed in Section 3.4. We discuss the approach for the setting where there is a single observation at each distinct LET value; only changes in notation are needed for the setting that involves repeated measurements at each LET value.

CPO statistics can be estimated across a wide variety of hierarchical models, using the MCMC output from the fit of the model to the full data vector (see, e.g., Chen, Shao and Ibrahim, 2000). By definition, under the parametric Weibull model, $\text{CPO}_i = \int \text{Poisson}(c_i; f_i\sigma_0\{1-\exp(-(\ell_i/w)^s)\}) p(\sigma_0, w, s \mid D_{(i)})\,\mathrm{d}\sigma_0\mathrm{d}w\mathrm{d}s$, where $D_{(i)} = \{(c_{i'}, f_{i'}, \ell_{i'}) : i' \neq i\}$. It is straightforward to show that

$$\text{CPO}_i = \left\{ \int \frac{1}{\text{Poisson}(c_i; f_i\sigma_0\{1 - \exp(-(\ell_i/w)^s)\})} p(\sigma_0, w, s \mid D)\,\mathrm{d}\sigma_0\mathrm{d}w\mathrm{d}s \right\}^{-1}, \; i = 1, ..., N.$$

A similar expression can be obtained for the parametric lognormal model.

Turning to the DP-based model of Section 3.1, for any specified $i = 1, ..., N$, let $p(\sigma_0, \boldsymbol{\theta}_{(i)}, \alpha, \boldsymbol{\psi} \mid D_{(i)})$, where $\boldsymbol{\theta}_{(i)} = \{\theta_{i'} : i' \neq i\}$, be the posterior distribution based on reduced data vector $D_{(i)} = \{(c_{i'}, f_{i'}, \ell_{i'}) : i' \neq i\}$. Moreover, denote by J and K the normalizing constants for posterior distributions $p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D)$ and $p(\sigma_0, \boldsymbol{\theta}_{(i)}, \alpha, \boldsymbol{\psi} \mid D_{(i)})$, respectively.

Using a derivation similar to the one in Appendix B (the interpolation approach for ℓ_i with $i = 2, ..., N-1$, and the extrapolation approach for ℓ_1 and ℓ_N), the cross-validation

predictive density for the i-th observation can be written as

$$p(z_i \mid (f_i, \ell_i), D_{(i)}) = \int \text{Poisson}(z_i; f_i\sigma_0\theta_i)p(\theta_i \mid \boldsymbol{\theta}_{(i)}, \alpha, \boldsymbol{\psi})p(\sigma_0, \boldsymbol{\theta}_{(i)}, \alpha, \boldsymbol{\psi} \mid D_{(i)})\mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi},$$

where $p(\theta_i \mid \boldsymbol{\theta}_{(i)}, \alpha, \boldsymbol{\psi})$ is a rescaled Beta distribution.

By definition, $\text{CPO}_i = p(c_i \mid (f_i, \ell_i), D_{(i)})$, and therefore we have

$$\begin{aligned}
\text{CPO}_i &= K^{-1} \int \text{Poisson}(c_i; f_i\sigma_0\theta_i)p(\theta_i \mid \boldsymbol{\theta}_{(i)}, \alpha, \boldsymbol{\psi}) \left(\textstyle\prod_{\{i' \neq i\}} \text{Poisson}(c_{i'}; f_{i'}\sigma_0\theta_{i'})\right) \\
&\qquad\qquad p(\boldsymbol{\theta}_{(i)} \mid \alpha, \boldsymbol{\psi})p(\sigma_0)p(\alpha)p(\boldsymbol{\psi})\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi} \\
&= K^{-1} \int \left(\textstyle\prod_{j=1}^{N} \text{Poisson}(c_j; f_j\sigma_0\theta_j)\right) p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})p(\sigma_0)p(\alpha)p(\boldsymbol{\psi})\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi} \\
&= K^{-1}J.
\end{aligned}$$

Next, for $\text{CPO}_i^{-1} = KJ^{-1}$, we obtain

$$\begin{aligned}
\text{CPO}_i^{-1} &= J^{-1} \int \left(\textstyle\prod_{\{i' \neq i\}} \text{Poisson}(c_{i'}; f_{i'}\sigma_0\theta_{i'})\right) p(\boldsymbol{\theta}_{(i)} \mid \alpha, \boldsymbol{\psi})p(\sigma_0)p(\alpha)p(\boldsymbol{\psi})\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}_{(i)}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi} \\
&= J^{-1} \int \left(\textstyle\prod_{\{i' \neq i\}} \text{Poisson}(c_{i'}; f_{i'}\sigma_0\theta_{i'})\right) \left(\int p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})\,\mathrm{d}\theta_i\right) \\
&\qquad\qquad p(\sigma_0)p(\alpha)p(\boldsymbol{\psi})\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}_{(i)}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi} \\
&= J^{-1} \int \left(\textstyle\prod_{\{i' \neq i\}} \text{Poisson}(c_{i'}; f_{i'}\sigma_0\theta_{i'})\right) p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})p(\sigma_0)p(\alpha)p(\boldsymbol{\psi})\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi} \\
&= J^{-1} \int (\text{Poisson}(c_i; f_i\sigma_0\theta_i))^{-1} \left(\textstyle\prod_{j=1}^{N} \text{Poisson}(c_j; f_j\sigma_0\theta_j)\right) \\
&\qquad\qquad p(\boldsymbol{\theta} \mid \alpha, \boldsymbol{\psi})p(\sigma_0)p(\alpha)p(\boldsymbol{\psi})\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi} \\
&= \int (\text{Poisson}(c_i; f_i\sigma_0\theta_i))^{-1} p(\sigma_0, \boldsymbol{\theta}, \alpha, \boldsymbol{\psi} \mid D)\ \mathrm{d}\sigma_0\mathrm{d}\boldsymbol{\theta}\mathrm{d}\alpha\mathrm{d}\boldsymbol{\psi}.
\end{aligned}$$

Finally, noting that in the expression above all parameters other than σ_0 and θ_i can be marginalized from the joint posterior distribution, we obtain

$$\text{CPO}_i = \left\{ \int \frac{1}{\text{Poisson}(c_i; f_i\sigma_0\theta_i)} p(\sigma_0, \theta_i \mid D)\,\mathrm{d}\sigma_0\mathrm{d}\theta_i \right\}^{-1}, \quad i = 1, ..., N,$$

which can be computed using a Monte Carlo estimate from samples of the parameters already generated by MCMC.

References

Albert, J.H. (1988). Computational methods using a Bayesian hierarchical generalized linear model. *Journal of the American Statistical Association*, 83, 1037-1044.

Antoniak, C.E. (1974). Mixtures of Dirichlet processes with applications to nonparametric problems. *Annals of Statistics*, 2, 1152-1174.

Chen, M.-H., Shao, Q.-M. and Ibrahim, J.G. (2000). *Monte Carlo Methods in Bayesian Computation*. Springer-Verlag: New York.

Damien, P., Wakefield, J. and Walker, S. (1999). Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society, Series B*, 61, 331-344.

Dodd, P.E. and Massengill, L.W. (2003). Basic mechanisms and modeling of single-event upset in digital microelectronics. *IEEE Transactions on Nuclear Science*, 50, 583-602.

Dunson, D.B. (2005). Bayesian semiparametric isotonic regression for count data. *Journal of the American Statistical Association*, 100, 618-627.

Ferguson, T.S. (1973). A Bayesian analysis of some nonparametric problems. *Annals of Statistics*, 1, 209-230.

Geisser, S. and Eddy, W.F. (1979). A predictive approach to model selection. *Journal of the American Statistical Association*, 74, 153-160.

Gelfand, A.E. and Kuo, L. (1991). Nonparametric Bayesian bioassay including ordered polytomous response. *Biometrika*, 78, 657-666.

Kottas, A., Branco, M.D. and Gelfand, A.E. (2002). A nonparametric Bayesian modeling approach for cytogenetic dosimetry. *Biometrics*, 58, 593-600.

Lacoe, R.C., Osborn, J.V., Koga, R., Brown, S. and Mayer, D.C. (2000). Application of hardness-by-design methodology to radiation-tolerant ASIC technologies. *IEEE Transactions on Nuclear Science*, 47, 2234-2341.

Ladbury, R., Reed, R.A., Marshall, P., LaBel, K.A., Anantaraman, R., Fox, R., Sanderson, D.P., Stolz, A., Yurkon, J., Zeller, A.F. and Stetson J.W. (2004). Performance of the High-Energy Single-Event Effects Test Facility (SEETF) at Michigan State University's National Superconducting Cyclotron Laboratory (NSCL). *IEEE Transactions on Nuclear Science*, 51, 3664-3668.

Lavine, M. and Mockus, A. (1995). A nonparametric Bayes method for isotonic regression. *Journal of Statistical Planning and Inference*, 46, 235-248.

Morris, R.D. and Foster, C.C. (2010). Cross-section estimation in the presence of uncertain dosimetry. *IEEE Transactions on Nuclear Science*, 57, 3528-3536.

Morris, R.D. and Foster, C.C. (2011). Characterizing the uncertainty in dosimetry using data from multiple tests. *IEEE Transactions on Nuclear Science*, 58, 3145-3151.

Mukhopadhyay, S. (2000). Bayesian nonparametric inference on the dose level with specified response rate. *Biometrics*, 56, 220-226.

Neal, R. M. (2003). Slice sampling (with discussion). *Annals of Statistics*, 31, 705-767.

Petersen, E.L. (1996). Cross section measurements and upset rate calculations. *IEEE Transactions on Nuclear Science*, 43, 2805-2813.

Petersen, E.L. (2010). Dosimetry comparisons between various single event test facilities. *IEEE Transactions on Nuclear Science*, 57, 3477-3482.

Petersen, E.L., Pickel, J.C., Adams, J.H. and Smith, E.C. (1992). Rate prediction for single event effects – a critique. *IEEE Transactions on Nuclear Science*, 39, 1577-1599.

Pickel, J.C. (1996). Single-event effects rate prediction. *IEEE Transactions on Nuclear Science*, 43, 483-495.

Swift, G., Carmichael, C. and Allen, G. (2008). Virtex-4QV static SEU characterization summary. *JPL Publication 08-164/08*.

Tylka, A.J., Adams, J.H., Boberg, P.R., Brownstein, B., Dietrich, W.F., Flueckiger, E.O., Petersen, E.L., Shea, M.A., Smart, D.F. and Smith, E.C. (1997). CREME96: A Revision of the Cosmic Ray Effects on Micro-Electronics Code. *IEEE Transactions on Nuclear Science*, 44, 2150-2160.

Warren, K.M., Weller, R.A., Mendenhall, M.H., Reed, R.A., Ball, D.R., Howe, C.L., Olson, B.D., Alles, M.L., Massengill, L.W., Schrimpf, R.D., Haddad, N.F., Doyle, S.E., McMorrow, D., Melinger, J.S. and Lotshaw, W.T. (2005). The contribution of nuclear reactions to heavy ion single event upset cross-section measurements in a high-density SEU hardened SRAM. *IEEE Transactions on Nuclear Science*, 52, 2125-2131.

Weller, R.A., Reed, R.A., Warren, K.M., Mendenhall, M.H., Sierawski, B.D., Schrimpf, R.D. and Massengill, L.W. (2009). General framework for single event effects rate prediction in microelectronics. *IEEE Transactions on Nuclear Science*, 56, 6, 3098-3108.

Xapsos, M.A., Weatherford, T.R. and Shapiro, P. (1993). The shape of heavy ion upset cross section curves. *IEEE Transactions on Nuclear Science*, 40, 1812-1819.

Table 1. Test data for the MC7447AT PowerPC. The counts are of the $0 \to 1$ upsets in the cache bits. Data courtesy JPL.

LET $(MeV \cdot cm^2/mg)$	fluence $(\times 10^3)$	counts
1.7	3000, 10000, 30000	14, 66, 263
3.5	30000, 30000, 30000, 30000, 27600, 36000	1721, 1722, 1291, 1342, 1007, 1350
9.5	450, 300, 300, 300, 300, 300, 2000, 2000, 2000, 1750, 2000, 2000	161, 109, 114, 114, 118, 131, 696, 766, 632, 524, 690, 688
13.9	2000, 2000, 2000	800, 691, 748

Table 2. Test data for the XQR4VLX200 FPGA device. The counts are for the POR SEFI failure mode. Data courtesy Xilinx.

LET $(MeV \cdot cm^2/mg)$	fluence $(\times 10^3)$	counts
2	100000	6
4	100000	11
6.8	50000	11
16.9	19400	10
22	10000	6
30	13400	19
90.3	9180	28

Table 3. Results for the LPML model comparison criterion for device A.

Device A, "leave-one-out" cross-validation Model Specification	LPML	Device A, "block" cross-validation Model Specification	LPML
Semiparametric, $G_0 =$ Weibull, $p(\alpha) \sim \Gamma$	-12.56	Semiparametric, $G_0 =$ Weibull, $p(\alpha) \sim \Gamma$	-6.50
Semiparametric, $G_0 =$ Weibull, $p(\alpha) \sim U$	-12.53	Semiparametric, $G_0 =$ Weibull, $p(\alpha) \sim U$	-6.44
Semiparametric, $G_0 =$ lognormal, $p(\alpha) \sim \Gamma$	-12.79	Semiparametric, $G_0 =$ lognormal, $p(\alpha) \sim \Gamma$	-6.39
Semiparametric, $G_0 =$ lognormal, $p(\alpha) \sim U$	-12.59	Semiparametric, $G_0 =$ lognormal, $p(\alpha) \sim U$	-6.26
Parametric Weibull	-13.50	Parametric Weibull	-12.69
Parametric lognormal	-18.31	Parametric lognormal	-103.45

Table 4. Results for the LPML model comparison criterion for device B.

Device B Model Specification	LPML
Semiparametric, $G_0 =$ Weibull, $p(\alpha) \sim \Gamma$	-2.89
Semiparametric, $G_0 =$ Weibull, $p(\alpha) \sim U$	-2.85
Semiparametric, $G_0 =$ lognormal, $p(\alpha) \sim \Gamma$	-2.99
Semiparametric, $G_0 =$ lognormal, $p(\alpha) \sim U$	-2.99
Parametric Weibull	-2.53
Parametric lognormal	-2.84

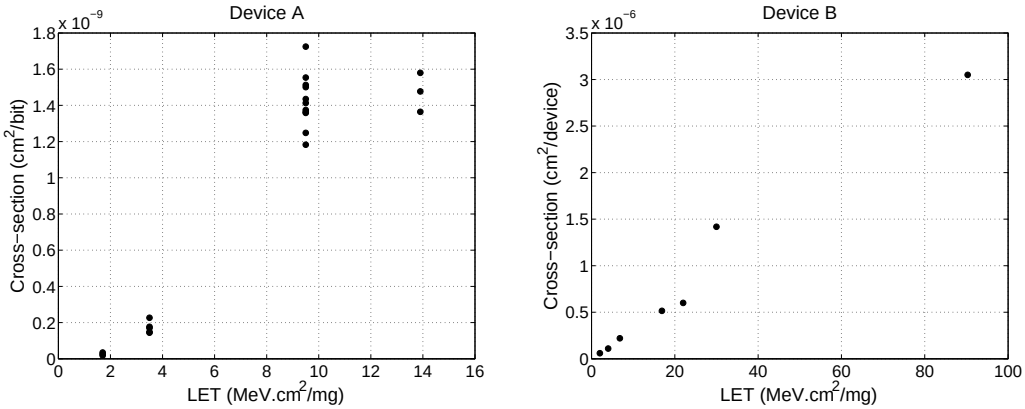

Fig. 1. Experimental data obtained from particle accelerator experiments performed on two different devices. Note that the cross-section for device A is per-bit, whereas for device B it is per-device.

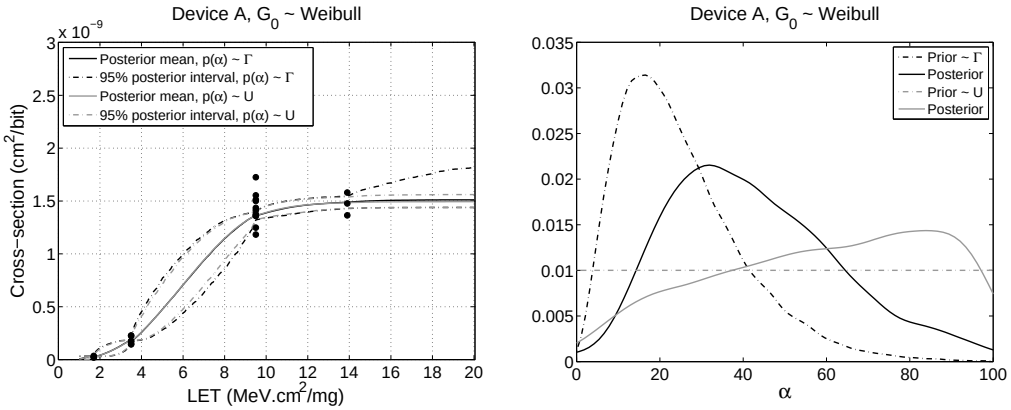

Fig. 2. Point estimates and 95% probability intervals, overlaid on the data, for the cross-section vs. LET curve for device A under the DP model with a Weibull centering distribution (left panel). Prior and posterior densities for α (right panel). Γ and U denote the gamma and uniform prior for α, respectively.

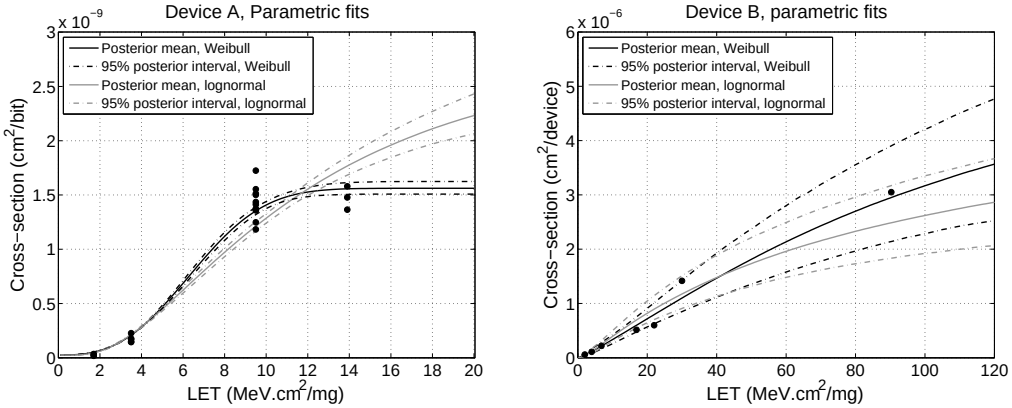

Fig. 3. Point estimates and 95% probability intervals, overlaid on the data, for the cross-section vs. LET curve under the parametric models for device A (left panel) and device B (right panel).

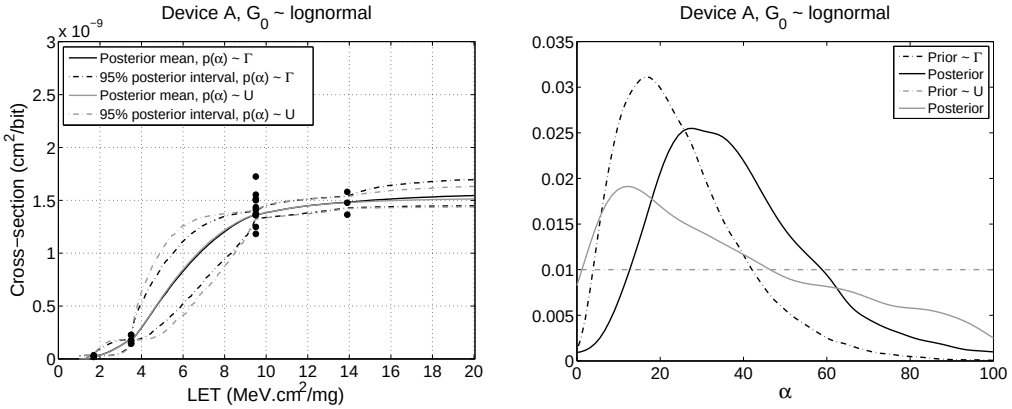

Fig. 4. Point estimates and 95% probability intervals, overlaid on the data, for the cross-section vs. LET curve for device A under the DP model using a lognormal centering distribution (left panel). Prior and posterior densities for α (right panel).

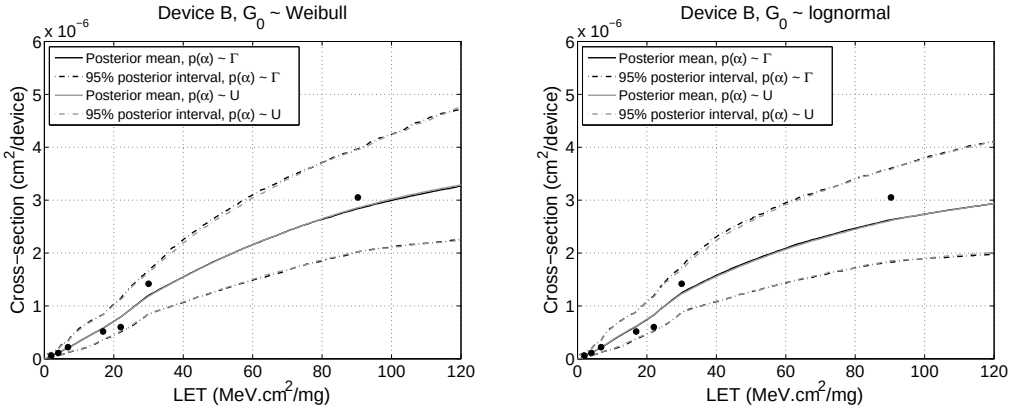

Fig. 5. Point estimates and 95% probability intervals, overlaid on the data, for the cross-section vs. LET curve for device B based on the DP model with a Weibull and a lognormal centering distribution (left and right panels, respectively).

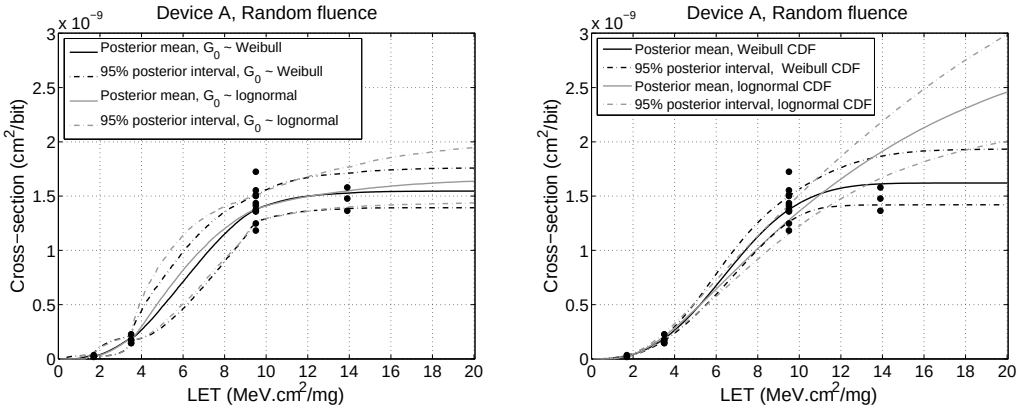

Fig. 6. Device A, random fluence case. Point estimates and 95% probability intervals, overlaid on the data, for the cross-section vs. LET curve under the DP model with a Weibull or lognormal centering distribution (left panel) and under the corresponding parametric models (right panel).

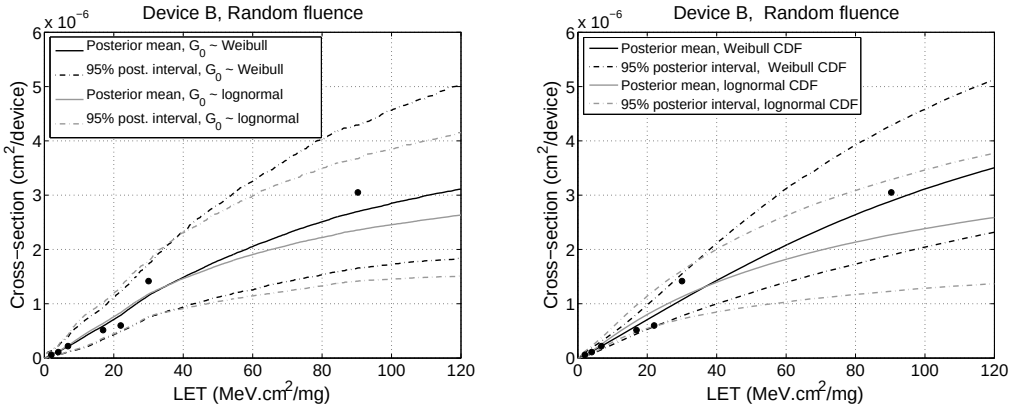

Fig. 7. Device B, random fluence case. Point estimates and 95% probability intervals, overlaid on the data, for the cross-section vs. LET curve under the DP model with the Weibull or lognormal centering distribution (left panel) and under the corresponding parametric models (right panel).

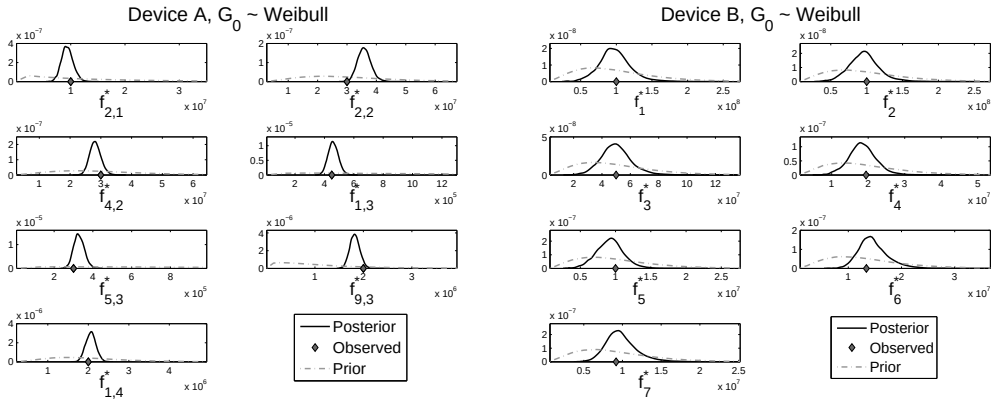

Fig. 8. Prior and posterior densities for random fluences under the DP semiparametric model with the Weibull centering distribution for device A (left panel) and device B (right panel). Each plot includes the observed fluence for the corresponding measurement.

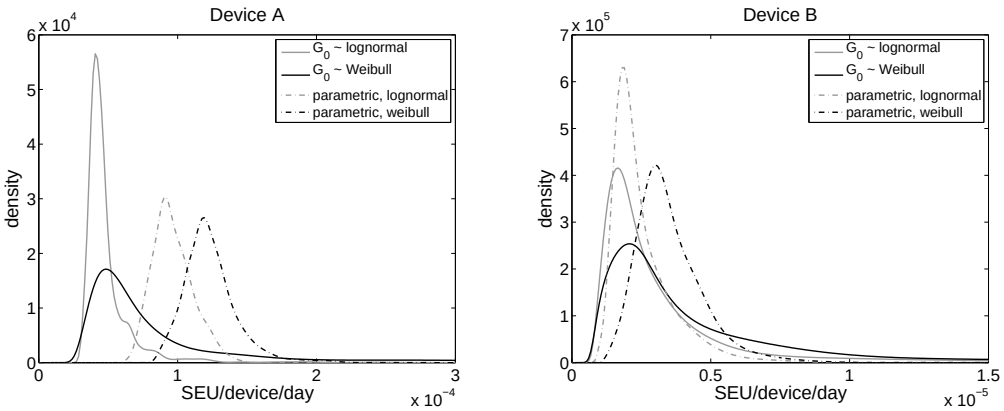

Fig. 9. Posterior densities of the predicted on-orbit upset rates for device A and B (left and right panels, respectively), using the two semiparametric models and the two parametric models with random fluences.