arXiv:1010.0632v1 [physics.data-an] 4 Oct 2010

Improved iterative Bayesian unfolding

G. D'Agostini
Università "La Sapienza" and INFN, Rome, Italy
(giulio.dagostini@roma1.infn.it, http://www.roma1.infn.it/~dagos)

Abstract

This paper reviews the basic ideas behind a Bayesian unfolding published some years ago and improves their implementation. In particular, uncertainties are now treated at all levels by probability density functions and their propagation is performed by Monte Carlo integration. Thus, small numbers are better handled and the final uncertainty does not rely on the assumption of normality. Theoretical and practical issues concerning the iterative use of the algorithm are also discussed. The new program, implemented in the R language, is freely available, together with sample scripts to play with toy models.

1 Introduction

Physicists like to think of 'true' distributions of physics quantities, i.e those distributions (graphically represented by histograms and here also referred as 'spectra') one would observe under idealized conditions that seldom – never, strictly speaking – happen in real live (ideal detector, no physical or instrumental background). The 'observed' distribution is then considered as a 'noisy distortion' of the true one. An important task of the experimental method is therefore to infer the true distribution from the observed one, i.e. to correct the observed spectrum for distortion and noise. This can be done by different methods that follow different approaches. Since this is not a review paper, I just outline the two different classes of strategies and then focus on the specific issues of this work.

In a first kind of approach a mathematical function for the true distribution is assumed (together with other functions to model the noise) and the

task becomes that of estimating the free parameters of the function(s). Indeed, we fall in the so called domain of *parametric inference* ('fits'), usually associated to names like 'least-squares' or 'maximum likelihood'.[1]

In parametric inference all information contained in the observed spectrum is, to say, 'distilled' into the model parameters. Parametric inference, especially when the conditions to use least-squares methods hold, is usually the best and fastest way to proceed, if we have good reasons to believe the hypothesized family of functions.

However, sometimes we wish to interpret the data as little as possible and just public 'something equivalent' to an experimental distribution, with the bin contents fluctuating according to an underlying multinomial distribution, but having possibly got rid of physical and instrumental distortions, as well as of background. In particle physics this second approach goes under the name of *unfolding* (or deconvolution, or restoration).

Several years ago a simple algorithm based on Bayes' theorem was presented [2] with which it was possible to achieve rather 'good' results ('good' compared with the difficulty of the task). The main improvements presented here concern the handling of small numbers and the evaluation of the uncertainty on the unfolded distribution, while the guiding ideas and the basic assumptions are substantially unchanged. To be more clear, the algorithm of Ref. [2] was relying on the underlying hypotheses of normality and 'small relative errors': 'best estimates' were provided, with uncertainty calculated from standard 'error propagation' formulas.[2] The new algorithm handles better small numbers, in the way it will be described in Sec. 3.2, and performs the propagation of uncertainty by sampling, i.e. by Monte Carlo (MC) integration. The algorithm has been implemented in a R language [4] code available on the author's web page [5].

The paper is organized in the following way. Section 2 gives a short introduction to Bayesian inference and to the specific application subject of this paper. Then the algorithm of Ref. [2] is reminded and the improvements are presented. The issue of iterative use of the algorithm is also discussed, although the program now gives also the possibility to use priors, provided

[1] However one can show that it is preferable to base the data analysis on more solid probabilistic ground, from which the mentioned methods might be derived (see e.g. Ref. [1]) under special conditions which fortunately hold in the large majority of practical cases (but it is better one knows when the conditions of validity hold and what to do when they do not!).

[2] It is worth remembering that these formulas rely on an hypothesis of linearity between *input* and *output quantities*, an hypothesis that holds approximately, in non-strictly linear problems, if the relative uncertainties of the input quantities are small enough, such that the dependence input$\rightarrow$output can be locally linearized [3].

by the user, over all possible true spectra. This option should avoid the need of iterations (but I anticipate that it might be not that easy to model such priors and then the iterative strategy remains a pragmatic solution). Finally, some results on toy models are presented.

Some technical issues concerning binomial and Dirichlet distribution, including the use of the latter as prior conjugate of the former, are reminded in Appendix A. A second appendix is dedicated to the handling of the zeros occurring in the evaluation of the smearing matrix (a cause that produces no event in some effect bins, as it usually happens, depending also on the statistics of the Monte Carlo simulation). As it happens with intermediate smoothing (or other kinds of regularization), this is a suggestion of how the problem can be approached, whose solution is delegated to the user, who is supposed to know the physics case under investigation.

2 Bayesian inference and Bayesian unfolding: from first principles to real life

The so called Bayesian inference is a way to learn about physical quantities and their relationships (all things we cannot directly see) from experimental data (what we can actually observe with our senses, usually mediated by more or less sophisticated detectors) using probability theory. This game is *conceptually* rather simple, proviso we accept that the intuitive meaning of probability – a scale to rank our beliefs that several events might happen, or that several hypotheses might be true – is suitable in scientific reasoning (see Refs. [6] and [7] and references therein for an introduction).

The starting point of a Bayesian inference is to build up a model for the deterministic ("*B follows from A*") and the probabilistic ("*B might follow from A*") connections that relate the several entities that enter the problem. This model has the interesting graphical representation of a network, usually called *Bayesian network* or *belief network* (to get an idea of their meaning and their application *Google* these keywords, or browse the *Wikipedia*). For example, Fig. 1, taken from Ref. [1], shows a Bayesian network to model two physical quantities, μ_x and μ_y, connected to each other with a 'law', whose parameters are denoted by $\boldsymbol{\theta}$. In the very elementary case depicted in the left hand network of Fig. 1 the solution to the problem can be rather simple, under some assumptions that, fortunately, hold very often in routine cases. But, in general, the solution is not that simple. Nevertheless, as stressed in Ref. [1], after one has built the graphical model, one is often more than half the way to the solution, thanks to the great progresses recently made

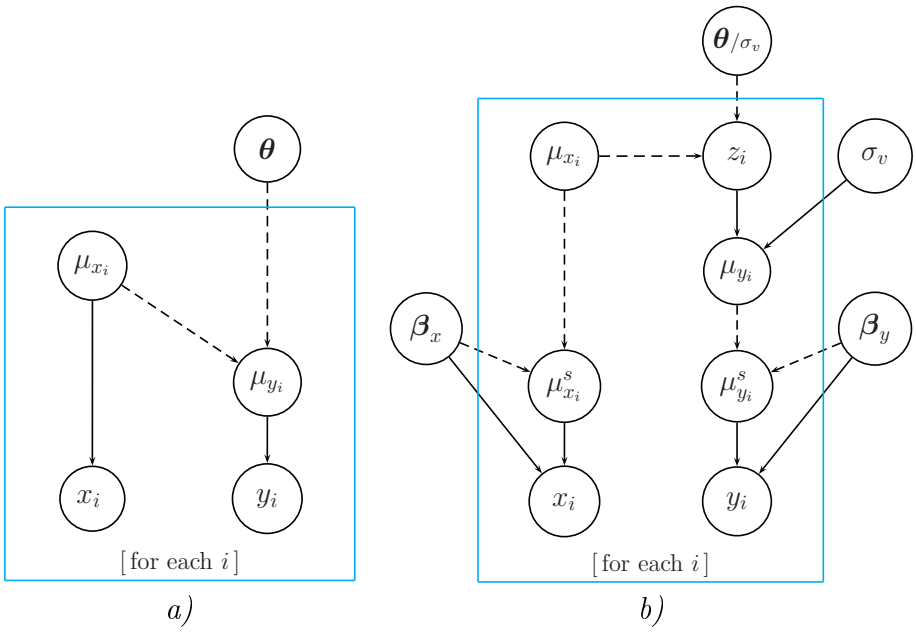

Figure 1: *Bayesian network describing a fit model [1]. x_i and y_i are the experimental observations, related to μ_{x_i} and μ_{y_i} by experimental errors. Instead, a deterministic 'law' connects the 'true' values μ_{y_i} to μ_{x_i} via the model parameters $\boldsymbol{\theta}$ (solid arrows stand for probabilistic links, dashed for deterministic). Network a) describes a simple model with errors only on the ordinate. Network b) takes into account errors on both axes, extra variability of the data points around the believed physical law and systematic effects.*

in Bayesian network computing.

The essence of the Bayesian unfolding of Ref. [2], that is the starting point also of this new version, is to make the problem discrete and to treat the 'cause' bins as independent degrees of freedom, i.e. without constraints among each other.[3] For this reason the algorithm can virtually handle any kind of 'smearing' and it is easily extensible, at least in principle and with the only limit due to computer power, to multidimensional problems. In fact, the core of the algorithm only knows about 'cause-cells' and 'effect-cells',[4] but it does not know the location of the cells in the configuration

[3] An important drawback of this feature will be discussed in section 4.

[4] 'Bin' and 'cell' are used here as synonyms, although 'cell' refers to a region on the configuration space and 'bin' to histograms representing them.

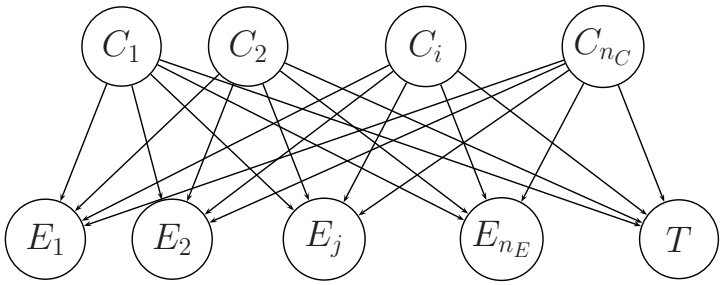

Figure 2: *Probabilistic links from causes to effects. The node indicated by T ('trash') stand for the inefficiency bin and corresponds to E_{n_E+1}*

space of the problem. For the same reason, the treatment of background, and even of several independent sources of background, can be easily embodied in the algorithm by just adding extra cause-cells, one cell per source of background. As a by-product, the algorithm also provides the number of events to be attributed to each source of background. (It is worth remembering that background might have an interesting physical meaning, and thus the estimation of the level 'noise' might provides indeed a physics measurement, as in the analysis of Ref. [8].)

Given the discretization of the problem, the Bayesian network relating causes and effects is that shown in Fig. 2, where we use the same notation of Ref. [2], with the addition of the effect bin T ('trash'), equivalent to E_{n_E+1}, to describe inefficiency (the reason to introduce this extra bin will become clear later).

Rephrasing the problem in probabilistic terms, the purpose of the unfolding is to find the 'true' number of events in each cause bin [$\#(C_i)$ in Fig. 3, indicated by $x(C_i)$ in the text], given the observed spectrum and assuming some knowledge about the smearing.

Since the links cause→effects have a probabilistic nature, it follows that also the links effect→causes will be probabilistic, and therefore it will be uncertain the number of events to be attributed to the cause-cells. We can only attempt to rank in probability all possible spectra that might have caused the observed one. In other words, the realistic goal of our analysis is not to determine the *the* true spectrum, but rather to assess

$$P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E,\, \Lambda,\, I)\,, \tag{1}$$

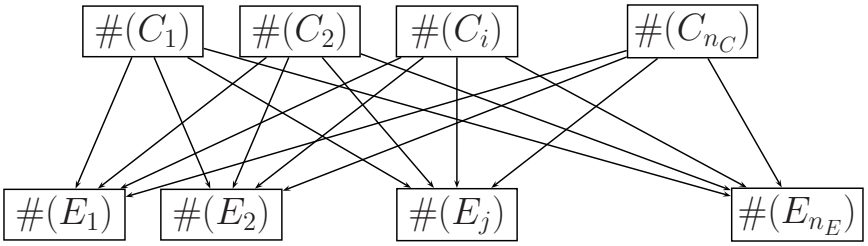

Figure 3: From the 'true' distribution (numbers of events in the cause-bins) to the observed distribution (numbers of events in the effect-bins). The number of events #( ) is indicated with $x(\,)$ in the text.

where:[5]

- $\boldsymbol{x}_C = \{x(C_1),\, x(C_2),\, \ldots,\, x(C_{n_C})\}$ is the number of events in each bin of the true distribution, i.e. *a* true spectrum.[6]
- $\boldsymbol{x}_E = \{x(E_1),\, x(E_2),\, \ldots,\, x(E_{n_E})\}$ is the observed spectrum.
- Λ stands for the smearing matrix, whose elements λ_{ji} (see remarks in footnote 4 concerning notation) are defined in probabilistic terms as

$$\lambda_{ji} \;\equiv\; P(E_j \,|\, C_i,\, I)\,. \tag{2}$$

[5] In Ref. [2] $\boldsymbol{x}_C$, $\boldsymbol{x}_E$ and Λ were respectively indicated by $\underline{n}_C$, $\underline{n}_E$ and $\boldsymbol{S}$. As in Fig. 3 and in Ref. [2], the symbols i and j are used to index causes and effects, respectively, no matter if this choice leads to the unusual convention of indexing the Λ rows by j and the columns by i, as in Eq. (2). Note that, when we refers to MC simulations to infer the smearing matrix, $\boldsymbol{x}_E$, we also need include the trash bin, as it will be reminded at the proper place. Later on [see Eq. (9)] it will be convenient to name $\boldsymbol{\lambda}_i$ the columns of the matrix Λ. In summary

$$\begin{aligned}
\Lambda \;&=\; \begin{pmatrix}
P(E_1 \,|\, C_1,\, I) & P(E_1 \,|\, C_2,\, I) & \ldots & P(E_1 \,|\, C_{n_c},\, I) \\
P(E_2 \,|\, C_1,\, I) & P(E_2 \,|\, C_2,\, I) & \ldots & P(E_2 \,|\, C_{n_c},\, I) \\
\ldots & \ldots & \ldots & \ldots \\
P(E_{n_E+1} \,|\, C_1,\, I) & P(E_{n_E+1} \,|\, C_2,\, I) & \ldots & P(E_{n_E+1} \,|\, C_{n_c},\, I)
\end{pmatrix} \\
&=\; (\boldsymbol{\lambda}_1,\, \boldsymbol{\lambda}_2, \ldots, \boldsymbol{\lambda}_{n_C})\,.
\end{aligned}$$

[6] For the use of the indefinite article in conjunction with 'true values', see the ISO *"Guide to the expression of uncertainty in measurement"* [9], according to which *a* true value is *"a value compatible with the definition of a given particular quantity."*

The knowledge of Λ comes usually from MC simulation and it is therefore affected by an uncertainty, described, in general terms, by a pdf $f(\Lambda \mid I)$.

- I stands for the state of information under which the analysis is performed (this underlying condition is often implicit in the pdf's).

Once we have stated clearly and in probabilistic terms our question (*"what is* $P(\boldsymbol{x}_C \mid \boldsymbol{x}_E,\, I)$*?"*), probability theory provides us the answer,[7] at least in principle. In fact,

[7] Instead, the idea of inverting the smearing matrix (assuming it square and not singular) is wrong in principle (i.e. besides the ascertainment that such a method yields unacceptable results). In fact, unfolding is a probabilistic problem and not a deterministic linear one (or, equivalently, a geometric problem of rotating vectors). Therefore it needs to be solved by probabilistic tools and not by linear algebra methods ('rotations'). It is easy to show that 'rotation' works only when we have an 'infinite' number of events, such that stochastic effects are negligible and the observations coincides with the expectations. In fact, each product $\lambda_{ji}\, x(C_i)$ is nothing but the expected number of events in the effect-cell E_j due to cause C_i alone:

$$\mathrm{E}[\, x(E_j)|_{x(C_i)}] \quad = \quad P(E_j \mid C_i,\, I)\, x(C_i) = \lambda_{ji}\, x(C_i)$$

Summing up the contributions from all cause-cells, we get the expected value in the effect-cell E_j

$$\mathrm{E}[\, x(E_j)|_{\boldsymbol{x}_C}] \quad = \quad \sum_i \lambda_{ji}\, x(C_i)\,,$$

that can be written in matrix form as

$$\boldsymbol{\mu}_E \equiv \mathrm{E}[\boldsymbol{x}_E] \quad = \quad \Lambda\, \boldsymbol{x}_C\,.$$

Then, if Λ is square and not singular, we get

$$\boldsymbol{x}_C \quad = \quad \Lambda^{-1}\, \boldsymbol{\mu}_E\,.$$

But this might be, at most, the solution of a text book exercise in mathematical probability theory, and does not help to solve real problems. This is the reason why, besides the fact that the matrix inversion gives notoriously bad results, *the very idea of inverting the smearing matrix is logically flawed*: we can certainly apply Λ^{-1} to a vector of numbers *already known* to be sums of expected values of binomials, but we cannot apply it to a vector of numbers that *might be* (we are not even sure of this, because some counts could be due to background we do not take into account!) sums of binomial random variables. If we do it, there is no guarantee that $\Lambda^{-1}\boldsymbol{x}_E$ yields a vector of 'valid numbers' of the n-parameters of binomials (the question that they might have the meaning of a physical spectrum for the problem under study is secondary at this level) and, in fact, even negative numbers can be obtained! It follows that unfolding methods which use the matrix inversion as starting point and try to cure its bad features with some kitchen are not appealing from a theoretical point view, although they might even provide acceptable results because of 'mysterious' reasons I do not want to enter into (cooks might be extremely clever!).

1. Bayes' theorem allows to calculate $P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E,\, \Lambda,\, I)$, given the observation $\boldsymbol{x}_E$ and the smearing matrix Λ, as

$$P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E,\, \Lambda,\, I) \;=\; \frac{P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I) \cdot P(\boldsymbol{x}_C \,|\, I)}{\sum_{\boldsymbol{x}_C} P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I) \cdot P(\boldsymbol{x}_C \,|\, I)} \tag{3}$$

   (the formula will be explained in a while).

2. We can take into account of the uncertainty about Λ using another theorem of probability theory, namely

$$P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E, I) \;=\; \int P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E, \Lambda,\, I)\, f(\Lambda \,|\, I)\, \mathrm{d}\Lambda\,. \tag{4}$$

Let us now go through the several ingredients needed to get $P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E, I)$ and try to understand were the problems arise.

- First of all, it is easy to realize that the denominator of Eq. (3) is just a normalization factor, and then we can rewrite the Bayes' formula in a way that highlights the main ingredients:

$$P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E,\, \Lambda,\, I) \;\propto\; P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I) \cdot P(\boldsymbol{x}_C \,|\, I)\,, \tag{5}$$

  where $P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I)$ is the so called *likelihood* and $P(\boldsymbol{x}_C \,|\, I)$ the *prior*. The left hand side of the Bayes' formula takes the name of *posterior*. As we see, in probabilistic inference the likelihoods have the role of updating probabilities.

- Although the presence of priors in the formula might cause anxiety in those who approach this kind of reasoning for the first time, it is a matter of fact that: 1) priors are logically necessary to get $P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E,\, \Lambda,\, I)$ starting from the likelihood, i.e. to perform the so-called 'probability inversion'; 2) they allow to plug into the model all relevant prior information, that might come from previous data or from theoretical prejudices; 3) priors are often so vague (or there are so many data – that is the same for the relative balance of prior and likelihood in the Bayes' formula) that they influence negligibly the posterior, and the inference is often dominated by the likelihood.

- Let us assume this last case of prior vagueness applies. This is equivalent to have

$$P(\boldsymbol{x}_C \,|\, I) \;=\; \textit{constant} \tag{6}$$

and the inference is then performed according to the rule

$$P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E,\, \Lambda,\, I) \quad \propto \quad P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I)\,. \tag{7}$$

It is then not a surprise that the most probable spectrum $\boldsymbol{x}_C$ is the one which maximizes the likelihood, *if* we have no other good reason to believe that this is not the case. (This is the meaning of the expression "recovering maximum likelihood estimators" from the Bayesian approach.)

- If we were able to write down a closed expression for the likelihood, or at least to provide a simple algorithmic expression of it, we could somehow scan all possible spectra $\boldsymbol{x}_C$, find the most probable one (or an 'average' spectrum that summarizes in some sense the variety of true spectra compatible with the data) and assess somehow the uncertainty about the result. But, unfortunately, this is not the case with $P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I)$ of our interest. Let us see why. Given a certain number of events in a cause-bin $x(C_i)$, the number of events in the effect-bins, included the 'trash' one, is described by a *multinomial distribution* (see Appendix A.1):

$$P(\boldsymbol{x}_E \,|\, x(C_i), \Lambda, I) \quad = \quad \frac{x(C_i)!}{\prod_j^{n_E+1} x(E_j)!} \prod_j^{n_E+1} \lambda_{ji}^{x(E_j)}\,. \tag{8}$$

It follows that $P(\boldsymbol{x}_E \,|\, \boldsymbol{x}_C,\, \Lambda,\, I)$ is the sum of independent multinomial distributions, for which, unfortunately, a closed formula does not exist.[8] This is the real serious technical problem that prevents a straight application of the Bayes' formula.

- The elements of the smearing matrix Λ are obtained by MC: we generate a large number of events in each cell C_i and count 'where they end' after a realistic simulation. Intuitively, we expect $\lambda_{ji} \approx x(E_j)^{MC}/x(C_i)^{MC}$, but we also know that this is just an estimate, with some uncertainty. Fortunately, in this case we can make direct

[8] It is well know that the sum of two Poisson variables is still a Poissonian. Similarly, the sum of two binomial variables having *the same* parameter 'p' is still a binomial. (This 'reproductive property' applies to few other distributions). But the sum of two binomial variables with different p does not have a closed expression (this can be understood either intuitively, just thinking to the meaning of a binomial, or, more formally, analyzing the product of the characteristic functions of each binomial). The same happens with a multinomial, that is just an extension of the binomial to more than two possible outcomes.

use of the Bayes' formula applied to MC events, if we model the prior using a convenient pdf. In fact, if we indicate by $\boldsymbol{\lambda}_i$ the i-th column of Λ (see footnote 4), i.e.

$$\boldsymbol{\lambda}_i = \{\lambda_{1,i},\, \lambda_{2,i}, \ldots,\, \lambda_{n_E+1,i}\}\,, \tag{9}$$

we have

$$f[\boldsymbol{\lambda}_i \,|\, \boldsymbol{x}_E^{MC},\, x(C_i)^{MC},\, I] \quad \propto \quad P[\boldsymbol{x}_E^{MC} \,|\, x(C_i)^{MC},\, \boldsymbol{\lambda}_i,\, I] \cdot f(\boldsymbol{\lambda}_i \,|\, I)\,. \tag{10}$$

Since $P[\boldsymbol{x}_E^{MC} \,|\, x(C_i)^{MC},\, \boldsymbol{\lambda}_i,\, I]$ is a multinomial, choosing a Dirichlet prior (an extension of the Beta distribution to many dimensions – see Appendix A.2), we get a Dirichlet posterior (Appendix A.3), i.e.

$$\boldsymbol{\lambda}_i \quad \sim \quad \mathrm{Dir}(\boldsymbol{\alpha}_{posterior_i})\,, \tag{11}$$

with

$$\boldsymbol{\alpha}_{posterior_i} \quad = \quad \boldsymbol{\alpha}_{prior_i} + \boldsymbol{x}_E^{MC}\big|_{x(C_i)}\,. \tag{12}$$

(Hereafter the symbols '$\sim$' stands for 'follows the probability distribution'.)

- Finally there is the integral (4), which is in principle not a problem (for example, thanks to modern technologies, it can easily performed by MC methods).

In conclusion, the main serious problem is $P(\boldsymbol{x}_C \,|\, \boldsymbol{x}_E, \Lambda,\, I)$. Thus being unable to tackle the problem from the main door, we need some 'tricks' to circumvent the obstacle, but still under the guidance of probability theory.

3 Practical algorithm to perform an independent-bin Bayesian unfolding

The basic trick to avoid the mentioned difficulty is to apply Bayes' theorem to causes and effects, instead than to true and observed spectrum. In practice, instead of using of Eq. (3) we start from

$$P(C_i|E_j,\, I) \quad = \quad \frac{P(E_j \,|\, C_i,\, I) \cdot P(C_i \,|\, I)}{\sum_i P(E_j \,|\, C_i,\, I) \cdot P(C_i \,|\, I)}\,, \tag{13}$$

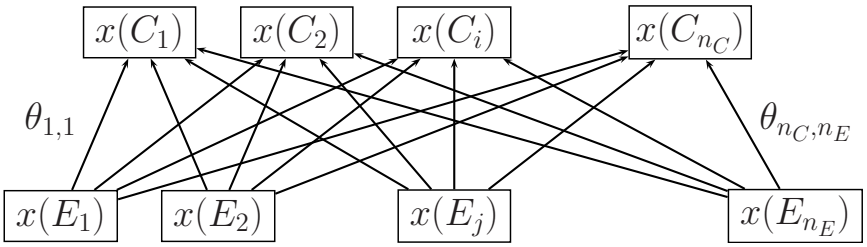

Figure 4: *Sharing counts observed in effect-cells among cause-cells according to* $\theta_{ij} = P(C_i \,|\, E_j,\, I)$.

or

$$\theta_{ij} \;\; = \;\; \frac{\lambda_{ji} \cdot P(C_i \,|\, I)}{\sum_i \lambda_{ji} \cdot P(C_i \,|\, I)}\,, \tag{14}$$

having defined $\theta_{ij} \equiv P(C_i \,|E_j,\, I)$ in analogy to $\lambda_{ji} \equiv P(E_j \,|\, C_i,\, I)$.

At this point a very important remark is in order. The prior in Eqs. (13)-(14) has a different meaning from that of Eq. (3). In Eq. (3) $P(\boldsymbol{x}_C \,|\, I)$ assigns different probabilities to all possible spectra. Instead, in Eqs. (13)-(14) $P(C_i \,|\, I)$ stands for *a single spectrum* (more precisely, all spectra that differ from each other just by normalization). This can be better understood analyzing the meaning of 'uniform' (or 'flat') referred to $P(\boldsymbol{x}_C \,|\, I)$ and to $P(C_i \,|\, I)$.

- $P(\boldsymbol{x}_C \,|\, I) =$ *constant* means *all spectra* are equally likely.
- $P(C_i \,|\, I) =$ *constant* means we consider the cause bins equally likely, i.e. the prior assess an initial belief in *flat spectra*.

In other words, while a flat $P(\boldsymbol{x}_C \,|\, I)$ means indifference about all possible spectra in order to 'let the data speak by themselves', a flat $P(C_i \,|\, I)$ is a strong assumption that usually does not correspond to our priors concerning the physics case. This implies that we have to tune somehow the algorithm in order to take into account of this gross approximation. We will come back to this issue in Sec. 4.

At this point, having evaluated $P(C_i \,|\, E_j,\, I)$, we can use it to share the counts observed in each effect-bin among all cause-bins (see Fig. 4). The estimate of the true spectrum is obtained repeating this sharing for all observed bins and taking into account inefficiency. An uncertainty on the unfolded spectrum is also evaluated. Let us see how all this was done in the old algorithm and how it has been improved.

3.1 Old algorithm (Ref. [2])

In the old algorithm the share of observed events among the causes was done only considering expectations and applying an inefficiency correction, going through the following steps:

- number of counts in C_i due to the observation in E_j:

$$x(C_i)|_{x(E_j)} \approx P(C_i\,|E_j,\,I)\cdot x(E_j) = \theta_{ij}\cdot x(E_j)\,; \tag{15}$$

- number of counts in C_i due to all observations:

$$x(C_i)|_{\boldsymbol{x}_E} \approx \sum_{j=1}^{n_E} P(C_i\,|E_j,\,I)\cdot x(E_j) = \sum_{j=1}^{n_E}\theta_{ij}\cdot x(E_j)\,; \tag{16}$$

- number of counts in C_i that also takes into account efficiency (ϵ_i):

$$x(C_i) \approx \frac{1}{\epsilon_i}\; x(C_i)|_{\boldsymbol{x}_E} = \frac{1}{\epsilon_i}\sum_{j=1}^{n_E} P(C_i\,|E_j,\,I)\cdot x(E_j) \tag{17}$$

$$= \frac{1}{\epsilon_i}\sum_{j=1}^{n_E}\theta_{ij}\cdot x(E_j) \tag{18}$$

  where

$$\epsilon_i = \sum_{j=1}^{n_E} P(E_j\,|\,C_i,\,I) = \sum_{j=1}^{n_E}\lambda_{ji} = 1-\lambda_{n_E+1,\,i}\,. \tag{19}$$

- Finally, we have to remember that the several λ_{ij} were estimated by MC simulation as

$$\lambda_{ji} \approx x(E_j)^{MC}/x(C_i)^{MC}\,, \tag{20}$$

  from which ϵ_i and θ_{ij} are calculated according to Eqs. (19) and (14).

Essentially, $x(C_i)$ of Eq. (17) was considered an 'estimator', whose 'error' was calculated by standard 'error propagation' (this was a pragmatic compromise between 'standard methods' I was accustomed at that time and the Bayesian approach I was in the process of learning – anyhow, not bad to begin, and not worst than other methods.)

At this point, there is still open the issue of the inappropriate priors, that we have encountered at the beginning of this section. Since this question remains in the improved algorithm, it will be discussed later in Sec. 4.

3.2 Improvements

As it is easy to understand from the previous description, weak points of the old algorithm are the treatment of small numbers and the calculation of uncertainty by 'error propagation formulas'. Let us see how we can improve them.

- Instead of just 'estimate' λ_{ij} as $x(E_j)^{MC}/x(C_i)^{MC}$, we can model their pdf by a Dirichlet (see Appendix A):

$$\boldsymbol{\lambda}_i \;\sim\; \mathrm{Dir}[\boldsymbol{\alpha}_{prior} \,+\, \boldsymbol{x}_E^{MC}\big|_{x(C_i)^{MC}}]\,, \tag{21}$$

  where

  - $\boldsymbol{\alpha}_{prior}$ is the initial set of Dirichlet parameters – typically unitary for uniform priors (see Appendix B for details);
  - $\boldsymbol{x}_E^{MC}\big|_{x(C_i)^{MC}}$ stands for the numbers of MC counts generated in C_i and which end in each of *all* possible $n_E + 1$ 'effects' (i.e. we also have to consider the inefficiency bin – see Fig. 2 – in order to satisfy the condition $\sum_i \lambda_i = 1$ of the Dirichlet variables).

- The uncertainty about Λ is taken into account by *sampling* its values, that is equivalent to perform the integral (4).[9] Therefore, for each sampling,

  - the smearing matrix elements λ_{ji} are extracted according to a Dirichlet distribution;
  - the efficiencies $\epsilon_i = \sum_{j=1}^{n_E} \lambda_{ji}$ are calculated;
  - the inverse probabilities θ_{ij} are obtained from the Bayes formula.

- The sharing of the number of counts $x(E_j)$ observed in an effect-cell among all cause-bins is performed using a multinomial distribution (see Appendix A.1):

$$\boldsymbol{x}_C|_{x(E_j)} \;\sim\; \mathrm{Mult}[x(E_j),\, \boldsymbol{\theta}_j]\,, \tag{22}$$

  with $\boldsymbol{\theta}_j = \{\theta_{1,j},\, \theta_{2,j},\, \ldots,\, \theta_{n_C,j},\, \}$.

[9] Note that if we are in doubt about several smearing matrices, because e.g. they are obtained from different parameters of the simulation, the ('systematic') effect of this uncertainty can be taken into account sampling from the different Λ's, with weights depending on our confidence on each of them.

- The contributions of all *observed* effect-bins are summed up, i.e.

$$\boldsymbol{x}_C|_{\boldsymbol{x}_E} \;=\; \sum_{j=1}^{n_E} \boldsymbol{x}_C|_{x(E_j)} \;. \tag{23}$$

  Inefficiency are taken into account as in the old algorithm:

$$\boldsymbol{x}_C \;=\; \frac{\boldsymbol{x}_C|_{\boldsymbol{x}_E}}{\boldsymbol{\epsilon}} \,, \tag{24}$$

  where the ratio in the r.h.s is done component by component.

However, we need to remember that the observed number of events in each effect bin, $x(E_j)$, comes from a Poisson distribution with unknown parameter μ_j, whose inference can be conveniently performed starting from a conjugate gamma distribution (see e.g. Ref. [6]). It follows that

$$\mu_j \;\sim\; \text{Gamma}[c_j + x(E_j),\, r_j + 1] \,, \tag{25}$$

where c_j and r_j are the initial parameters of the gamma. [The case of a flat prior corresponds to $c_j = 1$ and $r_j \to 0$, i.e. an exponential with vanishing rate.]

Therefore, what should be shared among the cause-bins is not $x(E_j)$ but μ_j, extracted in each sampling according to Eq. (25). However, μ_j is a continuous variable and therefore cannot be used as 'trials parameter' of a multinomial distribution. Nevertheless, fractional events can indeed be shared making use of a little trick:

- μ_j is extracted according to Eq. (25) and it is rounded to the closest *positive* integer, that we indicate here by m_j: m_j will be the number of trials used in the multinomial random generator.

- $\boldsymbol{x}_C|_{m_j}$ is extracted and then rescaled by the factor μ_j/m_j, i.e.

$$\boldsymbol{x}_C|_{m_j} \;\sim\; \text{Mult}(m_j,\, \boldsymbol{\theta}_j) \,, \tag{26}$$

$$\boldsymbol{x}_C|_{\mu_j} \;=\; \frac{\mu_j}{m_j}\; \boldsymbol{x}_C|_{m_j} \;. \tag{27}$$

This means that Eq. (23) has to be replaced by

$$\boldsymbol{x}_C|_{\boldsymbol{x}_E} \;=\; \sum_{j=1}^{n_E} \boldsymbol{x}_C|_{\mu_j} \;; \tag{28}$$

then the inefficiency corrections are applied as usual.

Each sampling provides a spectrum $\boldsymbol{x}_C^{(t)}$, where t is the 'time' index of the sampling. After N samplings we can calculate an average spectrum, variance and covariance for each cause-bin, as well as covariances among bins, and any other statistical summary; or we can inspect graphically each $x(C_i)$.

This algorithm has been implemented in the R language [4] and it is freely available on the web [5]. The R language has been chosen because it is a powerful scripting language,[10] open source, multi-platform, oriented towards statistics applications, well maintained and with tons of contributed extension packages. In practice a kind of programming *lingua franca*.

In order to test the algorithm, also a simple event generator has been included, that reproduces the toy models of Ref. [2]. Some results will be presented in Sec. 5

4 Iterations and smoothing

Let us now go back to the issue of the priors, that we have left open from the beginning of the previous section. The crucial thing to understand is that, as we stated above, *instead of using a prior flat over the possible spectra, we are using a particular, flat spectrum as prior*. Therefore, the posterior [i.e. the ensemble of $\boldsymbol{x}_C^{(t)}$ obtained by sampling] is affected by this quite strong assumption, that seldom holds in real cases.

Obviously, the first idea a practical physicist has, in order to fix the problem, is to do some *fine tuning*. In fact, simulations on toy models show that the unfolded distribution reproduces rather well the true one, even with a flat spectrum as prior. However, it still 'remembers the flat prior'. This effect can be cured iterating the procedure, using the posterior as prior in a

[10] I take the opportunity to make a point about the teaching of scripting languages, and in particular of *R*, in the physics courses. Surely physics students might need to learn C (and perhaps C++) at some point, but it is a matter of fact that writing some C code to solve little/medium problems that occur everyday, and that might also need some graphics, is a pain, not only for students. The consequence is that students usually do not write little programs in C to solve the problems they meet in the general physics or laboratory courses, continuing to use, instead, spread-sheets, learned in high school (and forget real programming until they need it for their theses, if they ever will need it). In my opinion, teaching a scripting language like R, perhaps in parallel to C, would be a good solution. (but if I really had to chose, I would opt for R, as the language to begin). (Personally, since I have discovered scripting languages, I use C only for though professional tasks. Indeed, I have even almost abandoned *Mathematica*, unless I really need symbolic calculation.)

subsequent unfolding. Empirically we learn that, in 'normal' cases, just two or three steps are sufficient to recover quite accurately the true spectrum.

As it has been discussed in length in Ref. [2], we cannot repeat the iterations for a very long time, otherwise there will be a kind of positive feedback, driven by fluctuations, that makes the asymptotically unfolded spectrum 'crazy' (it simply means that it is far away from all reasonable expectations). Again, as we understand it, it is a question of judging the outcome by our physics priors,[11] that, although rather vague, tell us that, given the kind of measurement, a spectrum with wild oscillations is far from our beliefs. This kind of problem happens also with other kind of unfolding algorithms and it is usually cured by *regularization*.

Regularization is based on the subjective scientific prejudice that 'wild' distributions are not physical (and, as the reader knows or might imagine, this very reasonable subjective ingredient is unavoidably also embedded in all non-Bayesian methods). In practice, regularization can be implemented constraining adjacent bins not to be 'too discontinuous', for example limiting the absolute value of the derivatives calculated numerically on the unfolded spectrum. Or one might fit the spectrum with some orthonormal functions, but dumping 'high frequencies'. And so on.

I do not have a strong opinion on the matter. My recommendation is that one should know what one is doing. My preferred regularization method for one-dimensional problems consists in smoothing the posterior before injecting it as prior of the next iteration (but not after the last step!). This smoothing can be performed by a low order polynomial fit, as it is implemented in the demo scripts accompanying the program. I find this technique quite simple, well consistent with the spirit of this unfolding method and with the reasons to go through the iterations. Moreover, it is easy to understand that the procedure unfolding-smoothing-unfolding converges rapidly.

In conclusion, although the idea of iteration does not seem very Bayesian (we cannot squeeze the same data twice!), it is in fact just an adaptive way to recover what could have been obtained by a more reasonable uniform prior over the possible spectra. (Stated with other words, it would be non-Bayesian to stick to the first iteration, obtained from a prior that rarely corresponds to physics spectra!) For this reason the idea of iteration has been kept, although some effort has been done to get rid of it. In fact, in the new code the user can, instead of providing a (typically flat) spectrum as prior, set up a function [`priorf()`] that returns a normalized spectrum

[11] Good physicists do have priors and always use them! (Only the perfect idiot has no priors.)

chosen at random according to a probability distribution describing prior knowledge. In this way, in each sampling also the prior is sampled. However, I anticipate that implementing the function `priorf()` might not be that trivial, and iterative procedure still remains a pragmatic compromise to slalom among the difficulties.

5 Results on toy models

The program has been checked with the same toy smearing matrices used in Ref. [2], reproduced here in Tab. 1. The results of the simulations are given in Fig. 5. The generated distributions are shown in black. The 'measured' distributions (in red) look completely different from the 'true', due to the very severe smearing matrices used. The figures also show the results after the first iteration (light blue) and the intermediate ones (yellow). Note that, although in all simulations twenty iterations have been performed, only a few intermediate iterations are visible, because the others overlap with the final step. This allows you to get a feeling about the speed of convergence of the algorithm, depending on the difficulty of the problem – let us remind that, thanks to the intermediate smoothing the algorithm *does* converge.

Other independent simulations are reported in Figs. 6 and 7.

6 Conclusions

The evaluation of uncertainties of the Bayesian unfolding of Ref. [2] has been improved taking the probability density functions of the quantities of interest. Simplifications are achieved using prior conjugates and performing the relevant integrations, needed to propagate uncertainties, by sampling.

The paper also discusses the issue of the iterations and the role of the *intermediate* smoothing to reach fast convergence. It is important to note that, contrary to other algorithms, the intermediate smoothing acts as a *regularization on the priors* and not on the unfolded spectrum. For this reason physical peaks should still appear in the unfolded spectrum.

The R code of the algorithm is available, together with a little simulation script to run it on toy models. This script also implements of simple intermediate prior regularization, but the reader should have clear in mind that this task is left to the user, who is supposed to understand well his/her physics case.

Smearing 1

	C_1	C_2	C_3	C_4	C_5	C_6	C_7	C_8	C_9	C_{10}
E_1	0.010	0	0	0	0	0	0	0	0	0
E_2	0.025	0	0	0	0	0	0	0	0	0
E_3	0.040	0	0	0	0	0	0	0	0	0
E_4	0.025	0.050	0	0	0	0	0	0	0	0
E_5	0	0.100	0	0	0	0	0	0	0	0
E_6	0	0.050	0.120	0	0	0	0	0	0	0
E_7	0	0	0.120	0	0	0	0	0	0	0
E_8	0	0	0.060	0.160	0.025	0	0	0	0	0
E_9	0	0	0	0.160	0.100	0.120	0.140	0	0	0
E_{10}	0	0	0	0.080	0.250	0.360	0.280	0.200	0	0
E_{11}	0	0	0	0	0.125	0.120	0.280	0.400	0.225	0
E_{12}	0	0	0	0	0	0	0	0.200	0.450	0.250
E_{13}	0	0	0	0	0	0	0	0	0.225	0.500
E_{14}	0	0	0	0	0	0	0	0	0	0.250
ϵ	*0.1*	*0.2*	*0.3*	*0.4*	*0.5*	*0.6*	*0.7*	*0.8*	*0.9*	*1*

Smearing 2

	C_1	C_2	C_3	C_4	C_5	C_6	C_7	C_8	C_9	C_{10}
E_1	0.045	0	0	0	0	0	0	0	0	0.225
E_2	0.090	0.054	0	0	0	0	0	0	0.225	0.405
E_3	0.045	0.072	0.045	0	0	0	0	0.225	0.360	0.225
E_4	0	0.054	0.045	0.045	0.225	0.180	0.360	0.405	0.225	0
E_5	0	0	0	0.180	0.360	0.450	0.315	0.225	0	0
E_6	0	0	0	0.315	0.180	0.180	0.180	0	0.045	0
E_7	0	0	0.135	0.270	0.045	0	0	0.045	0	0
E_8	0	0.045	0.270	0.045	0	0	0.045	0	0.045	0
E_9	0	0.135	0.270	0	0.045	0.045	0	0	0	0.045
E_{10}	0	0.360	0.045	0	0.045	0.045	0	0	0	0
E_{11}	0.180	0.180	0	0	0	0	0	0	0	0
E_{12}	0.270	0	0	0.045	0	0	0	0	0	0
E_{13}	0.270	0	0.045	0	0	0	0	0	0	0
E_{14}	0.090	0	0.045	0	0	0	0	0	0	0
ϵ	*0.9*	*0.9*	*0.9*	*0.9*	*0.9*	*0.9*	*0.9*	*0.9*	*0.9*	*0.9*

Table 1: *Smearing matrices $P(E_j \mid C_i)$ of the toy models*

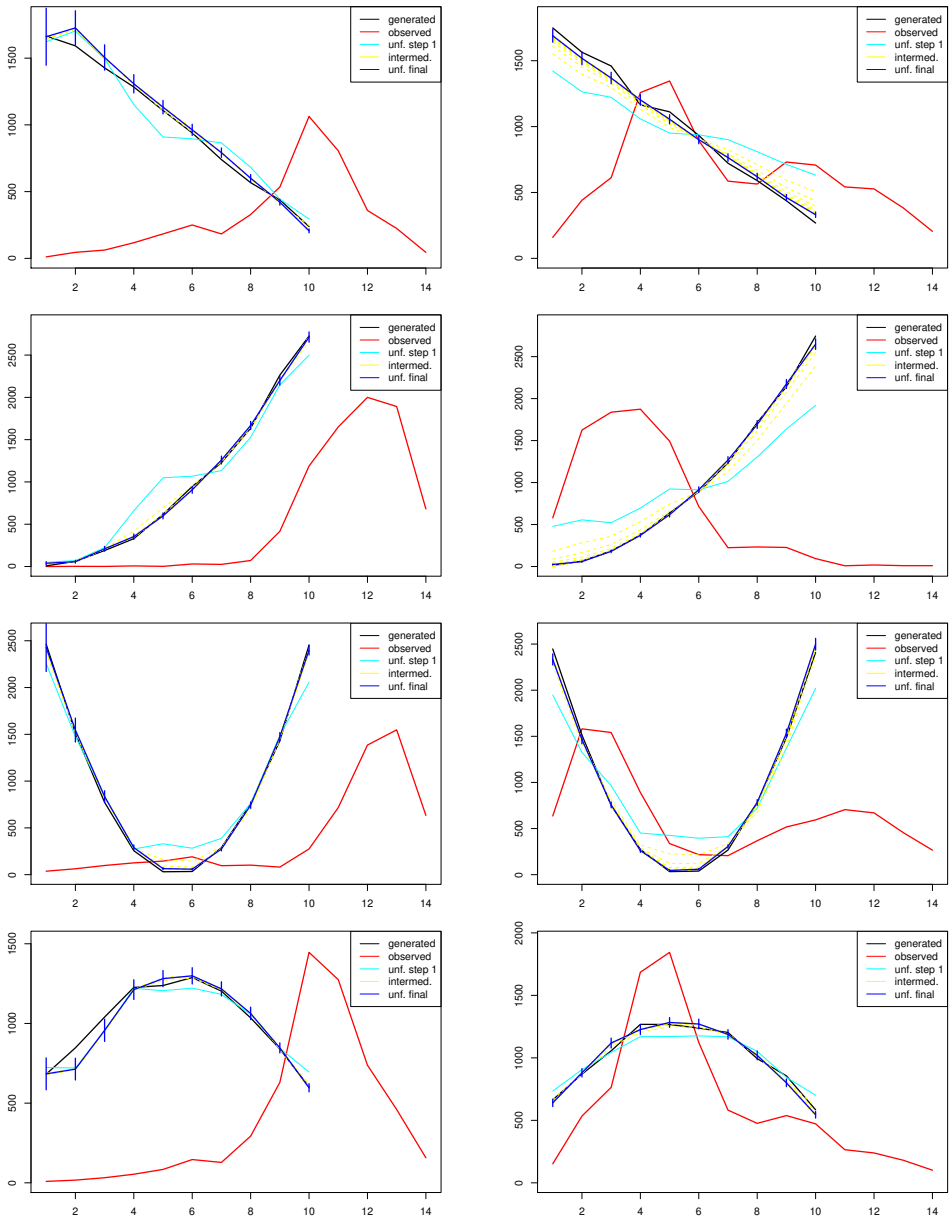

Figure 5: *Result of unfolding applied to four toy models* (data.func *parameter from 1 to 4 in the demo program) and two different smearing matrices (left and right figures correspond, respectively, to 'Smearing 1' and 'Smearing 2' of Tab. 1.).*

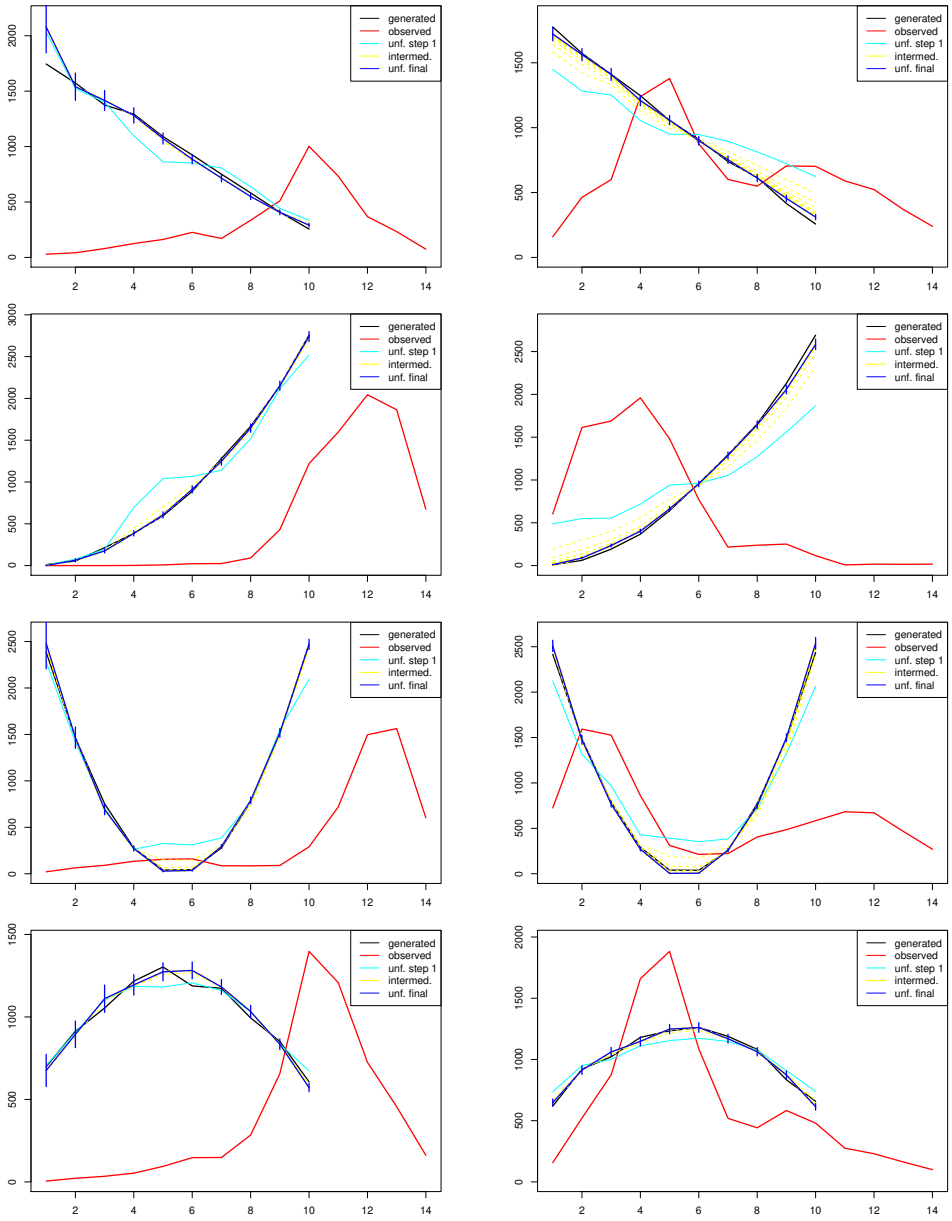

Figure 6: *Same as Fig. 5. Independent complete simulation.*

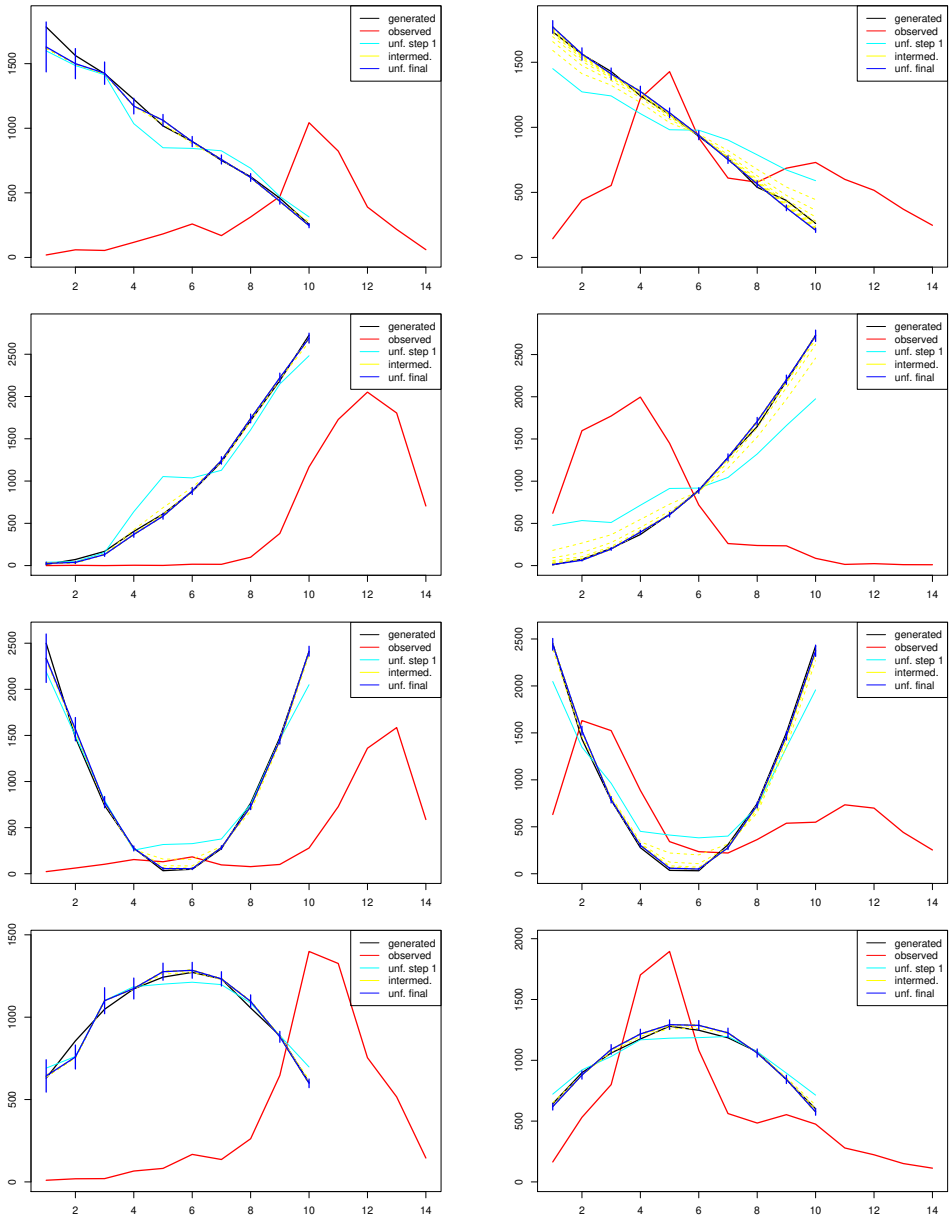

Figure 7: *Same as Fig. 5. Independent complete simulation.*

References

[1] G. D'Agostini, *Fits, and especially linear fits, with errors on both axes, extra variance of the data points and other complications*, physics/0511182.

[2] G. D'Agostini, *A multidimensional unfolding method based on Bayes' theorem*, Meth. in Phys. Res. **A362** (1995) 487.

[3] G. D'Agostini, *Asymmetric uncertainties: sources, treatment and possible dangers*, physics/0403086.

[4] R Development Core Team (2009) *R: A language and environment for statistical computing. R Foundation for Statistical Computing*, Vienna, Austria. ISBN 3-900051-07-0, `http://www.R-project.org`.

[5] `http://www.roma1.infn.it/~dagos/prob+stat.html#unf2`.

[6] G. D'Agostini, *Bayesian reasoning in data analysis: A critical introduction*, World Scientific Publishing, 2003.

[7] G. D'Agostini, *Bayesian inference in processing experimental data: principles and basic applications*, Rep. Prog. Phys. 66 (2003)1383, physics/0304102.
V. Dose, *Bayesian inference in physics: case studies*, Rep. Prog.. Phys. 66 (2003)1421.

[8] M. Raso, *A measurement of the σ_{ep} for $Q^2 < 0.3\ GeV^2$, as a by product of the F_2^{ep} measurement*, ZEUS-98-075, October 1998 (available at `http://www.roma1.infn.it/~dagos/prob+stat.html#Raso`)

[9] International Organization for Standardization (ISO), *"Guide to the expression of uncertainty in measurement"*, Geneva, Switzerland, 1993.

[10] M. Nakao, *Measurement of the proton structure function F_2 at HERA*, PhD Thesis, University of Tokyo, INS-J-178, 1994.
Scanned version available at
`http://ccdb4fs.kek.jp/cgi-bin/img_index?9501307`.

Appendix A – Multinomial and Dirichlet distributions

A.1 Multinomial distribution

The Multinomial distribution is the extension of the Binomial from 2 to k possible outcomes. If we assign the probability p_i to the i-th outcome and think of n trials, then we are interested in the joint probability to observe x_1 times the outcome 1, x_2 times the outcome 2, and so on. From basic rules of probability theory and some combinatorics we get

$$f[\,\boldsymbol{x}\,|\,\mathrm{Mult}(n,\,\boldsymbol{p})\,] \;=\; \frac{n!}{x_1!\,x_2!\cdots x_k!}\,p_1^{x_1}p_2^{x_2}\cdots p_k^{x_k}\,, \tag{29}$$

with

$$\begin{aligned}\boldsymbol{x} &= \{x_1, x_2, \ldots, x_k\} \qquad [\,x_i = 0, 1, \ldots, n;\;\; \sum_i x_i = n\,]\\ \boldsymbol{p} &= \{p_1, p_2, \ldots, p_k\} \qquad [\,0 < p < 1;\;\; \sum_i p_i = 1\,]\,.\end{aligned}$$

It is easy to see that the binomial distribution is recovered for $k = 2$. Note that, given the constraint $\sum_i x_i = n$, the multinomial has to be seen as a distribution for $k-1$ variables $x_1, x_2, \ldots, x_{k-1}$, as it is clear from the binomial case. Expected values, variances, covariances and correlation coefficients are given by

$$\mathrm{E}(X_i) \;=\; n\,p_i \tag{30}$$

$$\mathrm{Var}(X_i) \;=\; n\,p_i\,(1-p_i) \tag{31}$$

$$\mathrm{Cov}(X_i, X_j) \;=\; -n\,p_i\,p_j \qquad [i \neq j] \tag{32}$$

$$\rho(X_i, X_j) \;=\; \frac{-p_i\,p_j}{\sqrt{p_i\,(1-p_i)\,p_j\,(1-p_j)}} \qquad [i \neq j]\,. \tag{33}$$

Each marginal, i.e. $f[x_i\,|\,\mathrm{Mult}(n,\,\boldsymbol{p})]$, is a binomial with parameters n and p_i:

$$f[x_i\,|\,\mathrm{Mult}(n,\,\boldsymbol{p})] \;=\; \frac{n!}{x_i!\,(n-x_i)}\,p_i^{x_i}\,(1-p_i)^{n-x_i}\,. \tag{34}$$

For $k = 2$ the two variables X_1 and X_2 are 100% anti-correlated, as it is quite obvious (the knowledge of X_1 determines X_2, and vice versa). In fact, since $p_2 = 1 - p_1$, we get from Eq. (33)

$$\rho(X_1, X_2\,|\,k=2) \;=\; \frac{-p_1\,(1-p_1)}{\sqrt{p_1\,(1-p_1)\,(1-p_1)\,p_1}} = -1\,. \tag{35}$$

A.2 Dirichlet distribution

In this case we start from the formal definition. Given k continuous variables whose value can range between zero and one and that sum up to unity, the Dirichlet distribution is defined by the following joint pdf:

$$f[\,\boldsymbol{x}\,|\,\mathrm{Dir}(\boldsymbol{\alpha})\,] \;=\; \frac{\Gamma(\alpha_1+\alpha_2+\cdots+\alpha_k)}{\Gamma(\alpha_1)\,\Gamma(\alpha_2)\cdots\Gamma(\alpha_k)}\,x_1^{\alpha_1-1}\,x_2^{\alpha_2-1}\,\ldots\,x_k^{\alpha_k-1} \tag{36}$$

with

$$\begin{aligned}\boldsymbol{x} &= \{x_1,x_2,\,\ldots,\,x_k\} \qquad [\,0<x_i<1;\;\sum_i x_i=1\,]\\ \boldsymbol{\alpha} &= \{\alpha_1,\alpha_2,\,\ldots,\,\alpha_k\} \qquad [\,\alpha_i>0\,]\,,\end{aligned}$$

where $\Gamma(\,)$ is the usual gamma function. If $k=2$, then the Beta is recovered. Also in this case there is a constraint on the variables, $\sum_{i=1}^{k} x_i = 1$, that makes the distribution in fact $(k-1)$-dimensional. Expected values, variances and covariances are

$$\mathrm{E}(X_i) \;=\; \frac{\alpha_i}{\alpha} \tag{37}$$

$$\mathrm{Var}(X_i) \;=\; \frac{\alpha_i\,(\alpha-\alpha_i)}{\alpha^2\,(\alpha+1)} \tag{38}$$

$$\mathrm{Cov}(X_i,X_j) \;=\; \frac{\alpha_i\,\alpha_j}{\alpha^2\,(\alpha+1)} \qquad [i\neq j]\,, \tag{39}$$

where $\alpha = \sum_i \alpha_i$. Each marginal is a Beta with parameters $r=\alpha_i$ and $s=\alpha-\alpha_i$:

$$f[x_i\,|\,\mathrm{Beta}(r=\alpha_i,s=\alpha-\alpha_i)] \;=\; \frac{x_i^{\alpha_i-1}(1-x_i)^{\alpha-\alpha_i-1}}{\beta(\alpha_i,\,\alpha-\alpha_i)}, \tag{40}$$

where $\beta()$ is the beta function (hence the name to the distribution), defined as

$$\beta(r,s) = \int_0^1 z^{r-1}(1-z)^{s-1}\,\mathrm{d}z \;\left[=\frac{\Gamma(r)\,\Gamma(s)}{\Gamma(r+s)}\right]. \tag{41}$$

Note that the Beta pdf vanishes for $x_i=0$ if $r>1$, and for $x_i=1$ if $s>1$, respectively. This observation will be important for some considerations concerning the treatment of bins with zero counts (see Appendix B.1).

A.3 Dirichlet as prior conjugate of the multinomial

If the data sample $\boldsymbol{x}$, modelled by a multinomial distribution, has been observed in n trials and we are interested in inferring the multinomial parameters $\boldsymbol{p}$, applying Bayes' theorem we get

$$f[\,\boldsymbol{p}\,|\,\mathrm{Mult}(n),\,\boldsymbol{x}\,] \;\propto\; f[\,\boldsymbol{x}\,|\,\mathrm{Mult}(n,\,\boldsymbol{p})\,]\,\cdot\,f(\boldsymbol{p}) \tag{42}$$

$$\propto\; p_1^{x_1} p_2^{x_2}\cdots p_k^{x_k}\,\cdot\,f(\boldsymbol{p})\,, \tag{43}$$

having absorbed in the normalization all factors not depending on $\boldsymbol{p}$. If we model the prior $f(\boldsymbol{p})$ with a Dirichlet, i.e. $f(\boldsymbol{p}) \propto p_1^{\alpha_1-1}\,p_2^{\alpha_2-1}\,\ldots\,p_k^{\alpha_k-1}$, we get

$$f[\,\boldsymbol{p}\,|\,\mathrm{Mult}(n,\,\boldsymbol{x})\,] \;\propto\; (\,p_1^{x_1} p_2^{x_2}\cdots p_k^{x_k}\,)\,\cdot\,(\,p_1^{\alpha_1-1}\,p_2^{\alpha_2-1}\,\ldots\,p_k^{\alpha_k-1}\,) \tag{44}$$

$$\propto\; p_1^{\alpha_1+x_1-1}\,p_2^{\alpha_2+x_2-1}\,\ldots\,p_k^{\alpha_k+x_k-1}\,, \tag{45}$$

that is still a Dirichlet. In practice, the observation $\boldsymbol{x}$ updates the Dirichlet parameters according to the rule

$$\boldsymbol{\alpha}_{posterior} \;=\; \boldsymbol{\alpha}_{prior} + \boldsymbol{x}\,. \tag{46}$$

For this reason the Dirichlet is known as the *prior conjugate* of the multinomial. In the simplest case of $k = 2$ we have the Beta prior conjugate of the Binomial (see e.g. Ref. [6]). In particular, we can easily see that a flat prior is recovered if all α's of the Dirichlet prior are equal to one.

Appendix B – Handling the zeros

Bins with null counts might occur either in the Monte Carlo results used to infer the smearing matrix, or in the observed spectrum. As it is true that when we deal with large numbers we can neglect fluctuations and uncertainties, while small numbers are somewhat problematic, it is also true that handling zeros is particularly challenging and a physicist should mistrust abstract mathematical arguments of any kind. In fact, when dealing with zeros, subjective prior knowledge becomes crucial, because we need to have an idea of 'what a zero might mean': for some bins we could reasonably think that, if we would repeat the same experiment or the same simulation, the zero might turn into one count, or even in two o three; for other bins, instead, we are quite confident that we would need to repeat the experiment or the simulation many many times before a zero turns into one count, or

perhaps we have good reasons to believe this will never happen. In other words, experienced physicists think that 'not all zeros are the same'.

Take, for example, an histogram in which, after some bins with positive (and perhaps quite large) counts, suddenly there is a row of empty bins (like the second row of Tab. 2). Intuitively we do not think all the empty bins are equivalent, although for a mathematician their are. But we are physicists, and we have formed some opinion about how nature behaves. As a consequence, we do not only tend to smooth the positive numbers of the histogram, but we also tend to make some extrapolations to the region of zero counts, and judge the far zeros to be 'more zero' that the near ones. We need then to model somehow our opinion based on past experience.

In the following subsection we shall see, in a one-dimensional case, how it is possible to treat in a consistent way zeros that occur in the MC simulation needed to evaluate the smearing matrix, where it might be really relevant.[12] Indeed, as discussed in section 4 about smoothing, this is a task left to the user, which knows the physical meaning of the bins.[13]

B.1 – Zero counts in effect-bins of MC events used to infer smearing matrix

When we make Monte Carlo simulations in order to infer the smearing matrix we usually observe many zeros, because long range migrations are usually rather rare in typical detectors.[14]

[12] Indeed, as it it is common experience, 'row' smearing matrices (i.e. not yet expressed in terms of Dirichlet parameters) evaluated from MC simulations have many bins with zero counts (in most 'sane' situations most entries are null). Instead, there is not such a similar problem in the observed spectrum, where the best pragmatic solution is re-binning, in order to avoid zeros and even very small numbers. However, it is possible to imagine cases in which one is interested to keep one or a few bins separated from the others, even no counts have been observed in it/them. These special bins can be treated *ad hoc*, in a way similar to that discussed in the following subsection, playing on the parameters of the relevant gamma prior.

[13] But don't worry: the program gives you the option of treat the zeros as ... zeros, and you might want to skip the following two subsections (that might be a bit 'academic' and of little practical relevance in most cases). But you might want to reflect a while that, perhaps, considering, democratically, all zeros on equal foot might be against your beliefs, while probably you do not want to be incoherent...

[14] Nevertheless, when we are interested in reconstructing physics quantities computed on an ensemble of tracks and clusters, and the detector has not negligible inefficiencies, the migrations can be huge, has it was used to happen in 'two-photon physics' in e^+-e^- colliders, or in some kinematical regions of deep inelastic scattering, as impressively shown in Fig. 6.8 of page 77 of Ref. [10] (the choice of this reference is that similar figures, rather popular at the beginning of HERA physics, are usually only available on printed notes

As explained in Sec. 2, the columns of the smearing matrix are jointly described by vectors of Dirichlet random variables (one vector per cause-bin). The parameters of a Dirichlet pdf are updated according to Eq. (46), that we rewrite here:

$$\boldsymbol{\alpha}_{posterior} \;=\; \boldsymbol{\alpha}_{prior} + \boldsymbol{x}\,,$$

where $\boldsymbol{x}$ stands here for $\boldsymbol{x}_E^{MC}\big|_{x(C_i)}$ [see Eq. (12)]. A usual choice of the Dirichlet prior to model indifference is to set all α's to 1. The prior is then characterized [see Eqs. (40), (37) and (38)] by the the following marginals, expectations and variances (we use the symbol p for the variables, for their physical meaning in our contest, and the index j for numbering them, as we do in the text for the effect-bins):

$$f[p_j \,|\, \mathrm{Beta}(1,\, \nu-1)] \;=\; \frac{(1-p_j)^{\nu-2}}{\beta(1,\, \nu-1)} \tag{47}$$

$$\mathrm{E}(p_j) \;=\; \frac{1}{\nu} \tag{48}$$

$$\mathrm{Var}(p_j) \;=\; \frac{\nu-1}{\nu^2\,(\nu+1)} \xrightarrow[\nu\gg 1]{} \frac{1}{\nu^2}\,, \tag{49}$$

where, hereafter, ν is used in place of n_E+1 in order to improve the readability of the formulas. In the case of $\nu=2$ we recover a flat prior and hence an expectation of $1/2$.

If we throw n MC events in order to estimate the several p_j, the updated α parameters become $\alpha_j = x_j + 1$. For the marginals we have now:

$$f[p_j \,|\, \mathrm{Beta}(x_j+1,\, n+\nu-x_j-1)] \;=\; \frac{p^{x_j}\,(1-p_j)^{n+\nu-x_j-2}}{\beta(x_j+1,\, n+\nu-x_j-1)}\,, \tag{50}$$

with expected values and variances

$$\mathrm{E}(p_j) \;=\; \frac{x_j+1}{n+\nu} \tag{51}$$

$$\mathrm{Var}(p_j) \;=\; \frac{(x_j+1)\,(n+\nu-x_j-1)}{(n+\nu)^2\,(n+\nu+1)}\,. \tag{52}$$

In the limit of large numbers ($n \gg \nu$ and $x_j \gg 1$) we get an expected value equal to x_j/n (that is the fraction of events in the cell j, in agreement with a naive evaluation of the elements of the smearing matrix), with a variance of $(x_j/n)\,(1-x_j/n)\,/\,n$.

and conference proceedings, of difficult accessibility).

In the case *zero counts and large* n the pdf is given by

$$f(p_j \,|\, x_j = 0) = (n+1)\,(1-p_j)^n\,, \tag{53}$$

having a sharp peak at $p_j = 0$, with sharpness growing with n. Expected number and standard deviation are $\mathrm{E}[p_j \,|\, x_j] = \sigma[p_j \,|\, x_j] = 1/n$ [moreover, in this case the cumulative has the easy expression $F(p_j \,|\, x_j = 0) = 1 - (1 - p_j)^{n+1}$, from which we can calculate a 95% probability upper limit for p_j to be $1 - 0.05^{1/(n+1)}$ that is approximately $3/n$ for large n].

Solving a formal paradox with some good sense

The above conclusions seem reasonable when applied to an individual empty bin. But let us see, with the help of Tab. 2, what happens where several elements of $\boldsymbol{x}$ are null: one hundred events have been generated in a cause-cell and the second row of Tab. 2, whose elements are indicated by $x(E_j)$, shows the counts in each effect-bin. The third row shows the prior $\boldsymbol{\alpha}$ and the forth one the updated $\boldsymbol{\alpha}$. We also calculate in the following two rows expected values and standard deviation of the p_j values.

All elements of the smearing matrix that connect that cause-cell the effect-cells with zero counts have $\mathrm{E}[p_j] = \sigma[p_j] \approx 0.9\%$. This looks reasonable if we consider each cell individually. But, observing the values of $\mathrm{E}[p_j]$ and $\sigma[p_j]$ in the fifth and sixth row of Tab. 2, any physicist would be uneasy, for a couple of reasons. First, we are have strong prejudices towards regularity of physical laws, including whatever causes the smearing. If the smearing is peaked around bin E_3 and drops in both sides, we do not expect the smearing probability $p_{12} = P(E_{12} \,|\, C_i\,, I)$ to be equal to $p_7 = P(E_7 \,|\, C_i\,, I)$. Second, if we sum up the probabilities of all bins with no counts, i.e. $\sum_{j=7}^{12} \mathrm{E}[p_j]$, the total is not really negligible. The situation becomes worse if we add more bins in the right side.

Another way to understand the problem is starting from the probability density function of $p_{\mathcal{M}}$, the parameter of a binomial distribution that describes the occurrence of an event in the set $\mathcal{M}$ of m selected bins (not necessarily adjacent, although in Tab. 2 they are). Just extending a well known property of the Dirichlet distribution [see Appendix A.2, and in particular Eq. (40)], we get that $p_{\mathcal{M}}$ is described by a Beta pdf, with parameter $\alpha_{\mathcal{M}} = \sum_{k \subset \mathcal{M}} \alpha_k$ and $\alpha - \alpha_{\mathcal{M}}$, where $\alpha = \sum_j \alpha_j$:

$$f[p_{\mathcal{M}} \,|\, \mathrm{Beta}(\alpha_{\mathcal{M}},\, \alpha - \alpha_{\mathcal{M}})] \;=\; \frac{p_{\mathcal{M}}^{\alpha_{\mathcal{M}}-1}(1-p_{\mathcal{M}})^{\alpha-\alpha_{\mathcal{M}}-1}}{\beta(\alpha_{\mathcal{M}},\, \alpha - \alpha_{\mathcal{M}})}\,. \tag{54}$$

E_j	1	2	3	4	5	6	7	8	9	10	11	12
$x(E_j)$	10	20	40	15	10	5	0	0	0	0	0	0
	Standard uniform priors											
α_{pr_j}	1	1	1	1	1	1	1	1	1	1	1	1
α_{post_j}	11	21	41	16	11	6	1	1	1	1	1	1
$\mathrm{E}[p_j]$	9.8	18.8	36.6	14.3	9.8	5.4	0.9	0.9	0.9	0.9	0.9	0.9
$\sigma[p_j]$	2.8	3.7	4.5	3.3	2.8	2.1	0.9	0.9	0.9	0.9	0.9	0.9
	$\mathrm{E}[p_{1-6}] = 94.6\%$						$\mathrm{E}[p_{7-12}] = 5.4\%$					
	$\sigma[p_{1-6}] = 2.1\%$						$\sigma[p_{7-12}] = 2.1\%$					
	$\alpha_j = 2/\nu$											
α_{pr_j}	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6
α_{post_j}	10.2	20.2	40.2	15.2	10.2	5.2	0.2	0.2	0.2	0.2	0.2	0.2
$\mathrm{E}[p_j]$	10.0	19.8	39.4	14.9	10.0	5.1	0.02	0.02	0.02	0.02	0.02	0.02
$\sigma[p_j]$	3.0	3.9	4.8	3.5	3.0	2.2	0.4	0.4	0.4	0.4	0.4	0.4
	$\mathrm{E}[p_{1-6}] = 99.0\%$						$\mathrm{E}[p_{7-12}] = 1.0\%$					
	$\sigma[p_{1-6}] = 1.0\%$						$\sigma[p_{7-12}] = 1.0\%$					
	Regularized α's in bins with zero counts											
α_{pr_j}	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6
α_{post_j}	10.2	20.2	40.2	15.2	10.2	5.2	0.2	0.2	0.2	0.2	0.2	0.2
f	1	1	1	1	1	1	1	1/2	1/4	1/8	1/16	1/32
$\alpha^{(r)}_{post_j}$	10.2	20.2	40.2	15.2	10.2	5.2	0.2	0.1	0.04	0.02	0.01	0.005
$\mathrm{E}[p_j]$	10.0	19.9	39.6	15.0	10.0	5.1	0.2	0.1	0.04	0.02	0.01	0.005
$\sigma[p_j]$	3.0	3.9	4.8	3.5	3.0	2.2	0.4	0.3	0.2	0.1	0.1	0.1
	$\mathrm{E}[p_{1-6}] = 99.7\%$						$\mathrm{E}[p_{7-12}] = 0.3\%$					
	$\sigma[p_{1-6}] = 0.6\%$						$\sigma[p_{7-12}] = 0.6\%$					
	Regularized and rescaled α's $[\,* : \sum_{k\subset\mathcal{M}} \alpha_k = m/\nu\,]$											
α_{pr_j}	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6	1/6
α_{post_j}	10.2	20.2	40.2	15.2	10.2	5.2	0.2	0.2	0.2	0.2	0.2	0.2
f	1	1	1	1	1	1	1	1/2	1/4	1/8	1/16	1/32
$\alpha^{(r\&r)}_{post_j}$ *	10.2	20.2	40.2	15.2	10.2	5.2	0.5	0.3	0.13	0.06	0.03	0.015
$\mathrm{E}[p_j]$	10.0	19.8	39.4	14.9	10.0	5.1	0.5	0.2	0.1	0.06	0.03	0.02
$\sigma[p_j]$	3.0	3.9	4.8	3.5	3.0	2.2	0.7	0.5	0.3	0.2	0.2	0.1
	$\mathrm{E}[p_{1-6}] = 99.0\%$						$\mathrm{E}[p_{7-12}] = 1.0\%$					
	$\sigma[p_{1-6}] = 1.0\%$						$\sigma[p_{7-12}] = 1.0\%$					

Table 2: *MC events, Dirichlet parameters $\boldsymbol{\alpha}$ and expectations of Dirichlet variables p_j (value in %) with several treatments of the zeros (see text).*

If we take $\mathcal{M}$ to be the set of the m bins with zero counts ($m = 6$ in the case of Tab. 2) and for each of them we choose the initial value of α to be equal 1, we get:

$$f[p_{\mathcal{M}} \,|\, \mathrm{Beta}(m,\, n-m)] \;\; = \;\; \frac{p_{\mathcal{M}}^{m-1}(1-p_{\mathcal{M}})^{n-m-1}}{\beta(m,\, n-m)}\,. \tag{55}$$

Contrary to Eq. (53), this pdf vanishes for $p_{\mathcal{M}} \to 0$, thus stating *we are a priori sure that $p_{\mathcal{M}}$ cannot be zero*. This is a bit too much! We understand, then, that the required condition in order to have $f(p_{\mathcal{M}} = 0) = 0$ is to apply the following constraint to the α's: $\sum_{k\subset\mathcal{M}} \alpha_k \leq 1$, that becomes $\alpha_{k\subset\mathcal{M}} \leq 1/m$ in the case we consider all zero count bins on the same foot. Since the final α of the bins with many counts is influenced very little by the initial α, if this does not exceed $\mathcal{O}(1)$, then we could simply require $\alpha_j \approx 1/\nu$. But, just to recover the uniform prior when $\nu = 2$, which can be written as a Dirichlet with $\alpha_1 = 1$ and $\alpha_2 = 1$ (that is a Beta with $r = 1$ and $s = 1$), our starting point will be to set $\alpha = 2/\nu$ for all bins. Then this value will be modified due to the constraint $\sum_{k\subset\mathcal{M}} \alpha_k = m/\nu$ for the bins with zero counts. The result of this strategy is reported in Table 2 in the rows just below "$\alpha_j = 2/\nu$".

Regularization of the smearing matrix prior

There is still the problem that we do not believe all zeros are the same. We can solve it shaping in some way the α's in a suitable way. Obviously, in order to do this, we need to know the bin ordering in the parameter space, that was not taken into account till now. We show here a simple way of how this regularization can be performed in a 1-dimensional spectrum unfolding. We set all initial α's to $2/\nu$ in all bins that have non-zero counts and in the adjacent ones. Then, as we move far away from the bins with non-zero counts, we decrease the value of alpha, for example with some geometric law (i.e. exponential decrease). In Tab. 2 we have done this exercise taking the geometric factor equal 1/2 (see the rows indicated by 'f'). That is perhaps a too conservative choice that still takes in too much considerations the tails (indeed, in the program the geometric factor is given as an option in the program, whose default value is equal to $1/e$, i.e. a 'natural' exponential decrease that gives roughly one order of magnitude suppression every two bins). The α's shaped in this way are finally rescaled with the condition $\sum_{k\subset\mathcal{M}} \alpha_k = m/\nu$.

To conclude, these are the steps made in the default practical method to handle zeros in 1-dimensional unfolding:

- All bins:
  - set $\alpha_j = 2/\nu + x_j$ in all bins (α_{pr_j} in the table).
- Only bins with zero counts:
  - reshape the α's with an exponentially law;[15]
  - rescale the α's by the condition $\sum_{k\subset\mathcal{M}} \alpha_k = m/\nu$ [$\alpha^{(r\&r)}_{post_j}$ in the table].

The last rows of Tab. 2 show the results obtained applying this rule, which seems a reasonable compromise between the several pieces of information and beliefs as they have been stated here. Needless to say, you are free to adjust the treatments of the zero according to the best knowledge of your physical case. The program offers you the possibility to return your preferred α's using the function `my.alphas()`, or just to decide that zeros are simply zeros, and *basta*.

A warning about the Monte Carlo checks of Bayesian methods

A last comment about the outcomes of the simulations performed to prove the correctness of the above procedure is in order. In fact you might want to check the unfolding performing a large number of simulated experiments. Most likely you will arrive to the conclusion that the strategy to treat the zeros discussed above will provide in average a biased result. This is not a surprise, since you are setting some smearing probabilities to positive values, whereas they are most likely exactly zero in the simulation. This effect is discussed in section 10.6 of Ref. [6] and you should pay attention to it and make the simulation in the correct way, before stating that "Bayesian methods yield biased estimations".

B.2 – Zero counts in real data effect-bins

Having explained in length *a* strategy to treat the bins with zero counts in the evaluation of the smearing matrix, one could thing of doing something similar with the empty bins of the experimental spectrum. However it is

[15] In the case of a sequence of zero counts bins surrounded on both sides by non-zero bins, the two exponential starting at each side are summed up, to mimic the fact that there might be an incoherent sum of the two tails;

clear that in this second case the idea of handling in a sophisticate way the zeros might sound purely academic, because an appropriate re-binning is in most cases the simplest pragmatic solution.

Anyway, since a possible strategy follows from what it has been discussed in Appendix B.1, let us sketch it here briefly with the benefit of the inventory.

As we have seen in section 3.2, the default initial c and r parameters of the gamma pdf [see Eq. (25)] are 1 and 0, respectively, corresponding to an *improper* flat prior (extending to infinite, with infinite expected value). The Bayesian updating described by Eq. (25) makes the posterior (mathematically) nice (i.e. proper) and (physically) meaningful. In particular, empty bin j yields

$$f[\mu_j \,|\, x(E_j) = 0] \quad = \quad e^{-\mu_j} \,, \tag{56}$$

having null mode and unity expected value, that is in order.

The problem occurs when there are several empty bins, because if we treat them all in this way, we get that the pdf of the sum of the respective μ's (let us call it Σ) vanishes when its value goes to zero. Again, a possible solution is to constrain to zero the mode of Σ. This can be done playing with the c and r parameters of the empty bins, as done, for example, by the *custom* function (real cases require intelligent inputs by the user!) `my.gamma.par()` of the distribution code [5]. [Note that, contrary to `my.alphas()`, which tries to treat with a certain continuity adjacent zero bins, `my.gamma.par()` is more primitive, because, as said above, it is not considered important for practical applications.]