

Statistical Models with Uncertain Error Parameters

(G. Cowan, arXiv:1809.05778)



LAL Orsay
19 December 2018



Glen Cowan
Physics Department
Royal Holloway, University of London
g.cowan@rhul.ac.uk
www.pp.rhul.ac.uk/~cowan

Outline

Using measurements with “known” systematic errors:

Least Squares (BLUE)

Allowing for uncertainties in the systematic errors

Estimates of sys errors \sim Gamma

Single-measurement model

Asymptotics, Bartlett correction

Curve fitting, averages

Confidence intervals

Goodness-of-fit

Sensitivity to outliers

Discussion and conclusions

Details in: G. Cowan, *Statistical Models with Uncertain Error Parameters*, arXiv:1809.05778 [physics.data-an]

Introduction

Suppose measurements \mathbf{y} have probability (density) $P(\mathbf{y}|\boldsymbol{\mu},\boldsymbol{\theta})$,

$\boldsymbol{\mu}$ = parameters of interest

$\boldsymbol{\theta}$ = nuisance parameters

To provide info on nuisance parameters, often treat their best estimates \mathbf{u} as indep. Gaussian distributed r.v.s., giving likelihood

$$\begin{aligned} L(\boldsymbol{\mu}, \boldsymbol{\theta}) &= P(\mathbf{y}, \mathbf{u}|\boldsymbol{\mu}, \boldsymbol{\theta}) = P(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\theta})P(\mathbf{u}|\boldsymbol{\theta}) \\ &= P(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\theta}) \prod_{i=1}^N \frac{1}{\sqrt{2\pi}\sigma_{u_i}} e^{-(u_i - \theta_i)^2/2\sigma_{u_i}^2} \end{aligned}$$

or log-likelihood (up to additive const.)

$$\ln L(\boldsymbol{\mu}, \boldsymbol{\theta}) = \ln P(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\theta}) - \frac{1}{2} \sum_{i=1}^N \frac{(u_i - \theta_i)^2}{\sigma_{u_i}^2}$$

Systematic errors and their uncertainty

Often the θ_i could represent a systematic bias and its best estimate u_i in the real measurement is zero.

The $\sigma_{u,i}$ are the corresponding “systematic errors”.

Sometimes $\sigma_{u,i}$ is well known, e.g., it is itself a statistical error known from sample size of a control measurement.

Other times the u_i are from an indirect measurement, Gaussian model approximate and/or the $\sigma_{u,i}$ are not exactly known.

Or sometimes $\sigma_{u,i}$ is at best a guess that represents an uncertainty in the underlying model (“theoretical error”).

In any case we can allow that the $\sigma_{u,i}$ are not known in general with perfect accuracy.

Gamma model for variance estimates

Suppose we want to treat the systematic errors as uncertain, so let the $\sigma_{u,i}$ be adjustable **nuisance parameters**.

Suppose we have estimates s_i for $\sigma_{u,i}$ or equivalently $v_i = s_i^2$, is an estimate of $\sigma_{u,i}^2$.

Model the v_i as independent and gamma distributed:

$$f(v; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} v^{\alpha-1} e^{-\beta v}$$
$$E[v] = \frac{\alpha}{\beta}$$
$$V[v] = \frac{\alpha}{\beta^2}$$



Set α and β so that they give desired relative uncertainty r in σ_u .

Similar to method 2 in W.J. Browne and D. Draper, Bayesian Analysis, Volume 1, Number 3 (2006), 473-514.

Distributions of v and $s = \sqrt{v}$

For α, β of gamma distribution, $\alpha_i = \frac{1}{4r_i^2}$, $\beta_i = \frac{1}{4r_i^2 \sigma_{u_i}^2}$

$$r_i \equiv \frac{1}{2} \frac{\sigma_{v_i}}{E[v_i]} = \frac{1}{2} \frac{\sigma_{v_i}}{\sigma_{u_i}^2} \approx \frac{\sigma_{s_i}}{E[s_i]} \quad \leftarrow \text{relative "error on error"}$$



Motivation for gamma model

If one were to have n independent observations u_1, \dots, u_n , with all $u \sim \text{Gauss}(\theta, \sigma_u^2)$, and we use the sample variance

$$v = \frac{1}{n-1} \sum_{i=1}^n (u_i - \bar{u})^2$$

to estimate σ_u^2 , then $(n-1)v/\sigma_u^2$ follows a chi-square distribution for $n-1$ degrees of freedom, which is a special case of the gamma distribution ($\alpha = n/2, \beta = 1/2$). (In general one doesn't have a sample of u_i values, but if this were to be how v was estimated, the gamma model would follow.)

Furthermore choice of the gamma distribution for v allows one to profile over the nuisance parameters σ_u^2 in closed form and leads to a simple profile likelihood.

Likelihood for gamma error model

$$L(\boldsymbol{\mu}, \boldsymbol{\theta}, \boldsymbol{\sigma}_u^2) = P(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\theta}) \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma_{u_i}^2}} e^{-(u_i - \theta_i)^2 / 2\sigma_{u_i}^2}$$

$$\times \frac{\beta_i^{\alpha_i}}{\Gamma(\alpha_i)} v_i^{\alpha_i - 1} e^{-\beta_i v_i} .$$

Treated like data: y_1, \dots, y_L (the primary measurements)
 u_1, \dots, u_N (estimates of nuisance par.)
 v_1, \dots, v_N (estimates of variances
of estimates of NP)

Parameters: μ_1, \dots, μ_M (parameters of interest)
 $\theta_1, \dots, \theta_N$ (nuisance parameters)
 $\sigma_{u,1}, \dots, \sigma_{u,N}$ (sys. errors = std. dev. of
of NP estimates)

Profiling over systematic errors

We can profile over the $\sigma_{u,i}$ in closed form

$$\widehat{\sigma}_{u_i}^2 = \operatorname{argmax}_{\sigma_{u_i}^2} L(\boldsymbol{\mu}, \boldsymbol{\theta}, \boldsymbol{\sigma}_{\mathbf{u}}^2) = \frac{v_i + 2r_i^2(u_i - \theta_i)^2}{1 + 2r_i^2}$$

which gives the profile log-likelihood (up to additive const.)

$$\begin{aligned} \ln L'(\boldsymbol{\mu}, \boldsymbol{\theta}) &= \ln L(\boldsymbol{\mu}, \boldsymbol{\theta}, \widehat{\boldsymbol{\sigma}}_{\mathbf{u}}^2) \\ &= \ln P(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\theta}) - \frac{1}{2} \sum_{i=1}^N \left(1 + \frac{1}{2r_i^2}\right) \ln \left[1 + 2r_i^2 \frac{(u_i - \theta_i)^2}{v_i}\right] \end{aligned}$$

In limit of small r_i , $v_i \rightarrow \sigma_{u,i}^2$ and the log terms revert back to the quadratic form seen with known $\sigma_{u,i}$.

Equivalent likelihood from Student's t

We can arrive at same likelihood by defining $z_i \equiv \frac{u_i - \theta_i}{\sqrt{v_i}}$

Since $u_i \sim \text{Gauss}$ and $v_i \sim \text{Gamma}$, $z_i \sim \text{Student's } t$

$$f(z_i | \nu_i) = \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)}{\sqrt{\nu_i\pi}\Gamma(\nu_i/2)} \left(1 + \frac{z_i^2}{\nu_i}\right)^{-\frac{\nu_i+1}{2}} \quad \text{with} \quad \nu_i = \frac{1}{2r_i^2}$$

Resulting likelihood same as profile $L'(\boldsymbol{\mu}, \boldsymbol{\theta})$ from gamma model

$$L(\boldsymbol{\mu}, \boldsymbol{\theta}) = P(\mathbf{y} | \boldsymbol{\mu}, \boldsymbol{\theta}) \prod_{i=1}^N \frac{\Gamma\left(\frac{\nu_i+1}{2}\right)}{\sqrt{\nu_i\pi}\Gamma(\nu_i/2)} \left(1 + \frac{z_i^2}{\nu_i}\right)^{-\frac{\nu_i+1}{2}}$$

Single-measurement model

As a simplest example consider

$$y \sim \text{Gauss}(\mu, \sigma^2),$$

$$v \sim \text{Gamma}(\alpha, \beta), \quad \alpha = \frac{1}{4r^2}, \quad \beta = \frac{1}{4r^2\sigma^2}$$

$$L(\mu, \sigma^2) = f(y, v|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(y-\mu)^2/2\sigma^2} \frac{\beta^\alpha}{\Gamma(\alpha)} v^{\alpha-1} e^{-\beta v}$$



Test values of μ with $t_\mu = -2 \ln \lambda(\mu)$ with $\lambda(\mu) = \frac{L(\mu, \widehat{\sigma^2}(\mu))}{L(\hat{\mu}, \widehat{\sigma^2})}$

$$t_\mu = \left(1 + \frac{1}{2r^2}\right) \ln \left[1 + 2r^2 \frac{(y - \mu)^2}{v}\right]$$

Distribution of t_μ



From Wilks' theorem, in the asymptotic limit we should find $t_\mu \sim \text{chi-squared}(1)$.

Here “asymptotic limit” means all estimators $\sim \text{Gauss}$, which means $r \rightarrow 0$. For increasing r , clear deviations visible:



Distribution of t_μ (2)

For larger r , breakdown of asymptotics gets worse:





Values of $r \sim$ several tenths are relevant so we cannot in general rely on asymptotics to get confidence intervals, p -values, etc.

Bartlett corrections

One can modify t_μ defining
$$t'_\mu = \frac{n_d}{E[t_\mu]} t_\mu$$



such that the new statistic's distribution is better approximated by chi-squared for n_d degrees of freedom (Bartlett, 1937).

For this example $E[t_\mu] \approx 1 + 3r^2 + 2r^4$ works well:





Bartlett corrections (2)

Good agreement for $r \sim$ several tenths out to $\sqrt{t_\mu'} \sim$ several, i.e., good for significances of several sigma:



68.3% CL confidence interval for μ




Curve fitting, averages

Suppose independent
 $y_i \sim \text{Gauss}$, $i = 1, \dots, N$, with

$$E[y_i] = \varphi(x_i; \boldsymbol{\mu}) + \theta_i,$$

$$V[y_i] = \sigma_{y_i}^2.$$



$\boldsymbol{\mu}$ are the parameters of interest in the fit function $\varphi(x; \boldsymbol{\mu})$,

$\boldsymbol{\theta}$ are bias parameters constrained by control measurements

$u_i \sim \text{Gauss}(\theta_i, \sigma_{u,i})$, so that if $\sigma_{u,i}$ are known we have

$$-2 \ln L(\boldsymbol{\mu}, \boldsymbol{\theta}) = \sum_{i=1}^N \left[\frac{(y_i - \varphi(x_i; \boldsymbol{\mu}) - \theta_i)^2}{\sigma_{y_i}^2} + \frac{(u_i - \theta_i)^2}{\sigma_{u_i}^2} \right]$$

Profiling over θ_i with known $\sigma_{u,i}$

Profiling over the bias parameters θ_i for known $\sigma_{u,i}$ gives usual least-squares (BLUE)

$$-2 \ln L'(\boldsymbol{\mu}) = \sum_{i=1}^N \frac{(y_i - \varphi(x_i; \boldsymbol{\mu}) - u_i)^2}{\sigma_{y_i}^2 + \sigma_{u_i}^2} \equiv \chi^2(\boldsymbol{\mu})$$

Widely used technique for curve fitting in Particle Physics.

Generally in real measurement, $u_i = 0$.

Generalized to case of correlated y_i and u_i by summing statistical and systematic covariance matrices.

Curve fitting with uncertain $\sigma_{u,i}$

Suppose now $\sigma_{u,i}^2$ are adjustable parameters with gamma distributed estimates v_i .

Retaining the θ_i but profiling over $\sigma_{u,i}^2$ gives

$$-2 \ln L'(\boldsymbol{\mu}, \boldsymbol{\theta}) = \sum_{i=1}^N \left[\frac{(y_i - \varphi(x_i; \boldsymbol{\mu}) - \theta_i)^2}{\sigma_{y_i}^2} + \left(1 + \frac{1}{2r_i^2}\right) \ln \left(1 + 2r_i^2 \frac{(u_i - \theta_i)^2}{v_i}\right) \right]$$

Profiled values of θ_i from solution to cubic equations

$$\begin{aligned} \theta_i^3 + [-2u_i - y_i + \varphi_i] \theta_i^2 + \left[\frac{v_i + (1 + 2r_i^2)\sigma_{y_i}^2}{2r_i^2} + 2u_i(y_i - \varphi_i) + u_i^2 \right] \theta_i \\ + \left[(\varphi_i - y_i) \left(\frac{v_i}{2r_i^2} + u_i^2 \right) - \frac{(1 + 2r_i^2)\sigma_{y_i}^2 u_i}{2r_i^2} \right] = 0, \quad i = 1, \dots, N, \end{aligned}$$

Goodness of fit

Can quantify goodness of fit with statistic





$$q = -2 \ln \frac{L'(\hat{\mu}, \hat{\theta})}{L'(\hat{\varphi}, \hat{\theta})}$$
$$= \min_{\mu, \theta} \sum_{i=1}^N \left[\frac{(y_i - \varphi(x_i; \mu) - \theta_i)^2}{\sigma_{y_i}^2} + \left(1 + \frac{1}{2r_i^2} \right) \ln \left(1 + 2r_i^2 \frac{(u_i - \theta_i)^2}{v_i} \right) \right]$$

where $L'(\varphi, \theta)$ has an adjustable φ_i for each y_i (the saturated model).





Asymptotically should have $q \sim \text{chi-squared}(N-M)$.

For increasing r_i , may need Bartlett correction or MC.

Distributions of q




Distributions of Bartlett-corrected q'



Example: average of two measurements

MINOS interval (= approx. confidence interval) based on

$$\ln L'(\mu) = \ln L'(\hat{\mu}) - Q_\alpha/2 \quad \text{with} \quad Q_\alpha = F_{\chi^2}^{-1}(1 - \alpha; n)$$




Increased discrepancy between values to be averaged gives larger interval.

Interval length saturates at ~level of absolute discrepancy between input values.

relative error on sys. error

Same with interval from $p_\mu = \alpha$ with nuisance parameters profiled at μ



Coverage of intervals

Consider previous average of two numbers but now generate for $i = 1, 2$ data values

$$y_i \sim \text{Gauss}(\mu, \sigma_{y,i}),$$


$$u_i \sim \text{Gauss}(0, \sigma_{u,i}),$$

$$v_i \sim \text{Gamma}(\sigma_{u,i}, r_i)$$

$$\sigma_{y,i} = \sigma_{u,i} = 1$$


and look at the probability that the interval covers the true value of μ .

Coverage stays reasonable to $r \sim 0.5$, even not bad for Profile Construction out to $r \sim 1$.




Sensitivity of average to outliers

Suppose we average 5 values, $y = 8, 9, 10, 11, 12$, all with stat. and sys. errors of 1.0, and suppose negligible error on error (here take $r = 0.01$ for all).



Sensitivity of average to outliers (2)


Now suppose the measurement at 10 was actually at 20:



Estimate pulled up to 12.0, size of confidence interval \sim unchanged (would be exactly unchanged with $r \rightarrow 0$).

Average with all $r = 0.2$


If we assign to each measurement $r = 0.2$,



Estimate still at 10.00, size of interval moves $0.63 \rightarrow 0.65$

Average with all $r = 0.2$ with outlier

Same now with the outlier (middle measurement $10 \rightarrow 20$)



Estimate $\rightarrow 10.75$ (outlier pulls much less).

Half-size of interval $\rightarrow 0.78$ (inflated because of bad g.o.f.).

Naive approach to errors on errors

Naively one might think that the error on the error in the previous example could be taken into account conservatively by inflating the systematic errors, i.e.,

$$\sigma_{u_i} \rightarrow \sigma_{u_i} (1 + r_i)$$

But this gives

$$\hat{\mu} = 10.00 \pm 0.70 \quad \text{without outlier (middle meas. 10)}$$

$$\hat{\mu} = 12.00 \pm 0.70 \quad \text{with outlier (middle meas. 20)}$$

So the sensitivity to the outlier is not reduced and the size of the confidence interval is still independent of goodness of fit.

Correlated uncertainties

The phrase “correlated uncertainties” usually means that a single nuisance parameter affects the distribution (e.g., the mean) of more than one measurement.

For example, consider measurements y , parameters of interest μ , nuisance parameters θ with

$$E[y_i] = \varphi_i(\mu, \theta) \approx \varphi_i(\mu) + \sum_{j=1}^N R_{ij} \theta_j$$

That is, the θ_i are defined here as contributing to a bias and the (known) factors R_{ij} determine how much θ_i affects y_i .

As before suppose one has independent control measurements $u_i \sim \text{Gauss}(\theta_i, \sigma_{ui})$.

Correlated uncertainties (2)

The total bias of y_i can be defined as
$$b_i = \sum_{j=1}^N R_{ij} \theta_j$$

which can be estimated with
$$\hat{b}_i = \sum_{j=1}^N R_{ij} u_j$$

These estimators are correlated having covariance

$$U_{ij} = \text{cov}[\hat{b}_i, \hat{b}_j] = \sum_{k=1}^N R_{ik} R_{jk} V[u_k]$$

In this sense the present method treats “correlated uncertainties”, i.e., the control measurements u_i are independent, but nuisance parameters affect multiple measurements, and thus bias estimates are correlated.

Discussion / Conclusions

Gamma model for variance estimates gives confidence intervals that increase in size when the data are internally inconsistent, and gives decreased sensitivity to outliers (known property of Student's t based regression).

Equivalence with Student's t model, $\nu = 1/2r^2$ degrees of freedom.

Simple profile likelihood – quadratic terms replaced by logarithmic:

$$\frac{(u_i - \theta_i)^2}{\sigma_{u_i}^2} \quad \longrightarrow \quad \left(1 + \frac{1}{2r_i^2}\right) \ln \left[1 + 2r_i^2 \frac{(u_i - \theta_i)^2}{v_i}\right]$$

Discussion / Conclusions (2)

Asymptotics can break for increased error-on-error, may need Bartlett correction or MC.

Model should be valuable when systematic errors are not well known but enough “expert opinion” is available to establish meaningful errors on the errors.

Could also use e.g. as “stress test” – crank up the r_i values until significance of result degrades and ask if you really trust the assigned systematic errors at that level.

Here assumed that meaningful r_i values can be assigned. Alternatively one could try to fit a global r to all systematic errors, analogous to PDG scale factor method or meta-analysis à la DerSimonian and Laird. (→ future work).

Extra slides

Gamma model for estimates of variance

Suppose the estimated variance v was obtained as the sample variance from n observations of a Gaussian distributed bias estimate u .

In this case one can show v is gamma distributed with

$$\alpha = \frac{n-1}{2} \quad \beta = \frac{n-1}{2\sigma_u^2}$$

We can relate α and β to the relative uncertainty r in the systematic uncertainty as reflected by the standard deviation of the sampling distribution of s , σ_s

$$r = \frac{\sigma_s}{E[s]} = \frac{1}{2} \frac{\sigma_v}{E[v]}$$

Exact relation between r parameter and relative error on error

r parameter defined as: $r \equiv \frac{1}{2} \frac{\sigma_v}{E[v]} \approx \frac{\sigma_s}{E[s]}$

$v \sim \text{Gamma}(\alpha, \beta)$ so $s = \sqrt{v}$ follows a Nakagami distribution

$$g(s|\alpha, \beta) = \left| \frac{dv}{ds} \right| f(v(s)|\alpha, \beta) = \frac{2\beta^\alpha}{\Gamma(\alpha)} s^{2\alpha-1} e^{-\beta s^2}$$

$$E[s] = \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)\sqrt{\beta}}$$

$$V[s] = \frac{\alpha}{\beta} - \frac{1}{\beta} \left(\frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)} \right)^2$$


Exact relation between r parameter and relative error on error (2)

The exact relation between the error and the error r_s and the parameter r is therefore

$$r_s \equiv \frac{\sqrt{V[s]}}{E[s]}$$
$$= \sqrt{\alpha \left(\frac{\Gamma(\alpha)}{\Gamma(\alpha + \frac{1}{2})} \right)^2 - 1}$$

$$\alpha = 1/4r^2$$

→ $r_s \approx r$ good for $r \approx 1$.



PDG scale factor

Suppose we do not want to take the quoted errors as known constants. Scale the variances by a factor ϕ ,

$$\sigma_i^2 \rightarrow \phi \sigma_i^2$$

The likelihood function becomes

$$L(\mu, \phi) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\phi\sigma_i^2}} \exp \left[-\frac{1}{2} \frac{(y_i - \mu)^2}{\phi\sigma_i^2} \right]$$

The estimator for μ is the same as before; for ϕ ML gives

$$\hat{\phi}_{\text{ML}} = \frac{\chi^2(\hat{\mu})}{N} \quad \text{which has a bias;} \quad \hat{\phi} = \frac{\chi^2(\hat{\mu})}{N-1} \quad \text{is unbiased.}$$

The variance of $\hat{\mu}$ is inflated by ϕ :

$$V[\hat{\mu}] = \frac{\phi}{\sum_{i=1}^N \frac{1}{\sigma_i^2}}$$

Bayesian approach

G. Cowan, *Bayesian Statistical Methods for Parton Analyses*, in *Proceedings of the 14th International Workshop on Deep Inelastic Scattering (DIS2006)*, M. Kuze, K. Nagano, and K. Tokushuku (eds.), Tsukuba, 2006.

Given measurements: $y_i \pm \sigma_i^{\text{stat}} \pm \sigma_i^{\text{sys}}, \quad i = 1, \dots, n,$

and (usually) covariances: $V_{ij}^{\text{stat}}, V_{ij}^{\text{sys}}.$

Predicted value: $\mu(x_i; \theta),$ expectation value $E[y_i] = \mu(x_i; \theta) + b_i$

control variable parameters bias

Frequentist approach: $V_{ij} = V_{ij}^{\text{stat}} + V_{ij}^{\text{sys}}$

Minimize $\chi^2(\theta) = (\vec{y} - \vec{\mu}(\theta))^T V^{-1} (\vec{y} - \vec{\mu}(\theta))$


Its Bayesian equivalent

Take $L(\vec{y}|\vec{\theta}, \vec{b}) \sim \exp \left[-\frac{1}{2}(\vec{y} - \vec{\mu}(\theta) - \vec{b})^T V_{\text{stat}}^{-1} (\vec{y} - \vec{\mu}(\theta) - \vec{b}) \right]$

$$\pi_b(\vec{b}) \sim \exp \left[-\frac{1}{2} \vec{b}^T V_{\text{sys}}^{-1} \vec{b} \right]$$

$$\pi_\theta(\theta) \sim \text{const.}$$

Joint probability
for all parameters



and use Bayes' theorem: $p(\theta, \vec{b}|\vec{y}) \propto L(\vec{y}|\theta, \vec{b})\pi_\theta(\theta)\pi_b(\vec{b})$

To get desired probability for θ , integrate (marginalize) over b :

$$p(\theta|\vec{y}) = \int p(\theta, \vec{b}|\vec{y}) d\vec{b}$$

→ Posterior is Gaussian with mode same as least squares estimator, σ_θ same as from $\chi^2 = \chi^2_{\text{min}} + 1$. (Back where we started!)

Bayesian approach with non-Gaussian prior $\pi_b(b)$

Suppose now the experiment is characterized by

$$y_i, \quad \sigma_i^{\text{stat}}, \quad \sigma_i^{\text{sys}}, \quad s_i, \quad i = 1, \dots, n,$$

where s_i is an (unreported) factor by which the systematic error is over/under-estimated.

Assume correct error for a Gaussian $\pi_b(b)$ would be $s_i \sigma_i^{\text{sys}}$, so

$$\pi_b(b_i) = \int \frac{1}{\sqrt{2\pi s_i \sigma_i^{\text{sys}}}} \exp \left[-\frac{1}{2} \frac{b_i^2}{(s_i \sigma_i^{\text{sys}})^2} \right] \pi_s(s_i) ds_i$$

Width of $\sigma_s(s_i)$ reflects
'error on the error'.

Error-on-error function $\pi_s(s)$


A simple unimodal probability density for $0 < s < 1$ with adjustable mean and variance is the Gamma distribution:

$$\pi_s(s) = \frac{a(as)^{b-1}e^{-as}}{\Gamma(b)}$$

$$\text{mean} = b/a$$

$$\text{variance} = b/a^2$$


Want e.g. expectation value of 1 and adjustable standard Deviation σ_s , i.e., $a = b = 1/\sigma_s^2$



In fact if we took $\pi_s(s) \sim \text{inverse Gamma}$, we could find $\pi_b(b)$ in closed form (cf. D'Agostini, Dose, von Linden). But Gamma seems more natural & numerical treatment not too painful.

Prior for bias $\pi_b(b)$ now has longer tails

$$\pi_b(b_i) = \int \frac{1}{\sqrt{2\pi s_i \sigma_i^{\text{sys}}}} \exp \left[-\frac{1}{2} \frac{b_i^2}{(s_i \sigma_i^{\text{sys}})^2} \right] \pi_s(s_i) ds_i$$



Gaussian ($\sigma_s = 0$) $P(|b| > 4\sigma_{\text{sys}}) = 6.3 \times 10^{-5}$


$\sigma_s = 0.5$ $P(|b| > 4\sigma_{\text{sys}}) = 0.65\%$

A simple test


Suppose a fit effectively averages four measurements.

Take $\sigma_{\text{sys}} = \sigma_{\text{stat}} = 0.1$, uncorrelated.

Case #1: data appear compatible



Posterior $p(\mu|y)$:





Usually summarize posterior $p(\mu|y)$
with mode and standard deviation:

$$\begin{aligned}\sigma_s = 0.0 : & \quad \hat{\mu} = 1.000 \pm 0.071 \\ \sigma_s = 0.5 : & \quad \hat{\mu} = 1.000 \pm 0.072\end{aligned}$$

Simple test with inconsistent data

Case #2: there is an outlier

Posterior $p(\mu|y)$:



$$\sigma_s = 0.0 : \quad \hat{\mu} = 1.125 \pm 0.071$$

$$\sigma_s = 0.5 : \quad \hat{\mu} = 1.093 \pm 0.089$$

→ Bayesian fit less sensitive to outlier. See also

G. D'Agostini, *Sceptical combination of experimental results: General considerations and application to epsilon-prime/epsilon*, arXiv:hep-ex/9910036 (1999).

Goodness-of-fit vs. size of error

In LS fit, value of minimized χ^2 does not affect size of error on fitted parameter.

In Bayesian analysis with non-Gaussian prior for systematics, a high χ^2 corresponds to a larger error (and vice versa).

