THE ROCKEFELLER UNIVERSITY

Science for the benefit of humanity

Measurements and Searches in High Energy Physics

Luc Demortier

Institute of Mathematical Statistics

73rd Annual Meeting, August 9-13, 2010, Gothenburg, Sweden

Outline

1. High energy physics in a nutshell;
2. Types of statistical analysis done by high energy physicists;
3. What do we expect from a statistical method?
4. How do our actual methods perform?
5. A recent development: Bayesian reference analysis.

Experimental High Energy Physics in a Nutshell (1/2)

- Goal of HEP: to identify the ultimate constituents of matter and understand their interactions.
- Method of HEP: to study the products of the particle collisions created at large accelerators such as the Tevatron at Fermilab (near Chicago, IL) and the LHC at CERN (near Geneva, Switzerland).

- Although accelerators typically produce millions of collisions per second, *interesting* collisions ("signal") are very rare. Thus, the process of selecting such collisions is very stringent, and the number of collision events selected for further study obeys Poisson statistics to a very good approximation.
- Each collision event is characterized by thousands of measurements of the energy, momentum, and charge of the particles produced as "debris" in the collision. Thus we are dealing with marked Poisson processes, and the distributions of the marks help discriminate interesting from uninteresting events.

Experimental High Energy Physics in a Nutshell (2/2)

The so-called standard model of particle physics describes all observations so far made, and predicts the existence of a particle known as the Higgs boson, to account for the masses of all particles. The Higgs boson has not been observed yet but is within range of experiments at both the Tevatron and the LHC.

More disturbingly, the standard model is known to be theoretically incomplete. However, there are some candidates for its completion, the most popular of which is "Supersymmetry" or "SuSy". There is at present no shred of evidence for SuSy. However, this theory makes many predictions, essentially doubling the spectrum of currently known particles, so that many searches to confirm it are planned or already under way.

Types of Statistical Analysis in HEP (1/2)

1. Point and Interval Estimation
   Measuring masses, production rates, and lifetimes of various elementary particles can be helpful both in testing standard model predictions and in making new predictions.
2. Hypothesis Testing
   Alternative hypotheses can vary widely in specificity, from very vague (a deviation in the tail of a distribution, or a "bump" somewhere) to very precise (charge of top quark is $-4/3$ versus $+2/3$).
3. Goodness-of-Fit Testing
   Most often used at the preparatory stage of an analysis to check one's understanding of the experimental apparatus.
4. Data Mining
   Given a general idea of what new physics should look like, one can construct a search engine to systematically scan all the data. This approach presents a number of issues: (1) in the presence of systematic effects one needs to introduce tuning parameters to cancel any observed discrepancy, within constraints; (2) proper accounting for the "look-elsewhere effect" is not trivial; and (3) a global search tends to be less than optimal against specific alternative hypotheses.

Types of Statistical Analysis in HEP (2/2)

Typical difficulties associated with HEP analyses include:

- modeling prior information;
- handling nuisance parameters;
- incorporating physical boundaries;
- choosing significance thresholds;
- inference after testing;
- accounting for the sensitivity of the experimental apparatus;
- interpreting the measurement results.

Choosing Significance Thresholds

Hypothesis testing is typically done with p values, which are converted into "numbers of σ's" according to the formula:

$$N_\sigma = \Phi^{-1}(1 - p).$$

When testing a hypothesis H_0, we claim evidence if $N_\sigma \geq 3$ ($p \leq 1.35 \times 10^{-3}$), and discovery if $N_\sigma \geq 5$ ($p \leq 2.87 \times 10^{-7}$).

The reason for the rather large discovery threshold dates back to a 1960's back-of-the-envelope calculation of the expected number of fake discoveries the field would claim given the number of tests performed by high energy physicists every year.

The 5σ discovery threshold is universal in our field, regardless of sample size, prior evidence, or even separate accounting for multiple testing effects. We would like to find a more flexible criterion for discovery, but it is difficult to come up with a rational prescription that works in all cases and with which everybody agrees.

Inference after Testing

Sometimes we calculate a confidence interval on a parameter after performing a test on it. Suppose for example that we are interested in the production rate μ of a new particle:

- First we test $H_0 : \mu = 0$ versus $H_1 : \mu > 0$.
- If H_0 is accepted, the experiment found no evidence for the existence of the particle, and an upper limit on its production rate is computed.
- If H_0 is rejected, there is evidence of the particle's existence, and a two-sided interval on μ is computed.

Note that the same data are used twice. At least in a frequentist setting, this type of procedure can lead to bias.

The Sensitivity Problem

This problem occurs when measuring parameters that are bounded by physical constraints, for example particle masses, production rates, efficiencies, etc.

To fix ideas, suppose that we wish to compute an interval on a particle mass μ, and that our measurement X is Gaussian:

$$X \sim \mathcal{N}(\mu, \sigma),$$

where σ is the measurement resolution. The length of a frequentist central interval for μ decreases to zero as X becomes negative. In that case the interval length no longer provides an estimate of the measurement uncertainty.

HEP Desiderata

1. Inferences should be done according to general, well-defined principles, that are neither arbitrary nor subjective;
2. Inferences should not depend on prior information about the parameter of interest, except for well-understood physical constraints;
3. Inferences should incorporate *all* such constraints;
4. Point and interval estimates should be equivariant under one-to-one transformations of the parameter of interest;
5. Experimental uncertainty should be adequately reflected in the reported results;
6. Inferences should be associated with a precise statement of probability.

Choice of Statistical Paradigm

1. Frequentism:
   Is preferred by many physicists because its definition of probability seems straightforward, and there is no need to construct priors for parameters about which nothing is known a priori. However, physical boundaries and nuisance parameters are difficult to handle.
2. Bayes:
   Solves the nuisance parameter problem and the sensitivity problem, but requires the elicitation of priors for *all* parameters. In addition, physicists are not very comfortable with the idea of probability as degree of belief.
3. Likelihood:
   Provides an approximation to frequentism or Bayes, but its interpretation is somewhat murky.

In practice, physicists often use a combination of these paradigms, for example a frequentist handling of the parameter of interest combined with a Bayesian handling of nuisance parameters.

In some experiments analysts are encouraged to report results in more than one paradigm.

Performance of HEP Methods (1/3)

Evolution of some interval measurements over time:

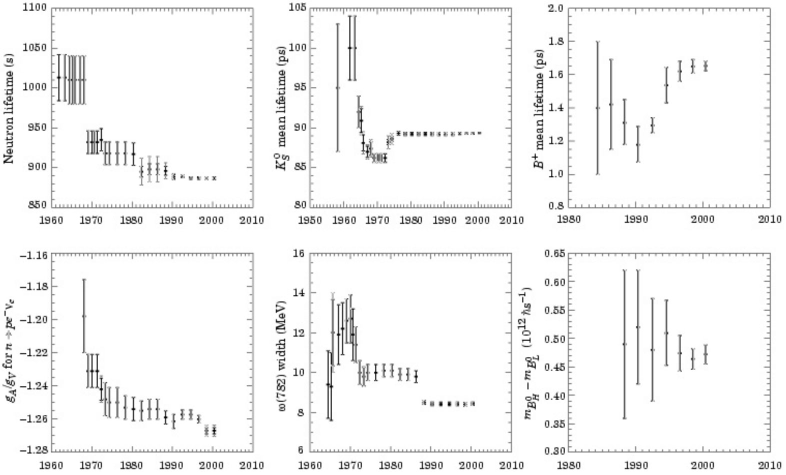

Performance of HEP Methods (2/3)

Example of a discovery claim that eventually fizzled out. . .

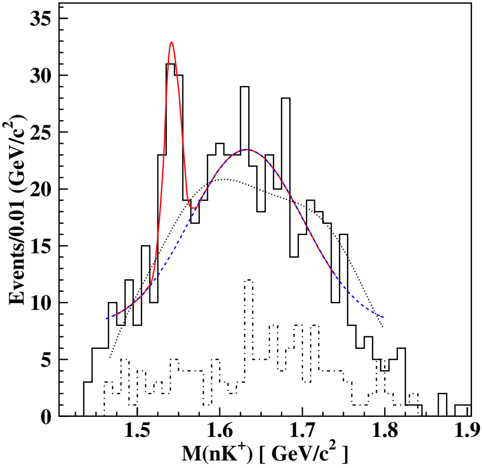

Performance of HEP Methods (3/3)

. . . but it took many tries to get this right:

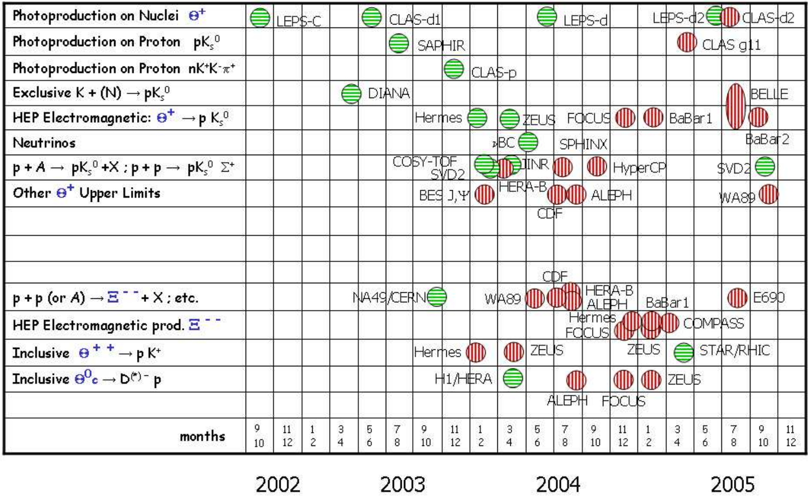

(from Reinhard A. Schumacher, "The Rise and Fall of Pentaquarks in Experiments," arXiv:nucl-ex/0512042v1, 27 Dec 2005.)

A Further Look at Objective Bayesian Methods

Physicist interest in Bayesian methods is growing, especially objective Bayesian methods. . .

An Example: Measurement of a Signal Cross Section

Suppose we are interested in measuring an upper limit on the production rate of some type of "signal" events. After selecting events with potentially interesting marks, we can write for their total number N:

$$N \sim \frac{(\mu + \epsilon\sigma)^n}{n!} \, e^{-\mu-\epsilon\sigma},$$

where:

- μ is the background contamination, i.e. the expected number of events that are not signal but are essentially indistinguishable from it;
- $\epsilon\sigma$ is the expected number of signal events, written as a product of the effective luminosity ϵ and a quantity of intrinsic physics interest, the signal cross section σ.

Usually there is prior information about μ and ϵ, from simulation studies, previous measurements, etc. As for σ, physicists prefer to act as if there is no prior information about it, because:

- There may not be universal agreement about the prior information;
- Even if there is, it is not always clear how to use it;
- In potential discovery situations, it is important for experiments to ignore tantalizing results obtained by competitors, in order to facilitate an objective assessment of the available evidence.

Measurement of a Signal Cross Section: Using a Flat Prior (1/3)

Suppose for the moment that the background contamination μ is known exactly. Then we only need to elicit a prior for the signal cross section σ and the effective luminosity ϵ. A typical choice would be a prior that is flat in σ and truncated Gaussian in ϵ:

$$\pi(\sigma, \epsilon) \;\propto\; e^{\frac{1}{2}\left(\frac{\epsilon-\epsilon_0}{\Delta\epsilon}\right)^2}$$

Unfortunately, with this choice of prior, posterior upper limits diverge to infinity. The reason for this is the flat improper prior for σ. Can we solve this problem by introducing a "cutoff" in the cross section prior?

Measurement of a Signal Cross Section: Using a Flat Prior (2/3)

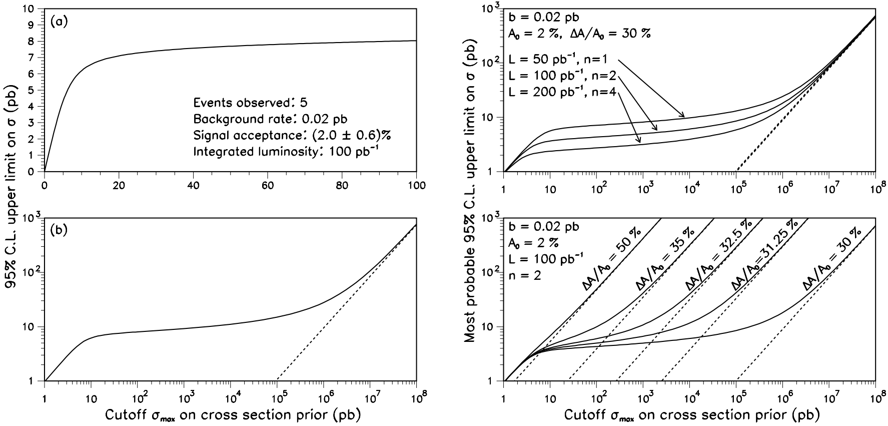

Bayesian upper limits at the 95% credibility level on a hypothetical cross section σ, as a function of the cutoff σ_{max} on the flat prior for σ.

Measurement of a Signal Cross Section: Using a Flat Prior (3/3)

Truncating the flat cross section prior yields finite upper limits; however:

- The resulting upper limits are very sensitive to the effective luminosity uncertainty;
- There is no obvious choice for the cutoff on the cross section prior that is general enough to be useful for everybody all of the time.

Using a gamma acceptance prior instead of a truncated Gaussian one yields finite upper limits, even with a flat improper cross section prior. This solution is not satisfactory however, because:

- It means that one has to give up on a fair robustness analysis.
  For a 10% effective luminosity uncertainty, the gamma and truncated Gaussian densities are not all that different. Therefore, the upper limit should not be sensitive to which one is chosen as prior, but it is.
- It leads to incoherence.
  With a flat improper prior for the cross section and a gamma prior $\pi(\epsilon)$ for the effective luminosity, the marginal *posterior* for the effective luminosity is proportional to $\pi(\epsilon)/\epsilon$, regardless of how much data is collected. In other words, information about the effective luminosity gets updated even if the experiment is not done.

Bayesian Reference Analysis

To avoid the problems introduced by the use of flat priors, some of us are considering a solution based on Bayesian reference analysis (see for example J. O. Berger, J. Bernardo, and D. Sun, "The formal definition of reference priors," Ann. Statist. **37**, 905 (2009)). There are many advantages to such an approach: inferences are invariant under parameter transformations and behave well under measurement replication, the method is very general and computationally tractable, it is "objective Bayesian" while avoiding the marginalization paradoxes, and it can be embedded in a subjective Bayesian framework.

In contrast with many applications of reference analysis, we will not require the full power of the method with regard to nuisance parameters. In high energy physics we usually have partial information about these, which simplifies some of the calculations. As we will see however, this also brings up some interesting new issues.

Reference Priors with Partial Information

Suppose that ϕ labels the nuisance parameter(s) and θ the parameter of interest. If we have prior information about ϕ, there are two ways to proceed (Sun & Berger 1998):

Method 1: We are given a marginal prior $\pi(\phi)$ for ϕ.
In this case we need the conditional reference prior $\pi_R(\theta \mid \phi)$ for θ given a fixed value of ϕ.

Method 2: We are given a conditional prior $\pi(\phi \mid \theta)$ for ϕ given θ.
In this case we can marginalize the probability model $p(x|\theta, \phi)$ with respect to ϕ in order to obtain $p(x|\theta) = \int p(x|\theta, \phi)\, \pi(\phi|\theta)\, d\phi$, and we can then compute the reference prior $\pi_R(\theta)$ for the marginalized model.

See D. Sun and J. O. Berger, "Reference priors with partial information," Biometrika **85**, 55 (1998).

Method 1 Applied to a Cross Section Measurement (1/5)

The single-count measurement is characterized by the likelihood:

$$p(n|\sigma, \epsilon, \mu) \;=\; \frac{(\epsilon\,\sigma + \mu)^n}{n!}\; e^{-\epsilon\,\sigma - \mu} \quad \text{with} \quad \sigma \geq 0 \quad \text{and} \quad \epsilon, \mu > 0.$$

In Method 1 we construct the conditional reference prior $\pi_R(\sigma \,|\, \epsilon, \mu)$. The first step consists in calculating Jeffreys' prior for σ while holding ϵ and μ fixed:

$$\pi_J(\sigma \,|\, \epsilon, \mu) \;\propto\; \left\{ \mathbb{E}\left[-\frac{\partial^2}{\partial \sigma^2} \ln p(n \,|\, \sigma, \epsilon, \mu) \right] \right\}^{\frac{1}{2}} \;\propto\; \frac{\epsilon}{\sqrt{\epsilon\,\sigma + \mu}}.$$

This prior is improper with respect to σ however, so that an additional step, known as the "compact support argument," is required:

Choose a nested sequence $\Lambda_1 \subset \Lambda_2 \subset \cdots$ of compact subsets of the parameter space Λ for (σ, ϵ, μ), such that $\cup_i \Lambda_i = \Lambda$ and the integral $K_i(\epsilon, \mu)$ of $\pi_J(\sigma \,|\, \epsilon, \mu)$ over $\Omega_i \equiv \{\sigma : (\sigma, \epsilon, \mu) \in \Lambda_i\}$ is finite. Then, on Ω_i:

$$\pi_{R,i}(\sigma \,|\, \epsilon, \mu) \;=\; \frac{\pi_J(\sigma \,|\, \epsilon, \mu)}{K_i(\epsilon, \mu)}\; I_{\Omega_i}(\sigma),$$

and on the whole parameter space:

$$\pi_R(\sigma \,|\, \epsilon, \mu) \;=\; \lim_{i \to \infty} \frac{\pi_{R,i}(\sigma \,|\, \epsilon, \mu)}{\pi_{R,i}(\sigma_0 \,|\, \epsilon_0, \mu_0)} \quad \text{with } (\sigma_0, \epsilon_0, \mu_0) \text{ a fixed point.}$$

Method 1 Applied to a Cross Section Measurement (2/5)

The theory of reference priors provides few guidelines for choosing the compact sets Λ_i. Try:

$$\Lambda_i \;=\; \Big\{(\sigma, \epsilon, \mu):\; \sigma \in [0, u_i],\; \epsilon \in [0, v_i],\; \mu \in [0, w_i]\Big\},$$

where $\{u_i\}$, $\{v_i\}$, and $\{w_i\}$ are increasing sequences of positive constants. Then:

$$\pi_{R,i}(\sigma \,|\, \epsilon, \mu) \;=\; \frac{1}{K_i(\epsilon, \mu)} \, \frac{\epsilon}{\sqrt{\epsilon\,\sigma + \mu}} \, I_{[0,u_i]}(\sigma),$$

where:

$$K_i(\epsilon, \mu) \;\equiv\; \int_{\Omega_i} \frac{\epsilon}{\sqrt{\epsilon\,\sigma + \mu}} \, d\sigma \;=\; 2\left[\sqrt{\epsilon\, u_i + \mu} - \sqrt{\mu}\right],$$

and therefore:

$$\pi_R(\sigma \,|\, \epsilon, \mu) \;\propto\; \sqrt{\frac{\epsilon}{\epsilon\,\sigma + \mu}}.$$

This is improper with respect to σ, but the ϵ dependence differs from that of the conditional Jeffreys' prior. There is a problem however. If

$$\pi(\epsilon, \mu) \;=\; \frac{e^{-\epsilon}}{\sqrt{\pi\,\epsilon/2}} \, \pi(\mu),$$

then the resulting posterior $p(n \,|\, \sigma, \epsilon, \mu)\, \pi_R(\sigma \,|\, \epsilon, \mu)\, \pi(\epsilon, \mu)$ is improper!

Method 1 Applied to a Cross Section Measurement (3/5)

The cause of this problem is the choice of compact sets. Note that the Jeffreys prior for this problem, $\pi_J(\sigma \,|\, \epsilon, \mu)\, d\sigma$, is invariant under scale transformations $\epsilon \to c\epsilon$, $\sigma \to \sigma/c$, where c is constant. Our initial choice of compact sets does not share this invariance, so we try instead:

$$\Lambda_i \;=\; \Big\{(\sigma, \epsilon, \mu):\; \sigma \in [0, u_i/\epsilon],\; \epsilon \in [1/v_i, v_i],\; \mu \in [0, w_i]\Big\},$$

where u_i, v_i, and w_i are as before. Repeating the same calculation as before, we now find:

$$\pi_{R1}(\sigma \,|\, \epsilon, \mu) \;\propto\; \frac{\epsilon}{\sqrt{\epsilon\sigma + \mu}},$$

which is identical to Jeffreys' prior for this case and yields well-behaved posteriors.

Method 1 Applied to a Cross Section Measurement (4/5)

To do further calculations with the Method 1 conditional reference prior we must specify a subjective prior for ϵ and μ. Here we take:

$$\pi(\epsilon, \mu) \;=\; \frac{a(a\epsilon)^{x-1/2}\, e^{-a\epsilon}}{\Gamma(x+1/2)} \quad \frac{b(b\mu)^{y-1/2}\, e^{-b\mu}}{\Gamma(y+1/2)},$$

A typical situation is that information about μ and ϵ is obtained from Monte Carlo simulations, subsidiary measurements, and theoretical beliefs. . . This information is summarized by point estimates $(\bar{\epsilon}, \bar{\mu})$ and relative uncertainties $(\delta\epsilon, \delta\mu)$, which are then identified with the corresponding means and coefficients of variation of the two component distributions of $\pi(\epsilon, \mu)$:

$$\bar{\epsilon} = \frac{x+\frac{1}{2}}{a}, \quad \delta\epsilon = \frac{1}{\sqrt{x+\frac{1}{2}}}, \quad \text{or} \quad x = \frac{1}{\delta\epsilon^2} - \frac{1}{2}, \quad a = \frac{1}{\bar{\epsilon}\,\delta\epsilon^2};$$

$$\bar{\mu} = \frac{y+\frac{1}{2}}{b}, \quad \delta\mu = \frac{1}{\sqrt{y+\frac{1}{2}}}, \quad \text{or} \quad y = \frac{1}{\delta\mu^2} - \frac{1}{2}, \quad b = \frac{1}{\bar{\mu}\,\delta\mu^2}.$$

Another possible interpretation of $\pi(\epsilon, \mu)$ is that it is the joint posterior of two Poisson measurements x and y of the effective luminosity and background, respectively, and where Jeffreys' prior was used for both ϵ and μ.

Method 1 Applied to a Cross Section Measurement (5/5)

For calculating posterior summaries in terms of intervals and upper limits it is convenient to express the marginal Method-1 posterior as a tail probability:

$$\int_{\sigma}^{\infty} \pi_{R1}(\tau \mid n)\, d\tau \;=\; \int_{\frac{\sigma}{a+\sigma}}^{1} \frac{u^{n+y}\,(1-u)^{x-\frac{1}{2}}}{B\left(n+y+1, x+\frac{1}{2}\right)}\; \frac{B_{\frac{b}{b+1}(1+\frac{u-1}{u}\,\frac{\sigma}{a})}\left(y+\frac{1}{2}, n+\frac{1}{2}\right)}{B_{\frac{b}{b+1}}\left(y+\frac{1}{2}, n+\frac{1}{2}\right)}\, du$$

where

$$B_z(u, v) \;\equiv\; \int_0^z t^{u-1}\,(1-t)^{v-1}\, dt$$

is the incomplete beta function, and $B(u, v) \equiv B_1(u, v) \;=\; \Gamma(u)\Gamma(v)/\Gamma(u+v)$.

Method 2 Applied to a Cross Section Measurement (1/2)

The first step in Method 2 is the calculation of the marginal data pdf. We use the same (ϵ, μ) prior as before, i.e. we assume that $\pi(\epsilon, \mu \,|\, \sigma) = \pi(\epsilon, \mu)$. Thus:

$$\begin{aligned}
p(n \,|\, \sigma) &= \iint p(n \,|\, \sigma, \epsilon, \mu)\; \pi(\epsilon, \mu \,|\, \sigma)\; d\epsilon\, d\mu, \\
&= \iint \frac{(\epsilon\sigma + \mu)^n}{n!}\; e^{-\epsilon\sigma - \mu}\; \frac{a(a\epsilon)^{x-1/2}}{\Gamma(x + 1/2)}\; e^{-a\epsilon}\; \frac{b(b\mu)^{y-1/2}}{\Gamma(y + 1/2)}\; e^{-b\mu}\; d\epsilon\, d\mu, \\
&= \left[\frac{a}{a+\sigma}\right]^{x+\frac{1}{2}} \left[\frac{b}{b+1}\right]^{y+\frac{1}{2}} \sum_{k=0}^{n} u_{nk} \left[\frac{\sigma}{a+\sigma}\right]^{k},
\end{aligned}$$

where

$$u_{nk} = \binom{x - \frac{1}{2} + k}{k} \binom{y - \frac{1}{2} + n - k}{n - k} \left[\frac{1}{b+1}\right]^{n-k},$$

and we used generalized binomial coefficients:

$$\binom{v}{w} \equiv \frac{\Gamma(v+1)}{\Gamma(w+1)\,\Gamma(v-w+1)}.$$

Method-2 Applied to a Cross Section Measurement (2/2)

We then compute Jeffreys' prior from $p(n \mid \sigma)$:

$$\pi_{R2}(\sigma) \;\propto\; \sqrt{\mathbb{E}\left\{\left[\frac{d}{d\sigma}\ln p(n \mid \sigma)\right]^2\right\}} \;\propto\; \sqrt{\sum_{n=0}^{\infty} \frac{\left[(x+1/2)\,S_n^0 \;-\; (a/\sigma)\,S_n^1\right]^2}{(a+\sigma)^{x+5/2}\;S_n^0}},$$

with

$$S_n^m \;\equiv\; \sum_{k=0}^{n} k^m\, u_{nk}\;\left[\frac{\sigma}{a+\sigma}\right]^k \quad \text{for} \quad m = 0, 1.$$

The posterior is simply:

$$\pi_{R2}(\sigma \mid n) \;\propto\; p(n \mid \sigma)\;\pi_{R2}(\sigma).$$

Its normalization must be obtained numerically.

Note that Method 2 does not require a compact support argument.

Marginal Cross Section Reference Priors and Posteriors

Example priors and posteriors for Methods 1 and 2:

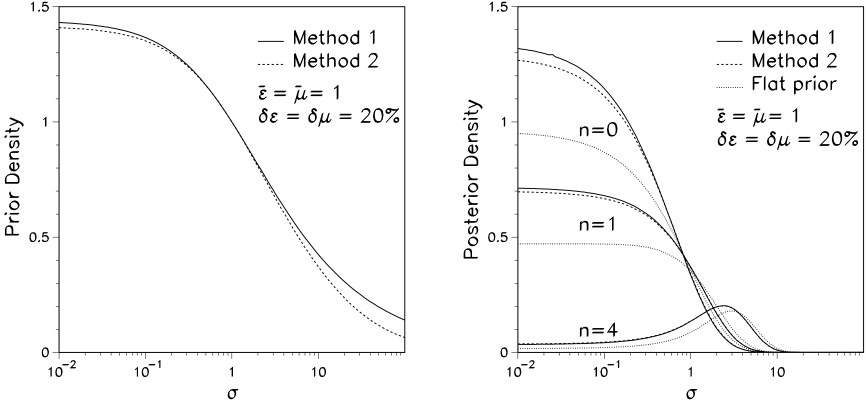

Generalizations

So far we have only considered single-count measurements. This model can be generalized in a number of ways, all of which are important in high energy physics:

1. The multiple-count model, based on the likelihood:

$$p(\vec{n} \mid \sigma, \vec{\epsilon}, \vec{\mu}) \;=\; \prod_{i=1}^{M} \frac{(\mu_i + \epsilon_i\,\sigma)^{n_i}}{n_i!}\, e^{-\mu_i - \epsilon_i\,\sigma}.$$

   The Method-1 reference prior for this model is:

$$\pi_J(\sigma \mid \vec{\epsilon}, \vec{\mu}) \;=\; \sqrt{\sum_{i=1}^{M} \frac{\epsilon_i^2}{\mu_i + \epsilon_i\,\sigma}}.$$

2. Unbinned likelihoods

See L. D., S. Jain, and H. B. Prosper, "Reference Priors for High Energy Physics," Phys. Rev. D **82**, 034002 (2010). There is software available to handle these more complicated problems. However we have not found a way to implement the compact support argument in all generality.

Upper Limits (1/4)

A common way to summarize posterior distributions is by computing intervals. Here we briefly look at upper limits:

- How much difference is there between Method-1 and Method-2 reference posterior upper limits?
- How do upper limits vary as a function of the mean of the background prior?
- How do upper limits vary with the uncertainties on background and effective luminosity?
- How do upper limits behave under experiment replication?

Upper Limits (2/4)

Difference between Method-1 and Method-2 upper limits at the 68% (left) and 95% (right) credibility levels:

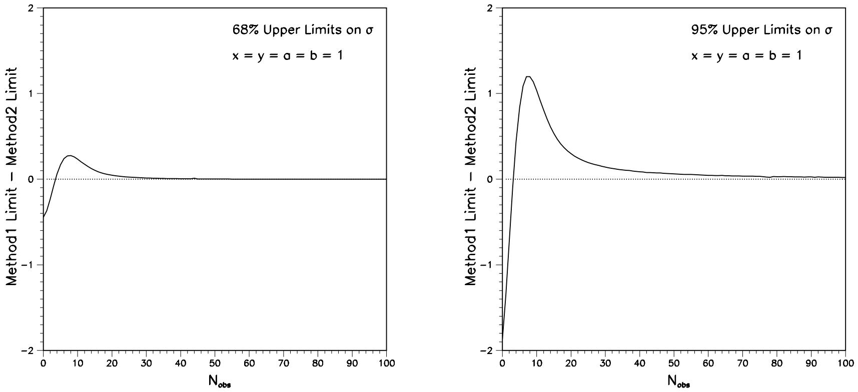

Upper Limits (3/4)

Dependence of various 90% C.L. upper limit constructions on expected background, when there are no uncertainties:

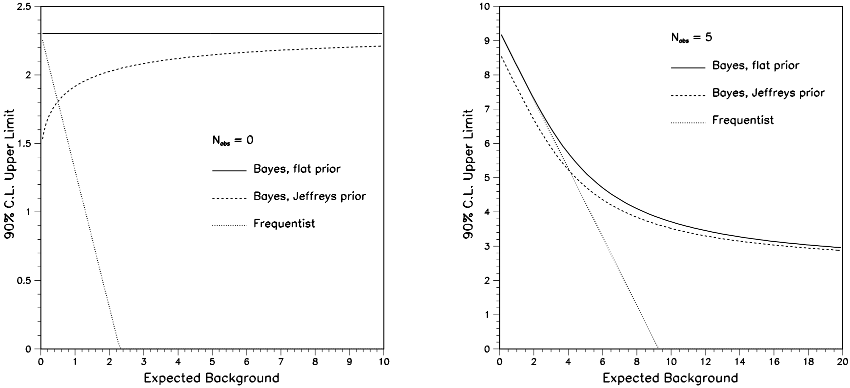

Upper Limits (4/4)

Variation of 95% C.L. Method-1 upper limit with mean background, for $\delta\epsilon = \delta\mu = 20\%$ (left) and $\delta\epsilon = \delta\mu = 50\%$ (right):

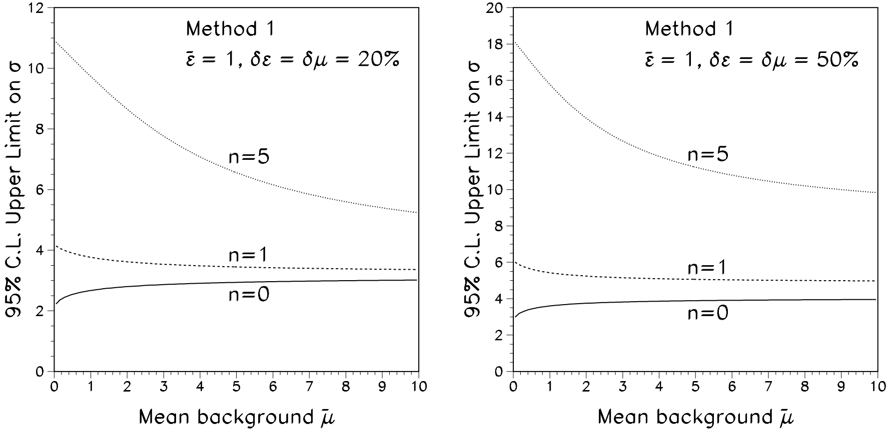

Behavior under Measurement Replication (1/2)

Frequentist coverage of Bayesian procedures:

- Bayesian intervals based on a subjective prior satisfy an "average coverage" theorem, according to which the average of the frequentist coverage over the prior equals the credibility of the interval.
- When the prior is objective, and especially when it is improper, there is no natural metric for calculating an average coverage. In this case, a useful criterion is the *pointwise* coverage of the intervals.
- When the prior is a mixture of objective and subjective components, a natural approach is to average the coverage over the subjective prior components and check pointwise coverage with respect to the remaining parameters.

In our case, adopting the latter approach means that we will calculate the coverage with respect to the marginalized data pdf:

$$p(n \,|\, \sigma) \;=\; \iint p(n \,|\, \sigma, \epsilon, \mu)\; \pi(\epsilon, \mu \,|\, \sigma)\; d\epsilon\, d\mu.$$

Behavior under Measurement Replication (2/2)

For the Method-1 reference prior (Method-2 plots are very similar):

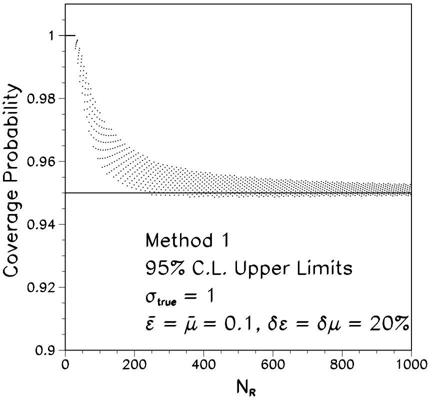

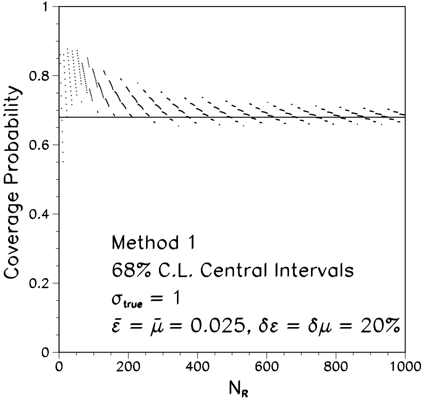

Conclusions

Due to the ever increasing size of high energy physics datasets, and the difficulty and sophistication of their analysis, physicists are paying more attention to correct statistical procedures. We have reviewed some of the issues that are currently being debated. It is unclear whether these have unambiguous solutions, so that the best approach may be to do the statistical analysis in more than one way and compare the results.

Interactions between physicists and statisticians take place regularly through the so called "PhyStat" conferences and workshops:

- Jan.2000: `http://doc.cern.ch/cernrep/2000/2000-005/2000-005.html`;
- Mar.2000: `http://conferences.fnal.gov/cl2k/`;
- Mar.2002: `http://www.ippp.dur.ac.uk/Workshops/02/statistics/`;
- Sep.2003: `http://www.slac.stanford.edu/econf/C030908/`;
- Sep.2005: `http://www.physics.ox.ac.uk/phystat05/proceedings/default.htm`;
- Jun.2007: `http://phystat-lhc.web.cern.ch/phystat-lhc/`.