BAYES versus FREQUENTISM

The Return of an Old Controversy

- The ideologies, with examples
- Upper limits
- Systematics

Louis Lyons, Oxford University

and CERN

PHYSTAT 2003

SLAC, STANFORD, CALIFORNIA

8^{TH} –11^{TH} SEPT 2003

Conference on:

STATISTICAL PROBLEMS IN:

PARTICLE PHYSICS

ASTROPHYSICS

COSMOLOGY

http://www-conf.slac.stanford.edu/phystat2003

It is possible to spend a lifetime analysing data without realising that there are two very different approaches to statistics:

Bayesianism and Frequentism.

How can textbooks not even mention Bayes/ Frequentism?

For simplest case $(m \pm \sigma) \leftarrow Gaussian$

with no constraint on $m(true)$ then

$$m - k\sigma < m(true) < m + k\sigma$$

at some probability, for both Bayes and Frequentist

(but different interpretations)

See Bob Cousins "Why isn't every physicist a Bayesian?" Amer Jrnl Phys 63(1995)398

We need to make a statement about

Parameters, Given Data

The basic difference between the two:

Bayesian : Probability (parameter, given data)
(an anathema to a Frequentist!)

Frequentist : Probability (data, given parameter)
(a likelihood function)

Bayesian versus Classical

Bayesian

$P(A \text{ and } B) = P(A;B) \times P(B) = P(B;A) \times P(A)$

e.g. A = event contains t quark

B = event contains W boson

or A = you are in CERN

B = you are at Workshop

Completely uncontroversial, provided….

$P(A;B) = P(B;A) \times P(A) / P(B)$

Bayesian

$$P(A;B) = \frac{P(B;A) \text{ x } P(A)}{P(B)}$$

Bayes Theorem

$$P(hyothesis;data) \propto P(data;hypothesis) \text{ x } P(hypothesis)$$

posterior ↑ likelihood ↑ prior ↑

Problems: P(hyp..) true or false

"Degree of belief"

Prior What functional form?

Coverage

Goodness of fit

P(hypothesis…..) True or False

"Degree of Belief"

credible interval

Prior: What functional form?

Uninformative prior:

flat? In which variable? e.g. $m, m^2, \ln m, ...?$

Unimportant if "data overshadows prior"

Important for limits

Subjective or Objective prior?

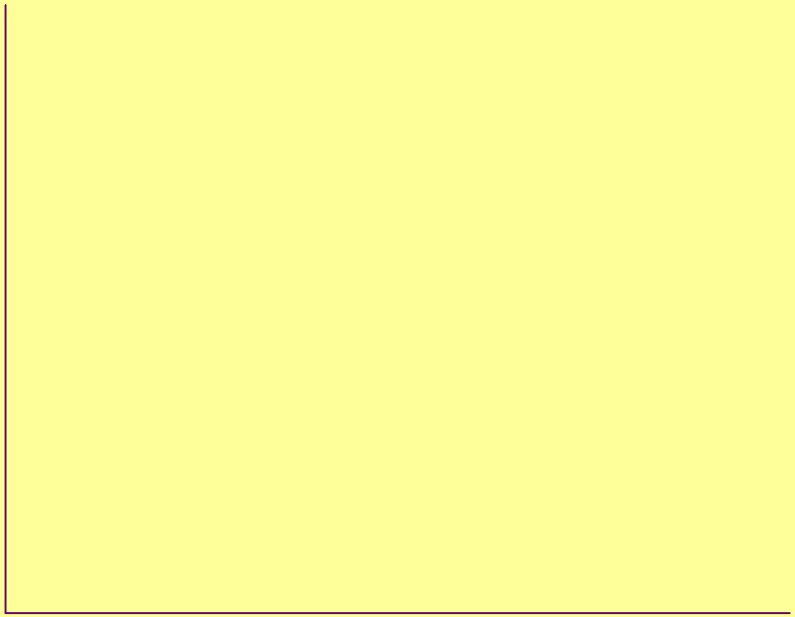

P (Data;Theory) $\neq$ P (Theory;Data)

HIGGS SEARCH at CERN

Is data consistent with Standard Model?

or with Standard Model + Higgs?

End of Sept 2000 Data not very consistent with S.M.

Prob (Data ; S.M.) < 1% valid frequentist statement

Turned by the press into: Prob (S.M. ; Data) < 1%

and therefore Prob (Higgs ; Data) > 99%

i.e. “It is almost certain that the Higgs has been seen”

P (Data;Theory) $\neq$ P (Theory;Data)

Theory = male or female

Data = pregnant or not pregnant

P (pregnant ; female) ~ 3%

but

P (female ; pregnant) >>>3%

Example 1 : Is coin fair ?

Toss coin: 5 consecutive tails

What is P(unbiased; data) ? i.e. p = ½

Depends on Prior(p)

If village priest prior ~ $\delta(1/2)$

If stranger in pub prior ~ 1 for 0<p<1

(also needs cost function)

Example 2 : Particle Identification

Try to separate π and protons

probability (p tag;real p) = 0.95

probability (π tag; real p) = 0.05

probability (p tag ; real (π) = 0.10

probability (π tag ; real π) = 0.90

Particle gives proton tag. What is it?

Depends on prior = fraction of protons

If proton beam, very likely

If general secondary particles, more even

If pure π beam, ~ 0

Hunter and Dog

1) Dog d has 50% probability of being 100 m. of Hunter h
2) Hunter h has 50% probability of being within 100m of Dog d

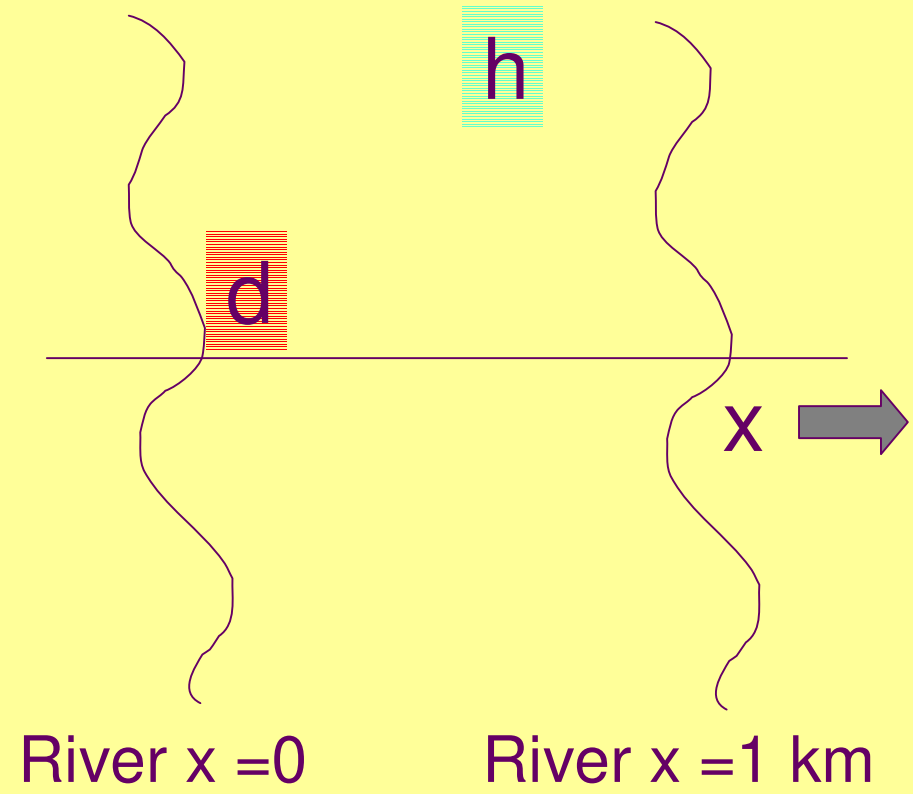

Given that: a) Dog d has 50% probability of being 100 m. of Hunter

Is it true that b) Hunter h has 50% probability of being within 100m of Dog d ?

Additional information

- Rivers at zero & 1 km. Hunter cannot cross them.

$$0 \le h \le 1\,\text{km}$$

- Dog can swim across river - Statement a) still true

If dog at –101 m, hunter cannot be within 100m of dog

Statement b) untrue

Classical Approach

Neyman "confidence interval" avoids pdf for μ

uses only P(x; μ)

Confidence interval $\mu_1 \rightarrow \mu_2$:

P($\mu_1 \rightarrow \mu_2$ contains μ) = α True for any μ

Varying intervals from ensemble of experiments

fixed

Gives range of μ for which observed value x_0 was "likely" (α)

Contrast Bayes : Degree of belief = α that μ_t is in $\mu_1 \rightarrow \mu_2$

COVERAGE

If true for all μ : "correct coverage"

P< α for some μ "undercoverage"
(this is serious !)

P> α for some μ "overcoverage"

Conservative

Loss of rejection power

$\mu_l \leq \mu \leq \mu_u$ at 90% confidence

Frequentist

μ_l and μ_u known, but random

μ unknown, but fixed

Probability statement about μ_l and μ_u

Bayesian

μ_l and μ_u known, and fixed

μ unknown, and random

Probability/credible statement about μ

Classical Intervals

- Problems

Hard to understand e.g. d'Agostini e-mail

Arbitrary choice of interval

Possibility of empty range

Over-coverage for integer observation

e.g. # of events

Nuisance parameters (systematic errors)

- Advantages

Widely applicable

Well defined coverage

Importance of Ordering Rule

Neyman construction in 1 parameter μ

2 measurements X_1 X_2

$$p(x;\mu) = G(x-\mu,1)$$

An aside: Determination of single parameter p via χ^2

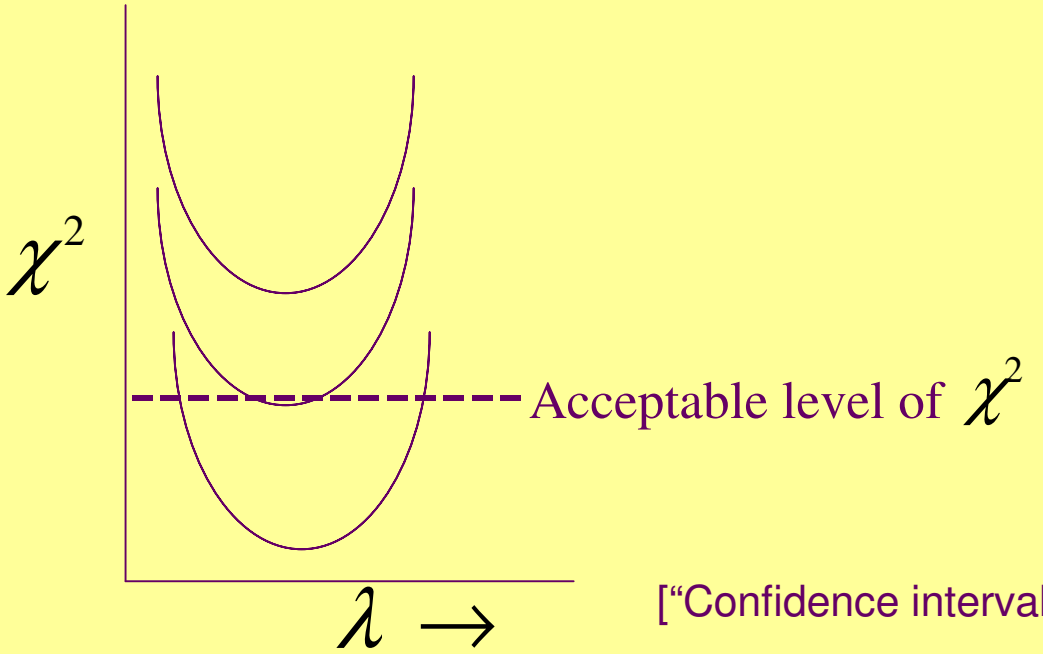

Range of parameters given by

1) Values of λ for which data is likely i.e. p(χ^2) is acceptable or

2) $\chi^2(\lambda) < \chi^2_{min}(\lambda) + 1$

2) is good

1) Range depends on χ^2_{min}

["Confidence interval coupled to goodness of fit"]

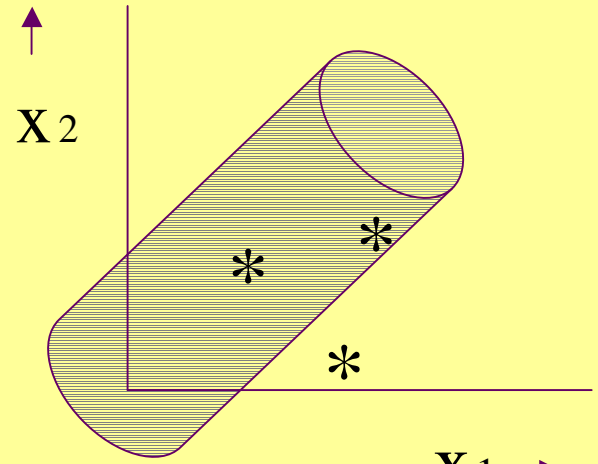

Neyman Construction

For given μ, acceptable (x_1 , x_2) satisfy

$$\chi^2 = (x_1 - \mu)^2 + (x_2 - \mu)^2 \leq Ccut$$

Defines cylinder in (μ, x_1, x_2) space

Experiment gives $(x_1, x_2) \rightarrow \mu$ interval

Range depends on $|x_1 - x_2|$

$$\mu = \frac{x_1 + x_2}{2} \pm \sqrt{2 - (x_1 - x_2)^2}\,/2$$

Range and goodness of fit are coupled

That was using Probability Ordering

Now change to Likelihood Ratio Ordering

For $x_1 \neq x_2$,no value of μ gives very good fit

For Neyman Construction at fixed μ , compare:

$$(x_1-\mu)^2+(x_2-\mu)^2 \quad \text{with} \quad (x_1-\mu_{best})^2+(x_2-\mu_{best})^2$$

where $\mu_{best} = (x_1 + x_2)/2$

giving $$2\left[\mu^2-\mu(x_1+x_2)+\frac{1}{4}(x_1+x_2)^2\right] = 2\left[\mu-\frac{1}{2}(x_1+x_2)\right]^2$$

Cutting on Likelihood Ratio Ordering gives:

$$\mu = \frac{x_1+x_2}{2} \pm \sqrt{\frac{C}{2}}$$

$$\mu = \frac{x_1 + x_2}{2} \pm \sqrt{\frac{C}{2}}$$

Therefore, range of μ is

Constant Width

Independent of $x_1 - x_2$

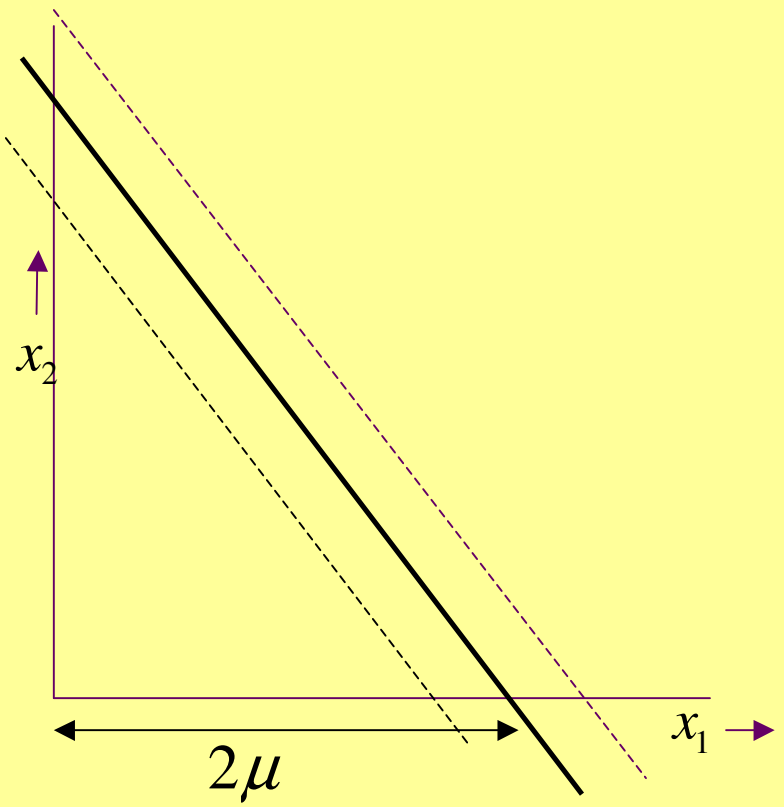

Confidence Range and Goodness of Fit are completely decoupled

Bayesian

Pros:

Easy to understand

Physical Interval

Cons:

Needs prior

Hard to combine

Coverage

Standard Frequentist

Pros:

Coverage

Cons:

Hard to understand

Small or Empty Intervals

Different Upper Limits

SYSTEMATICS

For example

$$N_{events} = \sigma LA + b$$

Observed

$N \pm \sqrt{N}$

for statistical errors

Physics parameter

we need to know these, probably from other measurements (and/or theory)

Uncertainties → error in σ

Some are arguably statistical errors

$$LA = LA_0 \pm \sigma_{LA}$$

$$b = b_0 \pm \sigma_b$$

Shift Central Value

Bayesian

Frequentist

Mixed

$$N_{events} = \sigma \, LA + b$$

Simplest Method

Evaluate σ_0 using LA_0 and b_0

Move nuisance parameters (one at a time) by their errors → $\delta\sigma_{LA}$ & $\delta\sigma_b$

If nuisance parameters are uncorrelated

Combine these contribution in quadrature

→ total systematic

Bayesian

Without systematics

$$p(\sigma; N) \propto p(N; \sigma)\Pi(\sigma)$$

↑ prior

With systematics

$$p(\sigma, LA, b; N) \propto p(N; \sigma, LA, b)\Pi(\sigma, LA, b)$$

↑

$$\sim \Pi_1(\sigma)\Pi_2(LA)\Pi_3(b)$$

Then integrate over LA and b

$$p(\sigma; N) = \iint p(\sigma, LA, b; N)\, dLA\, db$$

$$p(\sigma; N) = \iint p(\sigma, LA, b; N)\, dLA\ db$$

If $\Pi_1(\sigma)$ = constant and $\Pi_2(\mathrm{LA})$ = truncated Gaussian TROUBLE!

Upper limit on σ from $\int p(\sigma; N)\, \mathrm{d}\sigma$

Significance from likelihood ratio for $\sigma = 0$ and $\sigma_{\max}$

Frequentist

Full Method

Imagine just 2 parameters σ and LA

and 2 measurements N and M

↑ Physics ↑ Nuisance

Do Neyman construction in 4-D

Use observed N and M, to give

Confidence Region

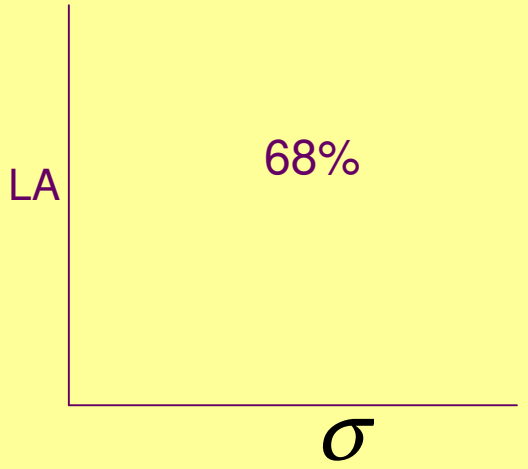

Then project onto σ axis

This results in OVERCOVERAGE

Aim to get better shaped region, by suitable choice of ordering rule

Example: Profile likelihood ordering

$$\frac{L(N_0 M_0; \sigma, LA_{best}(\sigma))}{L(N_0 M_0; \sigma_{best}, LA_{best}(\sigma))}$$

Full frequentist method hard to apply in several dimensions

Used in $\leq$3 parameters

For example: Neutrino oscillations (CHOOZ)

$$\sin^2 2\theta \, , \Delta m^2$$

Normalisation of data

Use approximate frequentist methods that reduce dimensions to just physics parameters

e.g. Profile pdf

i.e. $pdf_{profile}(N;\sigma) = pdf(N, M_0; \sigma, LA_{best})$

Contrast Bayes marginalisation

Distinguish "profile ordering"

Properties being studied by Giovanni Punzi

Talks at FNAL CONFIDENCE LIMITS WORKSHOP

(March 2000) by:

Gary Feldman

Wolfgang Rolke hep-ph/0005187 version 2

Acceptance uncertainty worse than Background uncertainty

$$\text{Limit of C.L. as } \sigma \rightarrow 0$$

$$\neq \text{C.L. for } \sigma = 0$$

Need to check Coverage

Method: Mixed Frequentist - Bayesian

Bayesian for nuisance parameters and

Frequentist to extract range

Philosophical/aesthetic problems?

Highland and Cousins

(Motivation was paradoxical behavior of Poisson limit when LA not known exactly)

Bayesian versus Frequentism

	Bayesian	Frequentist
Basis of method	Bayes Theorem --> Posterior probability distribution	Uses pdf for data, for fixed parameters
Meaning of probability	Degree of belief	Frequentist defintion
Problem of parameters?	Yes	Anathema
Needs prior?	Yes	No
Choice of interval?	Yes	Yes (except F+C)
Data considered	Only data you have	….+ more extreme
likelihood principle?	Yes	No

Bayesian versus Frequentism

	Bayesian	Frequentist
Ensemble of experiment	No	Yes (but often not explicit)
Final statement	Posterior probability distribution	Parameter values → Data is likely
Unphysical/ empty ranges	Excluded by prior	Can occur
Systematics	Integrate over prior	Extend dimensionality of frequentist construction
Coverage	Unimportant	Built-in
Decision making	Yes (uses cost function)	Not useful

Bayesianism versus Frequentism

"Bayesians address the question everyone is interested in, by using assumptions no-one believes"

"Frequentists use impeccable logic to deal with an issue of no interest to anyone"