Confidence Intervals in Particle (Astro-)physics

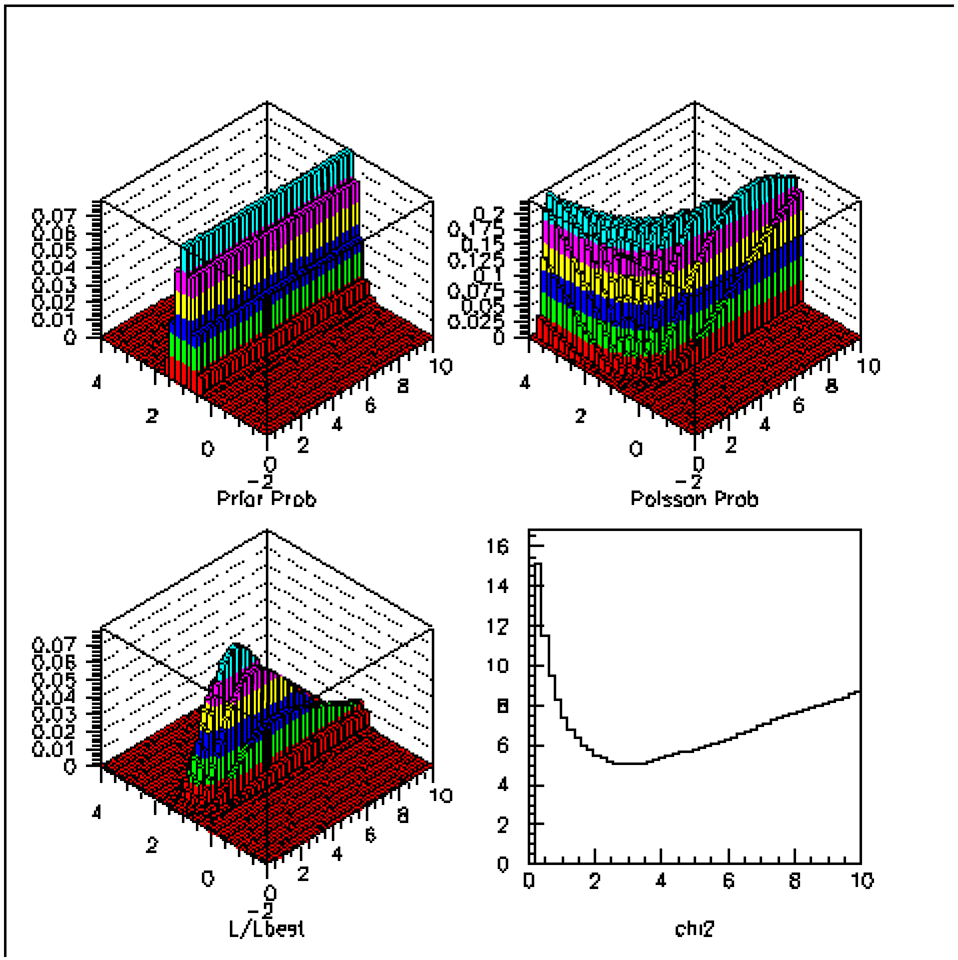

Jan Conrad (CERN)

Content

- Preliminaries
  - Definitions
  - Frequentist and Bayesian statistics
- Methods of calculation of Confidence Intervals
  - Bayesian
    - Example: $B^0_s \rightarrow \mu^+\mu^-$
  - Frequentist
    - Example: $B^0_s \rightarrow \mu^+\mu^-$
    - Example: KARMEN experiment
    - Example: CHOOZ experiment
    - Example: Higgs search at LHC
- Summary
- So, what would I do ?

Some Definitions

- **A physical measurement, n, is used to estimate the parameter, *s*, of a PDF (e.g. Poisson, or Poisson with "known" background b)**
- **The confidence interval [s1,s2] identifies a range in which we expect the true value of s**
- **The probability for the confidence interval to contain (cover) our hypothesis s is called the confidence level (usually symbolized by 1-α)**
- **Experimenter chooses a confidence level (usually 100(1-α) % = 90 %, 95 %, 99 % etc) and presents confidence intervals corresponding to this level**
- **Systematic and statistical uncertainties are uncertainties present in parameters which will affect the confidence interval, but which are not of prime interest (nuisance parameters), e.g. signal efficiency ε or background expectation, b)**

Note: the 68 % confidence interval on a physical quantity is what we usually call "error", e.g. in fit-parameters, number of entries in an histogram etc. pp.

Definitions cont'd

- We will in the remainder of this talk consider a Poisson with efficiency signal parameter s (e.g. cross-section), an efficiency ε (e.g luminosity x acceptance) and background b[1)]

$$P(n|\varepsilon_{true}s + b_{true})$$

- The nuisance parameters (say background) are usually determined in subsidiary measurements (side band) or by Monte Carlo simulation

The set of observations is thus $(n,\varepsilon_{meas},b_{meas})$ and the set of parameters $(s,\varepsilon_{true},b_{true})$

Bayesian vs. Frequentist

- Frequentist probability:
  - defined in terms of repeated experiment, statements are about **P(data | Hypothesis)**
  - In particular: Hypothesis is not a random variable

    → **P(Hypothesis)** is not defined
- Bayesian probability:
  - Degree of belief, statements are about **P(Hypothesis|data)**
- Connection: Bayes Theorem:
  - **P(Hypothesis | data) ≈ P(data|hypothesis) * P(Hypothesis)**

    (P(Hypothesis) ← Prior)

Key performance figure of frequentist confidence intervals: **coverage**

- A method is said to have **coverage** $(1-\alpha)$ if, in infinitely many repeated experiments the resulting confidence interval includes the true value with probability $(1-\alpha)$ irrespective of what the true value is

Bayesian dude: don't worry about coverage

Bayesian physicist: use Bayesian methods, but check coverage !

More on Coverage

- The Coverage requirement is necessary and sufficient for a construction method to be valid.

- Coverage can never be fulfilled exactly for Poisson process (discrete random variable)

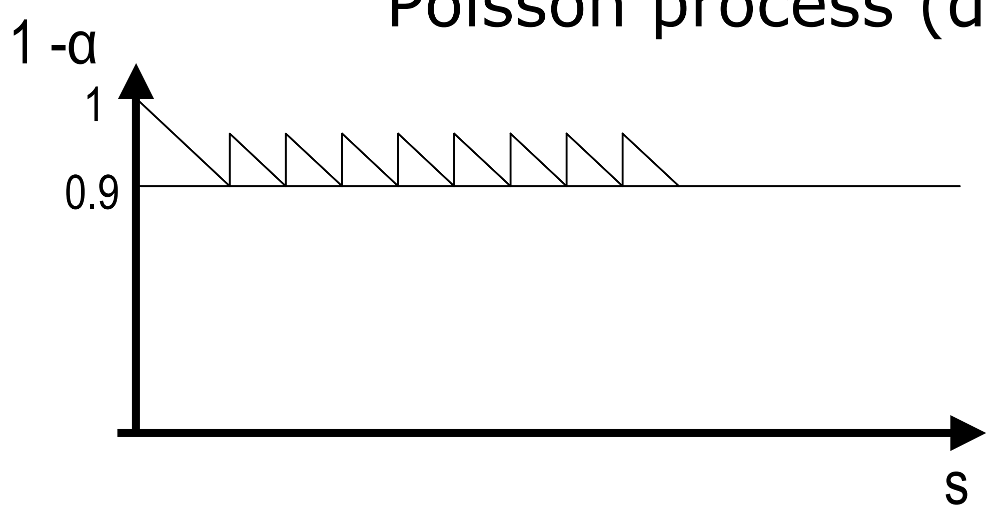

bayesian confidence intervals

Calculating CI in Bayesian way

Remember: Bayes treats s as random variable:

$$1 - \alpha = \int_{s1}^{s2} P(s|n)\, ds$$

Special case: Upper limit, Poisson with background:

$$1 - \alpha = \frac{\sum_{n=0}^{n=n_{obs}} P(n|s+b)}{\sum_{n=0}^{n=n_{obs}} P(n|b)}$$

Adding Nuisance Parameters

- <u>Very simple in the Bayesian approach:</u>

$$1-\alpha = \int_{s1}^{s2}\int_{0}^{\infty} P(\epsilon_{true} s|n) G(\epsilon_{true}|\epsilon_{meas})\, d\epsilon_{true}\, ds$$

Add PDF for uncertainties and average over them

CDF Example: CDF: $B^0_s \rightarrow \mu^+\mu^-$

- $B^0_s \rightarrow \mu^+\mu^-$ are strongly suppressed in the standard model (SM), but can be enhanced by beyond SM physics
- CDF (D0): look for oppositely charged muons
- Calculate branching ratio by normalizing to $B^+ \rightarrow J/\Psi\ K^+$

R. Bernhard, F. Lehner (and some other people from outside Zurich which I don't remember now) hep-ex/0508058, calculate upper limit for CDF and combination (more below)

Example: CDF(Bernhard et. al.): $B^0_s \rightarrow \mu^+\mu^-$

$$1-\alpha = \int_{-\infty}^{s_{true,UP}^{B_s\to\mu^+\mu^-}} \int_0^{\infty} P(\lambda' s_{true}^{B_s\to\mu^+\mu^-} | n^{obs}_{B_s\to\mu^+\mu^-}) G(\lambda'|\lambda) d\lambda' d s_{true}^{B_s\to\mu^+\mu^-}$$

$s_{true}^{B_s\to\mu^+\mu^-}$ → BR($B^0_s \rightarrow \mu^+\mu^-$)

.... neglecting BG

$$\lambda = \frac{acc_{B^+}\cdot\epsilon^{total}_{B^+}}{acc_{B_s}\cdot\epsilon^{total}_{B_s}} \cdot \frac{1}{n^{obs}_{B^+}} \cdot \frac{f_u}{f_s} \cdot BR(B^+ \to J/\Psi K^+)\cdot BR(J/\Psi \to \mu^+\mu^-)$$

$\frac{acc_{B^+}\cdot\epsilon^{total}_{B^+}}{acc_{B_s}\cdot\epsilon^{total}_{B_s}}$ → Acceptance x efficiency

$\frac{f_u}{f_s}$ → B-fragmentation

90 (5 ?) % UL (including uncertainties):

hep-ex/0508058 : **$3.5 \bullet 10^{-7}$**

nb: hep-ex/050836: 1.5e-7

hep-ex/050858: (2xCDF/DO): 1.2e-7

Prior distributions ?

- <u>Informative Priors:</u>
  - Put in what information you really have ….

- <u>Objective priors:</u>
  - Try to be "uninformative", or objective:
    - P(s) = c
      - Variant under translation, but not scale invariant
    - Jeffrey's rule: (scale invariant):

      P(s) = 1/s, 1/sqrt(s)

The choice of objective priors is a science in itself and some people think objective priors are bad anyway …

→ at least don't just blindly use flat priors → see next slides ….

coverage of Bayesian intervals ?

- The CDF statistics committee recommends Bayesian methods but likes to test their frequentist properties

- We will see how frequentist coverage can help to identify badly chosen priors

- Consider (Heinrich, PhyStat 2005):
  - Poisson process with $\varepsilon s + b$
  - ε and b determined by subsidiary measurements
  - Priors (to start with):
    - flat for $s > 0$
    - flat for $\varepsilon > 0$[1)]
    - flat for $b > 0$[1)]

1) Flat prior in Bayesian analysis of ε,b in subsidiary measurement

coverage for single experiments

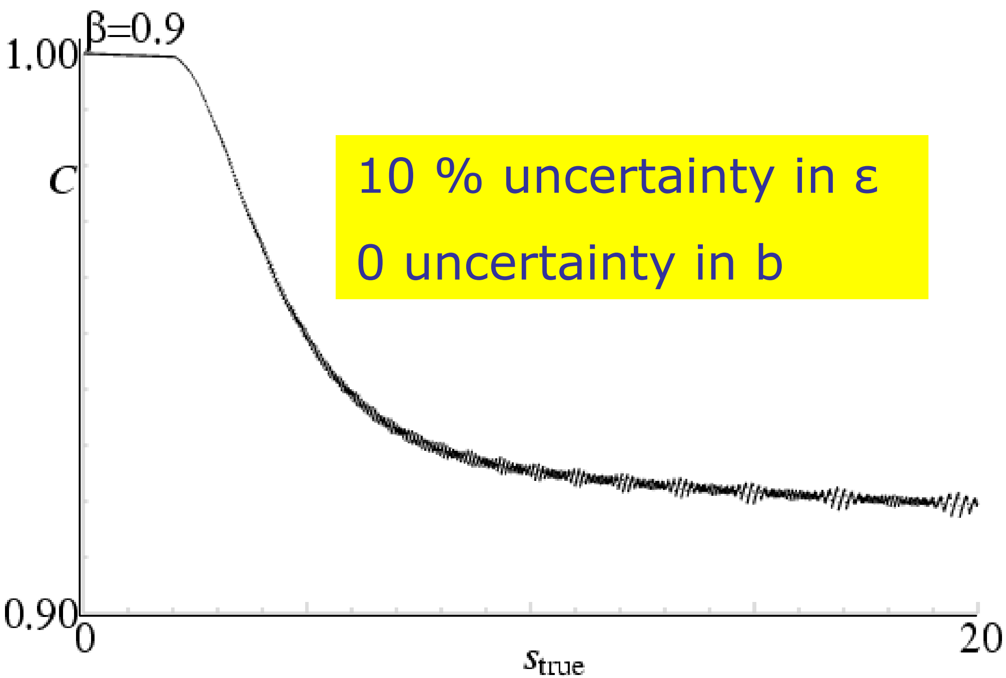

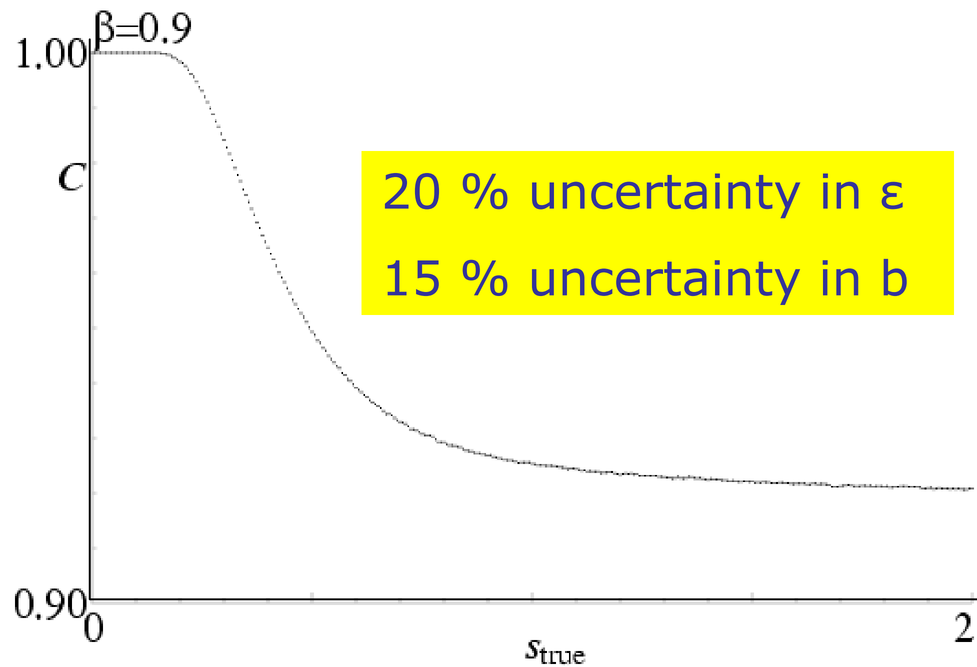

- over-coverage

- larger uncertainties lead to slightly larger coverage

J. Heinrich, Proceedings PhyStat 05

Combined experiments

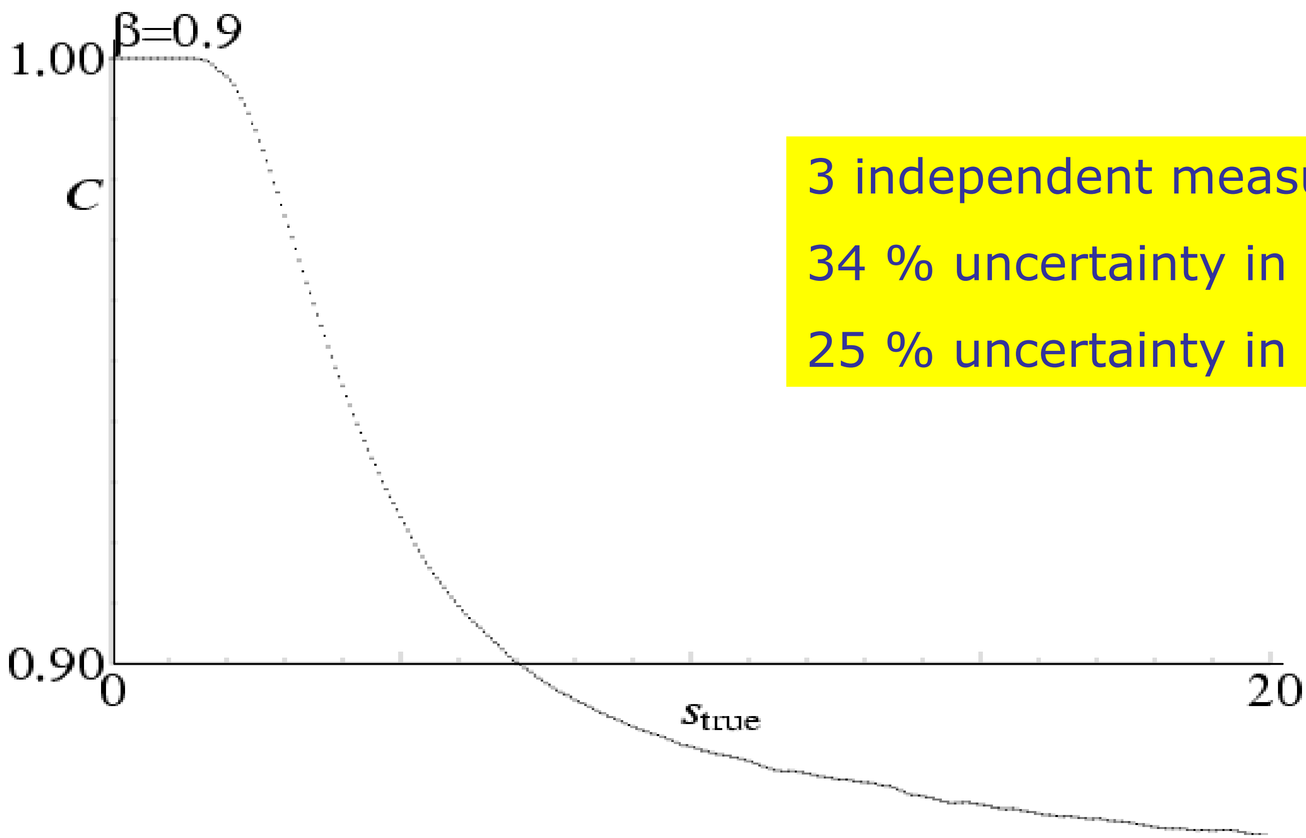

3 independent measurements

34 % uncertainty in ε (per channel)

25 % uncertainty in b (per channel)

J. Heinrich, Proceedings PhyStat 05

Combined Experiments cont'd

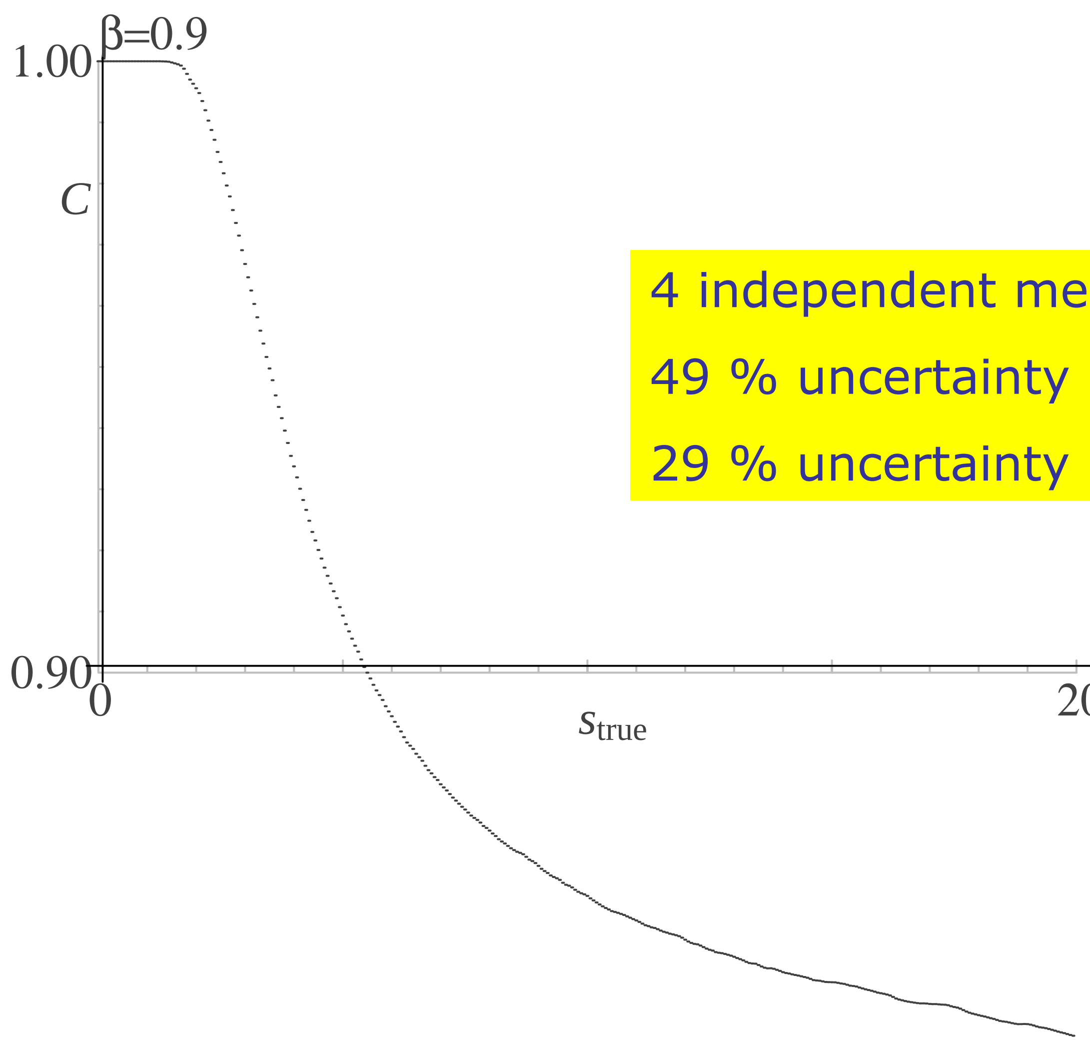

4 independent measurements

49 % uncertainty in ε (per channel)

29 % uncertainty in b (per channel)

J. Heinrich, Proceedings PhyStat 05

What happened ?

- <u>In hindisight, trivial:</u>
  - A flat prior for each of N channels leads to a
    - **ε^{N-1} prior for total acceptance**
- <u>So: different prior: in this case simple</u>
  - **$1/\varepsilon$ in each channel**

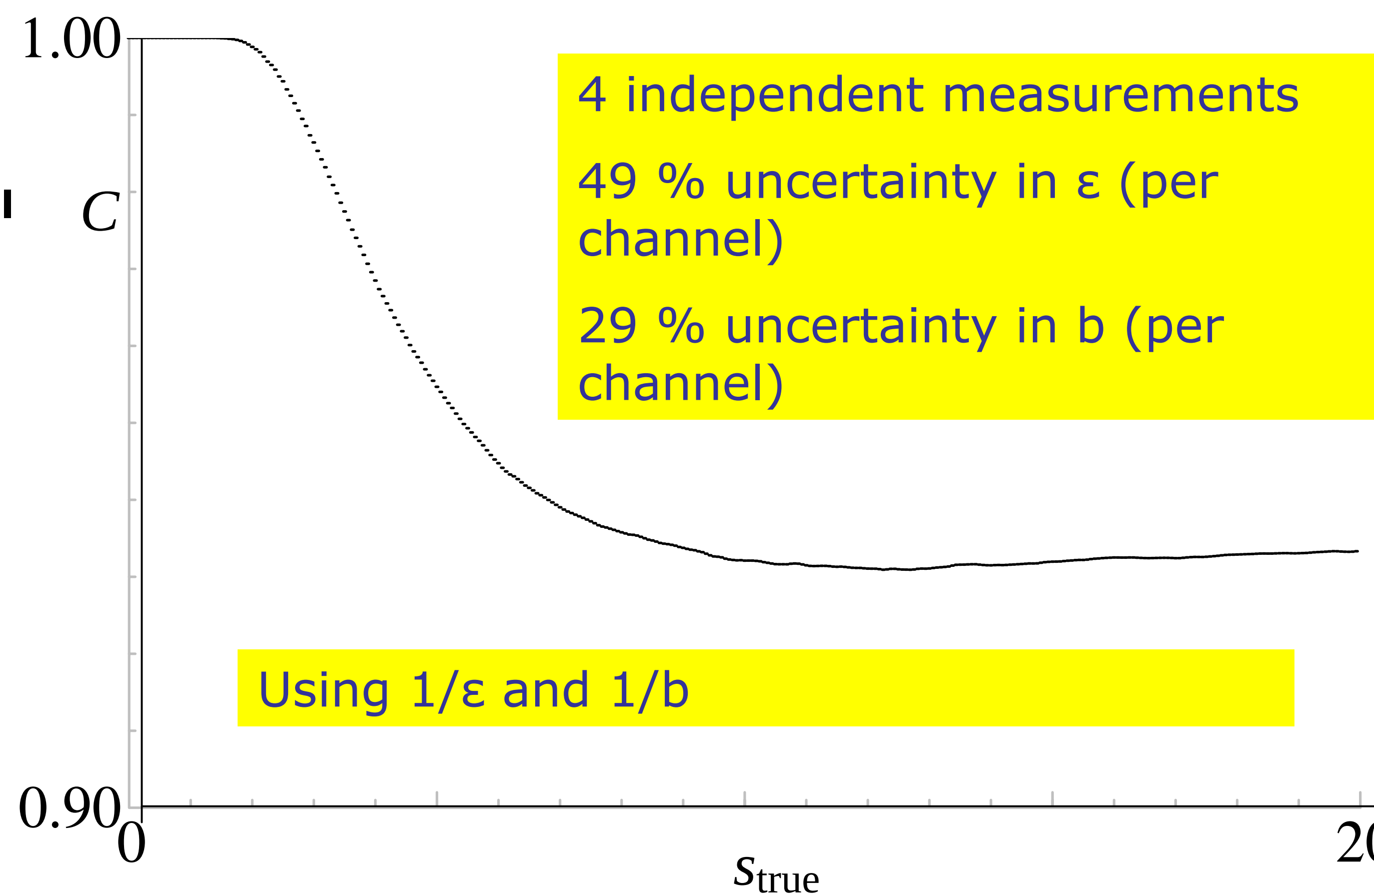

J. Heinrich, Proceedings PhyStat 05

frequentist confidence intervals

What is the problem ?

- We would like to make probability statements about the parameter of a pdf
- But remember, the parameter is not a random variable so what to do ?
- Neyman (1937): make probability statements about sets of confidence intervals, instead

Neyman construction

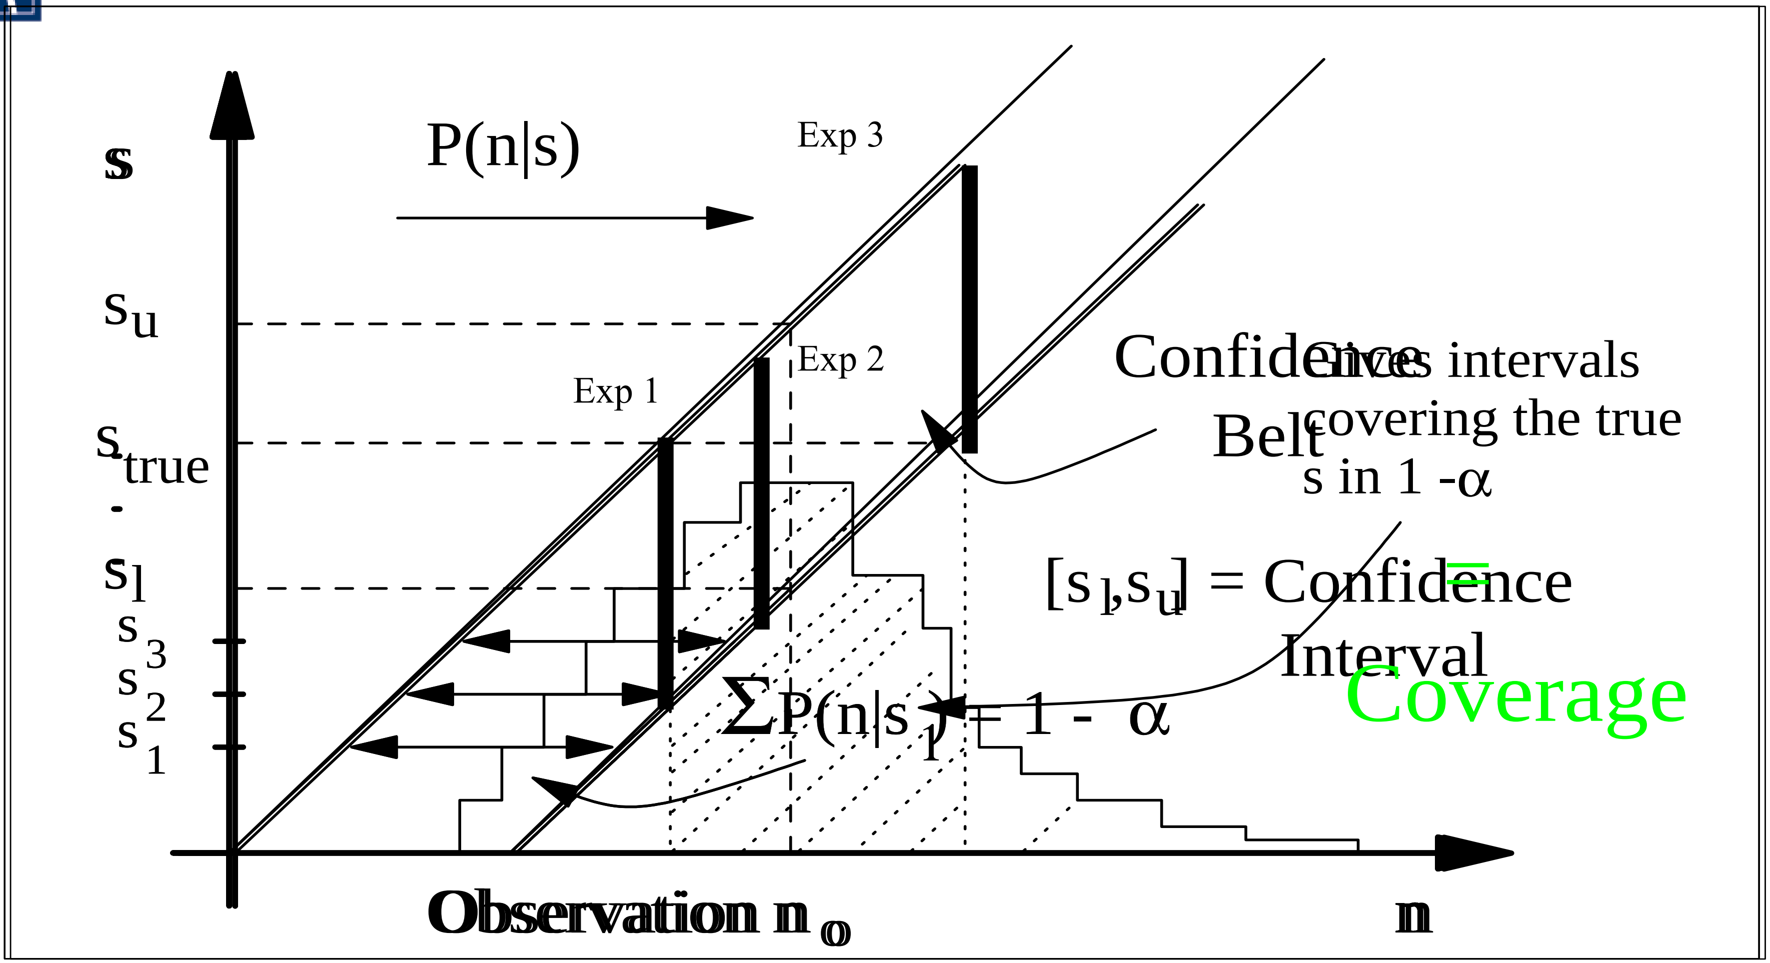

One additional degree of freedom: ORDER in which you inlcude the n into the belt

J. Neyman, Phil. Trans. Roy. Soc. London A, 333, (1937)

Which observations to include then ?

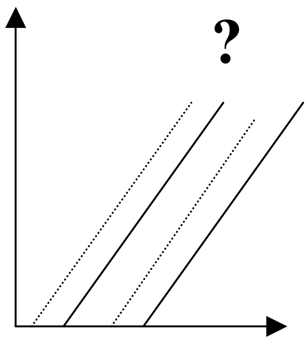

Neyman's traditional choices:

$$P(n < n_1(s)|s) = P(n > n_2(s)|s) = \frac{\alpha}{2}$$

For Central Intervals

$$P(n < n_1(s)|s) = \alpha;\ n_2 \to \infty$$

Upper limits

Likelihood ratio (FC) ordering

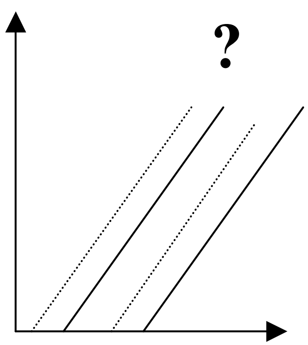

- Calculate likelihood ratio:

$$R = \frac{\mathcal{L}(n|b+s)}{\mathcal{L}(n|b+s_{best})}$$

- Rank n according to likelihood ratio
- Include n in descending order of the ratio until sum condition fulfilled
- Unifies “upper limit” and “central limit” and never gives empty intervals.

G. Feldman & R. Cousins, Phys. Rev D57:3873-3889

Why FC ordering ?

- We have to consider realistic ensemble !!!!

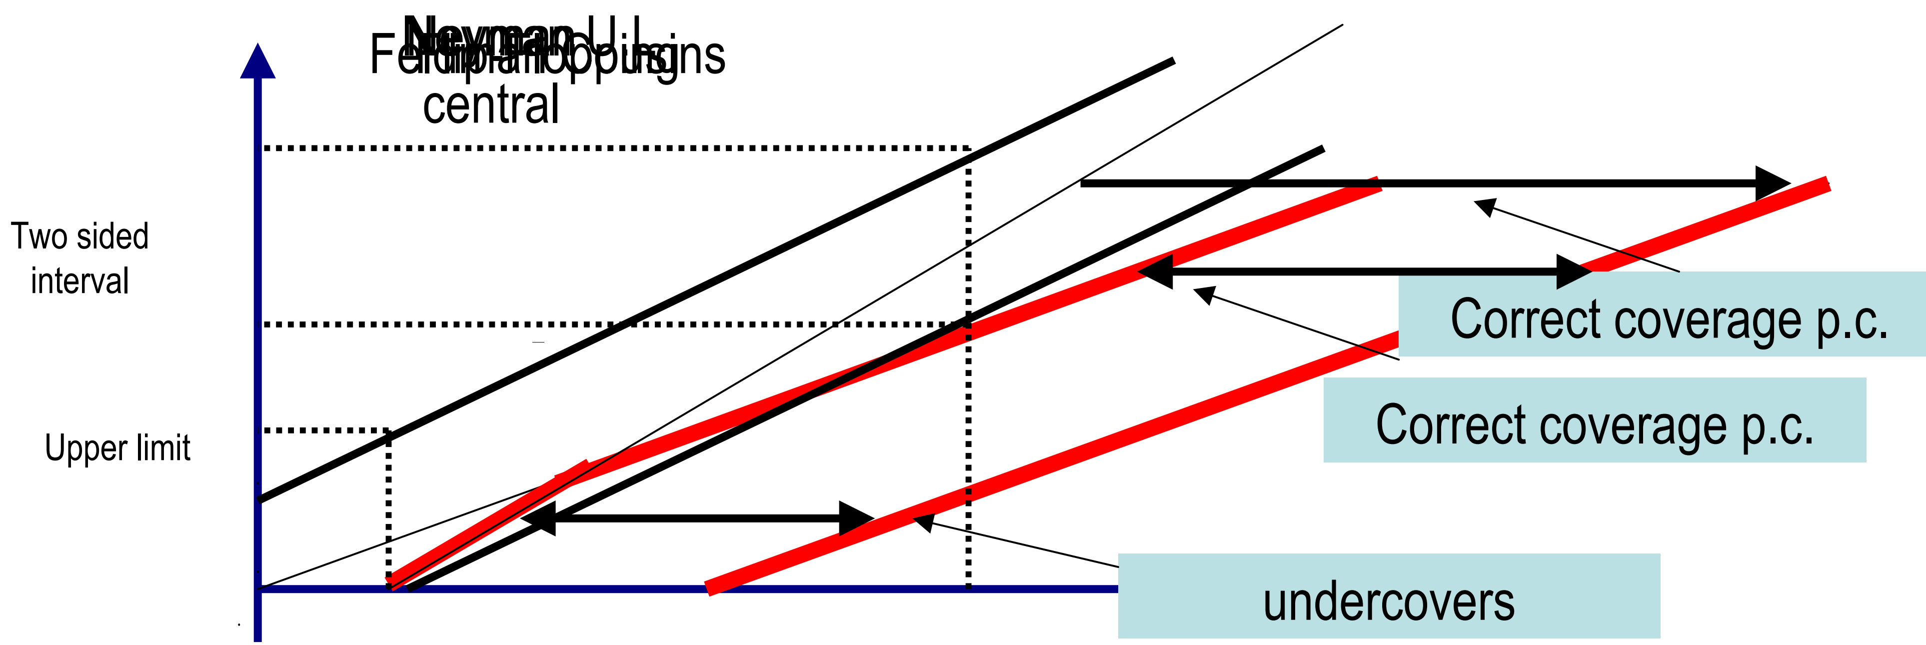

FC ordering: coverage

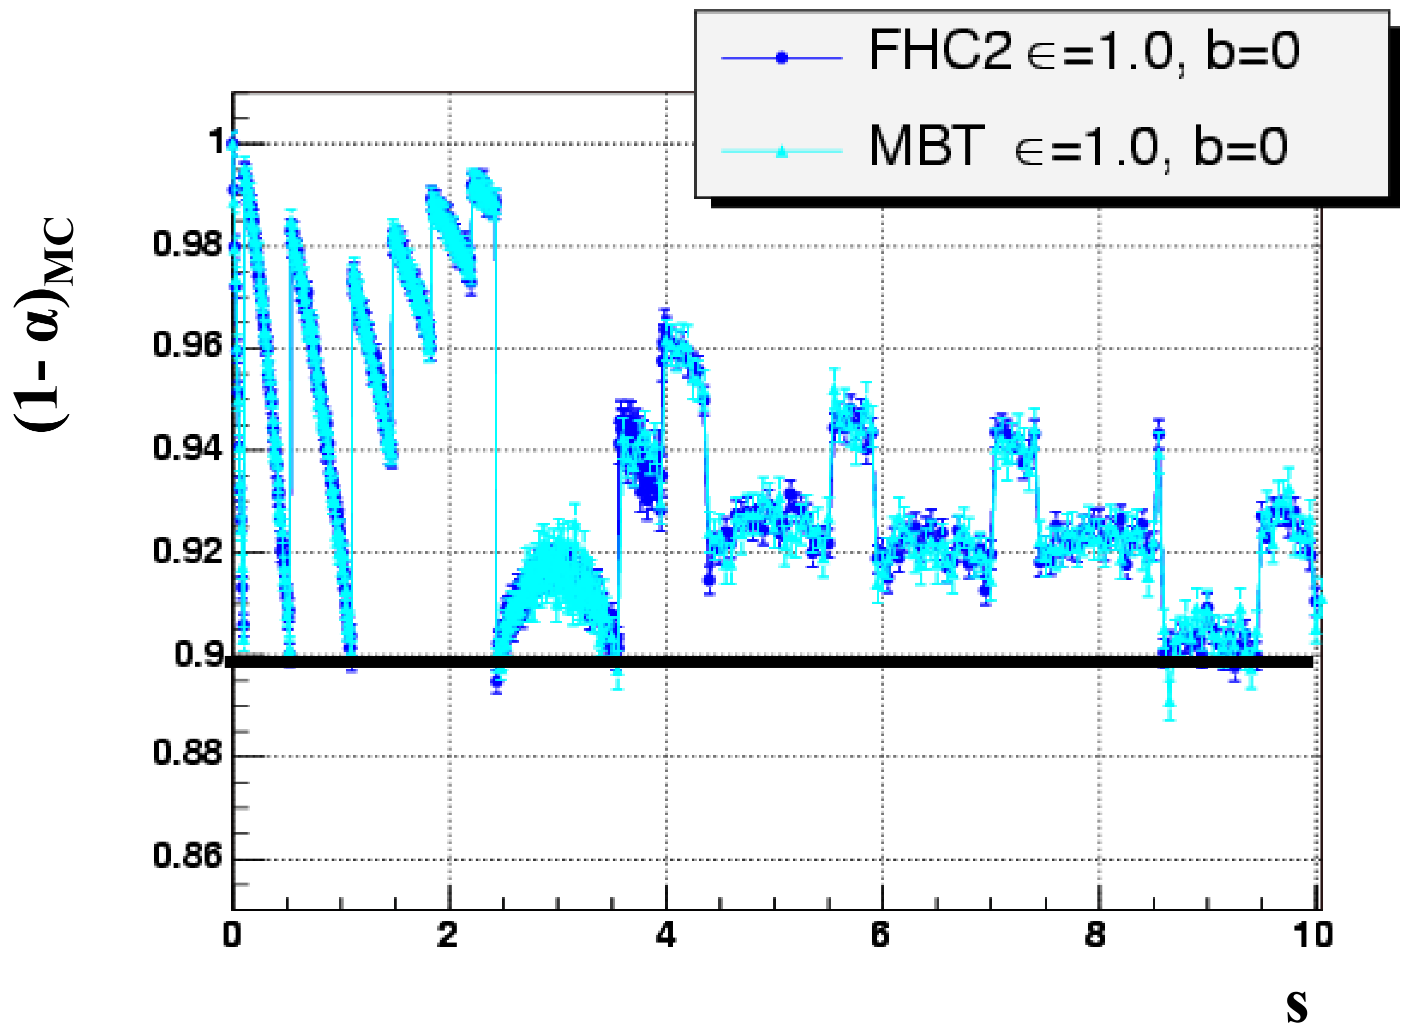

KARMEN anomaly

Liquid scintillation calorimeter and PMTs, beam-stop neutrinos

L~ 20 m

$$\nu_e + {}^{12}\mathrm{C} \rightarrow {}^{12}\mathrm{N}_{g.s.} + \boxed{e^-}$$
$$\downarrow$$
$${}^{12}\mathrm{C}_{g.s.} + \boxed{e^+} + \nu_e$$

<u>Neutrino experiment which</u>

→ sees no events
→ expects b = **2.9**
→ FC gives upper limit of **1.1**

<u>If experiment</u>

→ sees no events
→ expects b = **0**
→ FC gives upper limit of **2.4**

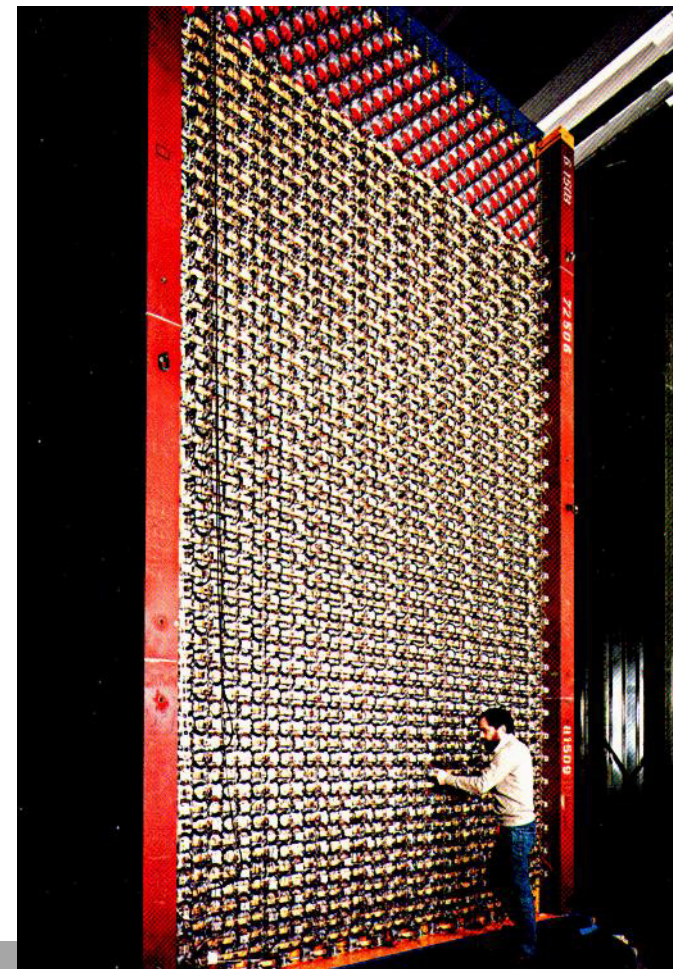

Solutions to the KARMEN anomaly: **none** generally accepted !

- Roe-Woodroofe
  - •Neyman Construction, FC ordering with renormalization (conditioning)

$$q^{n_o}_{s+b}(n) = \begin{cases} \dfrac{p(n)_{s+b}}{\sum_{n'=0}^{n_o} p(n')_b} & \text{if } n \le n_o, \\ \dfrac{\sum_{n'=0}^{n_o} p(n')_b p(n-n')_s}{\sum_{n'=0}^{n_o} p(n')_b} & \text{if } n > n_o. \end{cases}$$

B. Roe & M. Woodroofe, Phys.Rev.D60:053009,1999

Under-covers

B = 0	B =4
2.4	2.4

- Ciampolillo,Mandelkern & Schultz
  - based on MLE, including constraint (biased)

S. Ciampolillo, Nuovo Cim.A111:1415-1430, 1998
M. Mandelkern &. J. Schultz J.Math.Phys.41:5701-5709,2000

Over-covers, seriously

B = 0	B =3.0
2.6	4.7

- Strong confidence intervals
  - consider only subset of intervals of observational space

G. Punzi, Proceedings , Durham 2002

90 %
sCL

B = 0	B = 4
2.5	2.3

Generally recommended: present "sensitivity" (mean limit one would obtain in case of no signal

Inserting Nuisance Parameters

- The Bayesian Way
- The Likelihood Principle
  - Profile Likelihood with χ^2 approximation (the jack-knife)
- Full Neyman Construction with nuisance parameters
  - Profile Likelihood Ratio
  - Projection method

Including Systematics: the Bayesian way

- Main idea: fold PDF (for prime process) with a PDF describing the uncertainties

$$q(n)_{s+b} = \frac{1}{2\pi\sigma_b\sigma_\epsilon} \times$$
$$\int_0^\infty \int_0^\infty p(n)_{b'+\epsilon' s} \; e^{\frac{-(b-b')^2}{2\sigma_b^2}} \; e^{\frac{-(1-\epsilon')^2}{2\sigma_\epsilon^2}} \, db' \, d\epsilon'$$

- Perform Neyman-Construction with this new PDF

- In what follows we will assume FC ordering

Integral is performed in true variables → Bayesian

R. Cousins & V. Highland Nucl. Inst. Meth. A320:331-335,1992

J.C et. al. , Phys. Rev D67:012002,2003

Commercial break: pole++

- Extension of FORTRAN program pole which includes Bayesian treatment in FC ordering Neyman construction
  - treats $P(n|\varepsilon s + b)$
- Consists of C++ classes:
  - ***Pole*** calculate likits
  - ***Coverage*** coverage studies
  - ***Combine*** combine experiments
- Nuisance parameters
  - supports flat, log-normal and Gaussian uncertainties in efficiency and background
  - Correlations (multi-variate distributions and uncorrelated case)
- Code and documentation available from:
  - http://cern.ch/tegen/statistics.html

F.Tegenfeldt & J.C Nucl. Instr. Meth.A539:407-413, 2005
J.C et. al. , Phys. Rev D67:012002,2003

$B^0_s \rightarrow \mu^+\mu^-$ revisited: FC ordering with Bayesian treatment of uncertainties

	CDF1	CDF 2
Eff uncertainty [%]	18.2	16.0
Eff uncertainty [%]	20.3	19.2
Corr. eff. Uncertainty [%]	15.5	
95 % CI [10^{-7}]	2.5	4.3
95 % combined limit [10^{-7}]	**1.7 (2.0 in Bernhard et. al.)**	

Remember that flat prior leads to under-coverage for combined experiments (should be checked for Bernhard et. al., less a problem here (but should be checked)

J.C & F. Tegenfeldt , Proceedings PhyStat 05, physics/0511055

coverage of (FC ordering with Bayesian treatment of uncertainties)

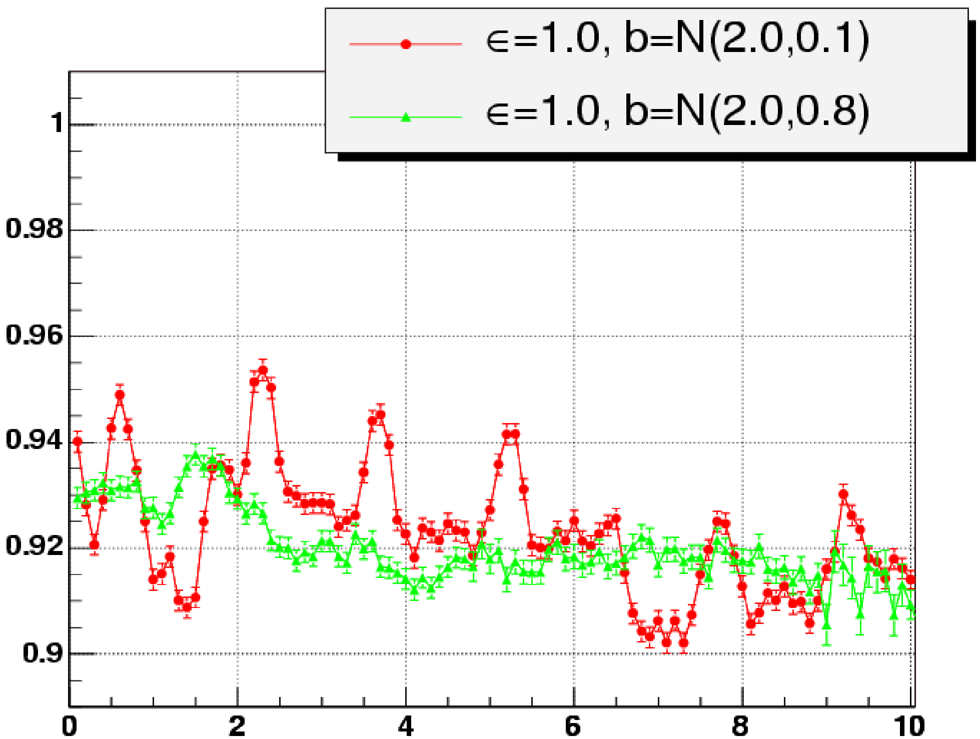

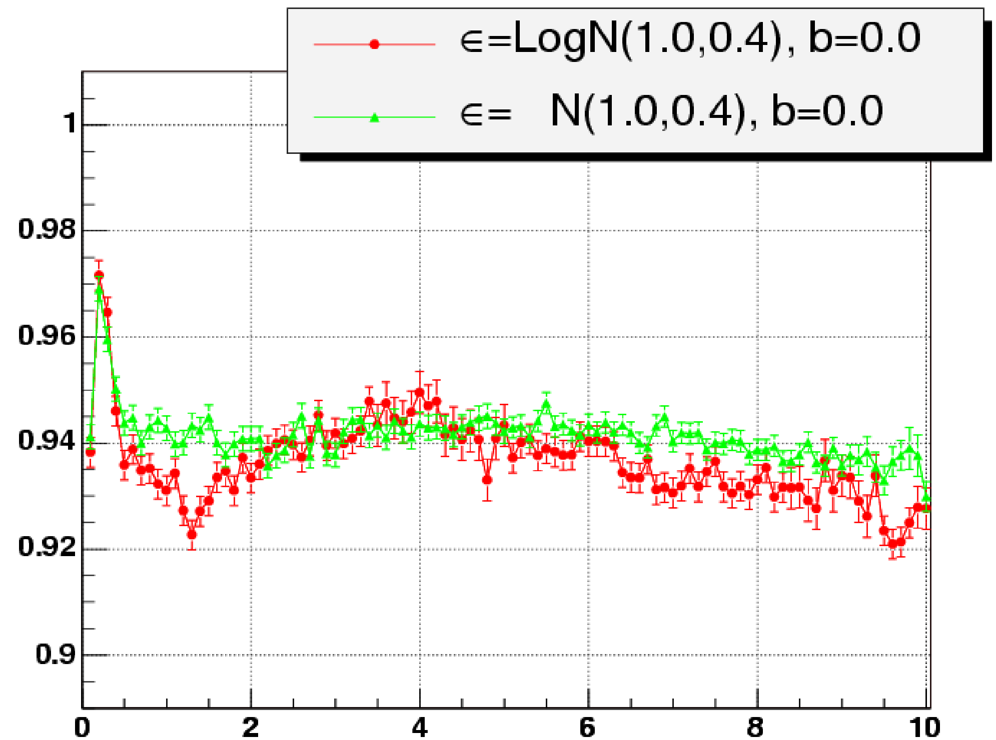

F.Tegenfeldt & J.C Nucl. Instr. Meth.A539:407-413, 2005

likelihood function in presence of nuisance parameters

- For example: add background as nuisance parameter

$$\mathcal{L}(n, b_{meas}|s, b_{true}) = P(n|s, b_{true}) \cdot P(b_{meas}|b_{true})$$

Subsidiary measurement

Profile Likelihood with χ^2 approximation: the jack-knife

Meas. back

MLE of b given s

$$\lambda(s) = \frac{\mathcal{L}(n, b_m | s, \hat{b}(s))}{\mathcal{L}(n, b_m | \hat{s}, \hat{b})}$$

MLE of b and s given observations

To extract limits:

$$-2\ln\lambda \approx \chi^2$$

$$\chi^2 - \chi^2_{min} = 2.706 \quad \equiv 90\% \text{ C.I.}$$

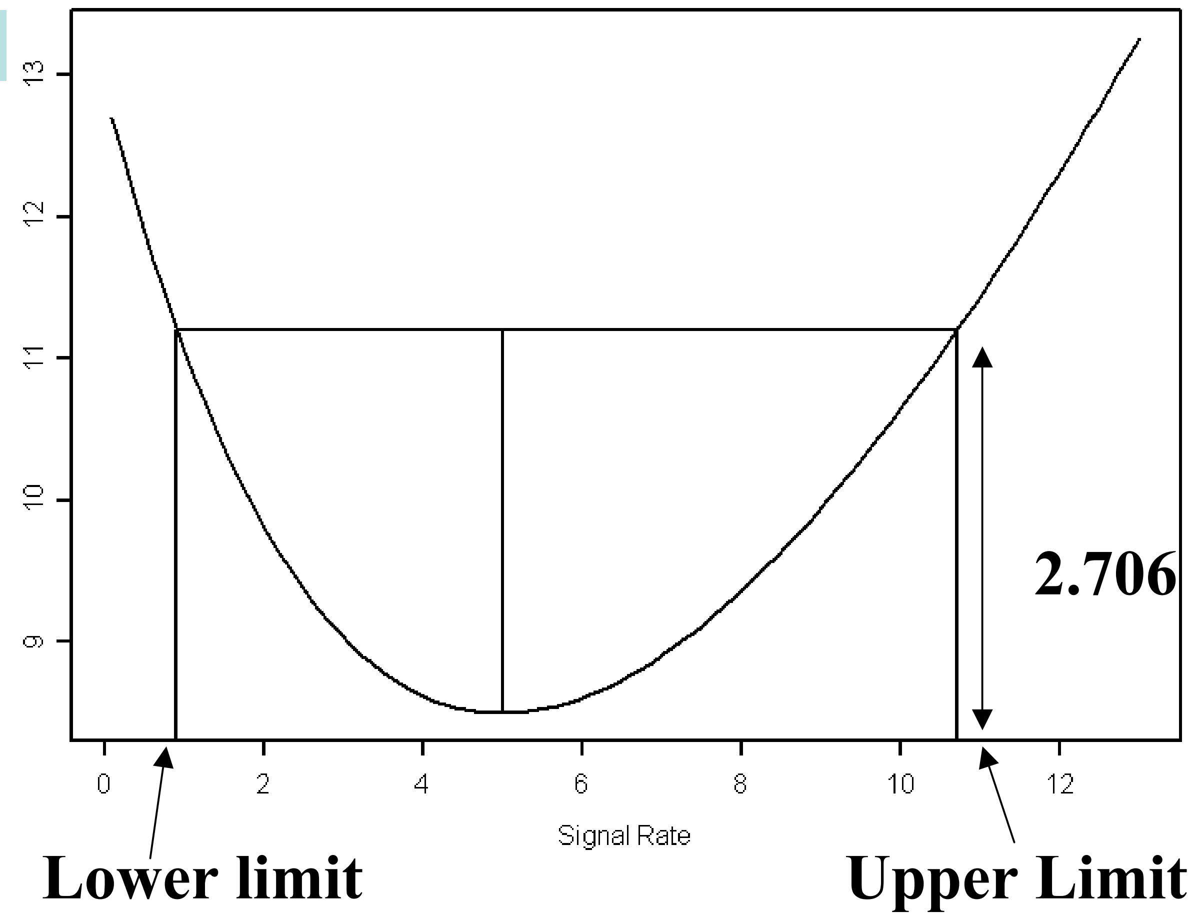

F. James, e.g. Computer Phys. Comm. 20 (1980) 29 -35
W. Rolke, A. Lopez, J.C. Nucl. Inst.Meth A 551 (2005) 493-503

From MINUIT manual

- See F. James, MINUIT Reference Manual, CERN Library Long Writeup D506, p.5:

„The MINOS error for a given parameter is defined as as the change in the value of the parameter that causes the F′ to increase by the amount UP, where F′ is the minimum w.r.t to all other free parameters".

Confidence Interval

$\Delta X^2 = 2.71$ (90%),

$\Delta X^2 = 1.07$ (70 %)

Profile Likelihood

Jackknife: coverage

Background: Poisson (unc ~ 20 % -- 40 %)

Efficiency: binomial (unc ~ 12%)

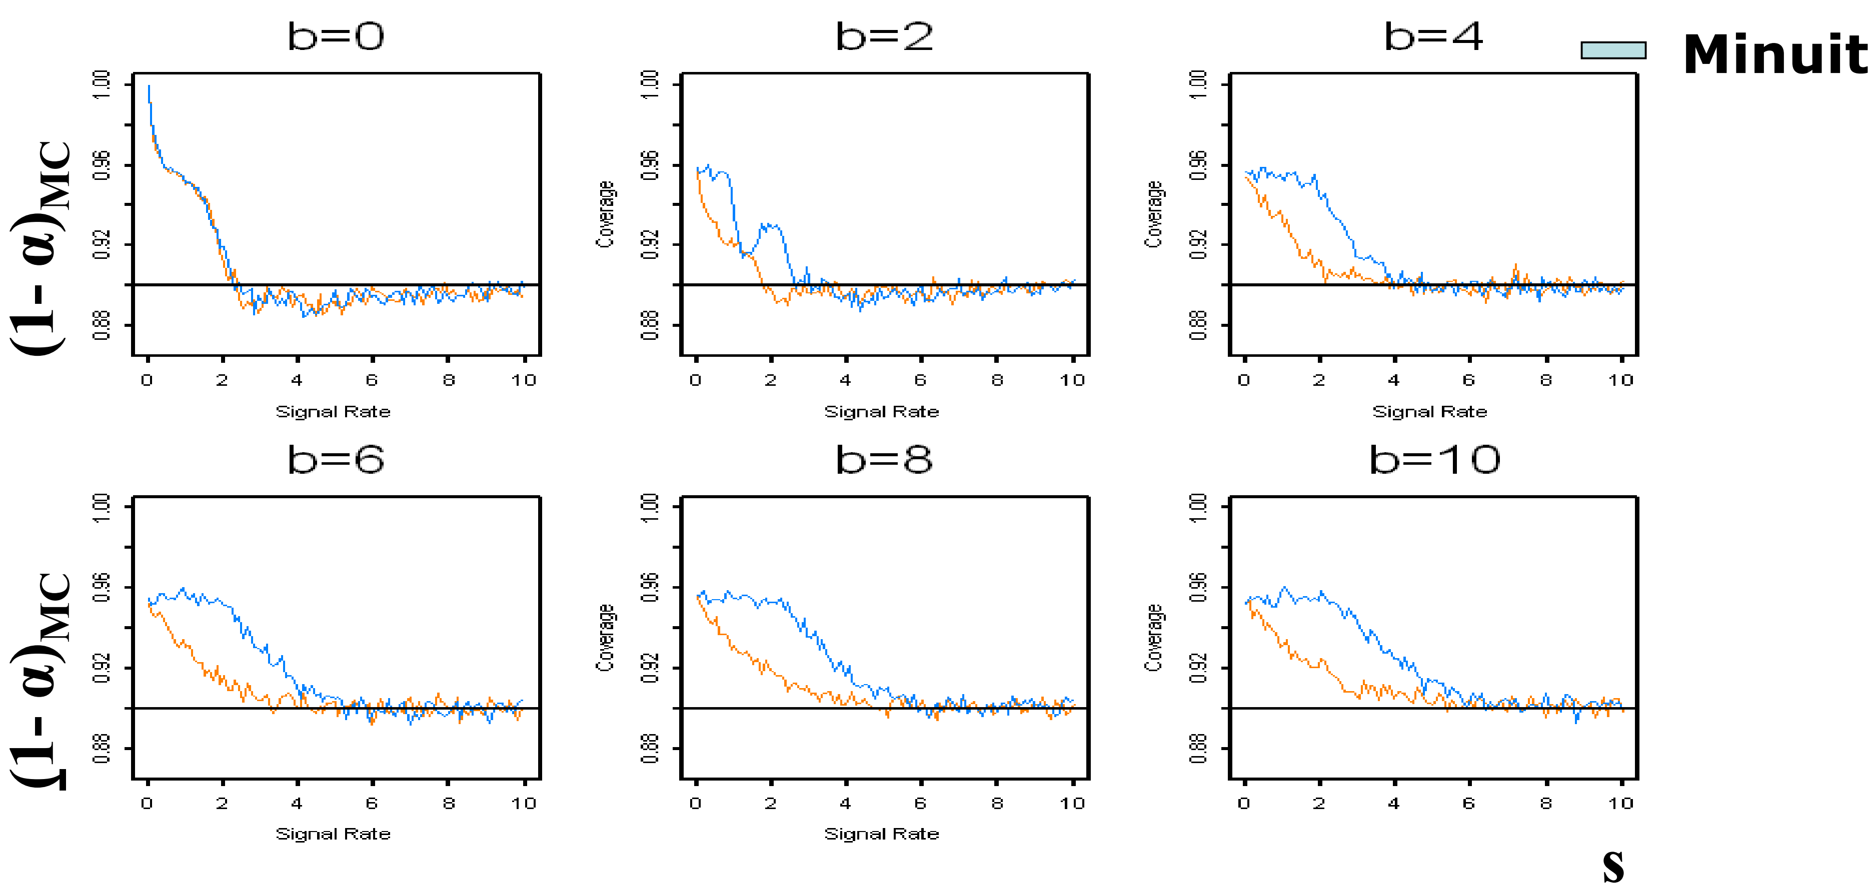

W. Rolke, A. Lopez, J.C. Nucl. Inst.Meth A 551 (2005) 493-503

Commercial break: TRolke

- Implemented with 7 different probability models

  - $P(n|\varepsilon s+b)$
  - Gaussian uncertainties in signal efficiency /background pred.
  - Poisson uncertainties in background
  - Binomial uncertainty in signal efficiency

Part of ROOT since v. 4.002

http://root.cern.ch/root/htmldoc/TRolke.html

N.B: MINUIT is useable for all kinds of problems as long as you can write down the likelihood function, but TRolke applies some improvements to the 7 probability models inlcuded

Full contruction: Profile Likelihood with Likelihood Ratio ordering (FC profile)

- Start with 4-dimensional space
- to reduce the dimensionality: maximize w.r.t ε_{true} and then order n as in FC

Maximize under cond s

$$R(n,s) = \frac{\mathcal{L}(n, \epsilon_{meas} | s, \epsilon_{max})}{\mathcal{L}(n, \epsilon_{meas} | s_{max}, \epsilon'_{max})}$$

Maximize unconditionally

- Do ordering as in case without nuisance parameters in (n,s) plane.

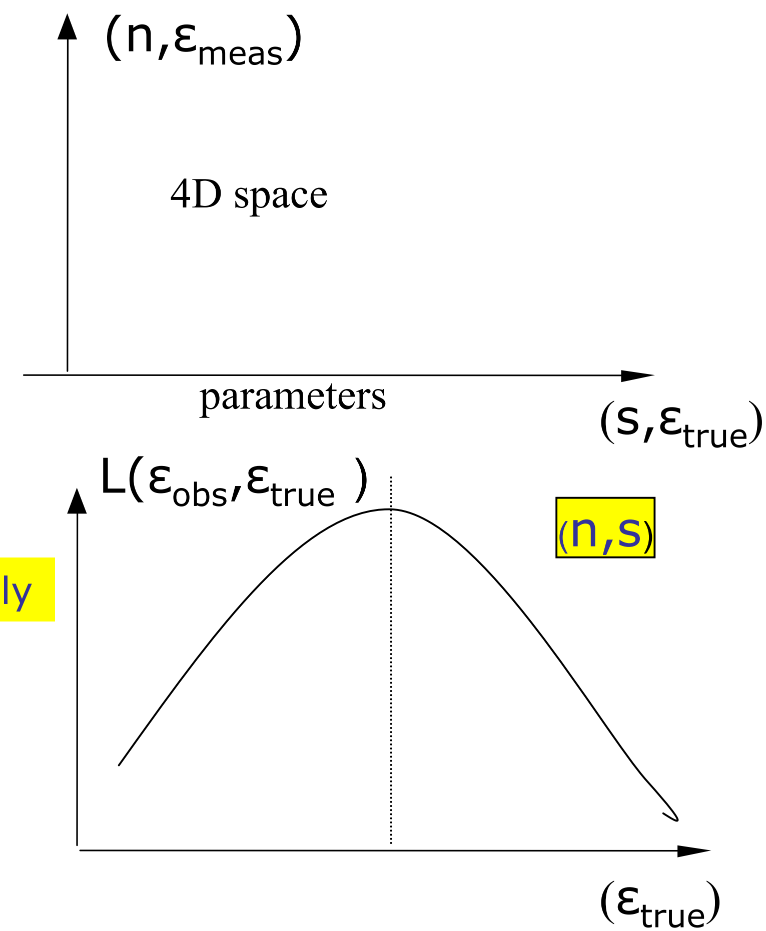

M. Kendall & A. Stuart, "The Advanced Theory of Statistics" vol 2, ch 24

Full Construction: projection method

- remember Likelihood function:

$$\mathcal{L}(n, \epsilon_{meas}|s, \epsilon_{true}) = P(n|\epsilon_{true}, s) \cdot P(\epsilon_{meas}|\epsilon_{true})$$

- define (>) 4 D confidence volume

$$\int_{r(n,\epsilon_{meas};s,\epsilon_{true})>c} P(n, \epsilon_{meas}|s, \epsilon_{true})\, dn\, d\epsilon_{meas} > 1 - \alpha$$

dering
nction
g like
nks in
)

- upon measurement:(n0,ε0) you get 2D hypersurface
- project 2D hypersurface on s axis to get limit on s only
- generally leads to over-coverage
- but can be avoided if appropriate ordering principle is applied in 4(n)D (see next slides)

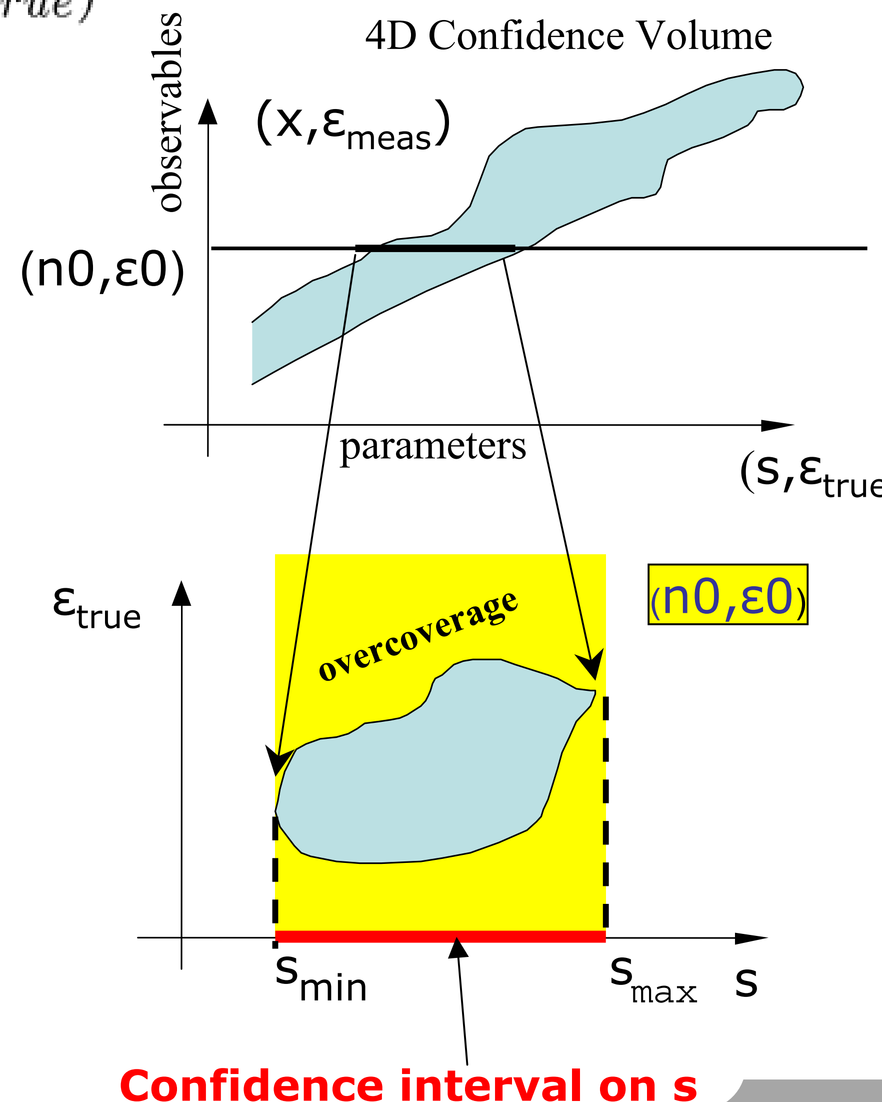

CHOOZ experiment

Liquid scintillator calorimeter and PMTs, reactor neutrinos

L~ 1 km, inverse β decay

- Reactor neutrino experiment located in France
- Observable: $x = \frac{\bar{\nu}_e^{meas}}{\bar{\nu}_e^{exp}}$
- Paremeters: $\mu = (\sin^2(2\theta), \Delta m^2).$
- Likelihood function:

$$p(x|\mu,\alpha) \propto \exp -1/2 \left(\frac{x - \alpha \mathcal{P}(\mu)}{\sigma_{stat}}\right)^2 \cdot \exp -1/2 \left(\frac{\alpha - 1}{\sigma_{sys}}\right)$$

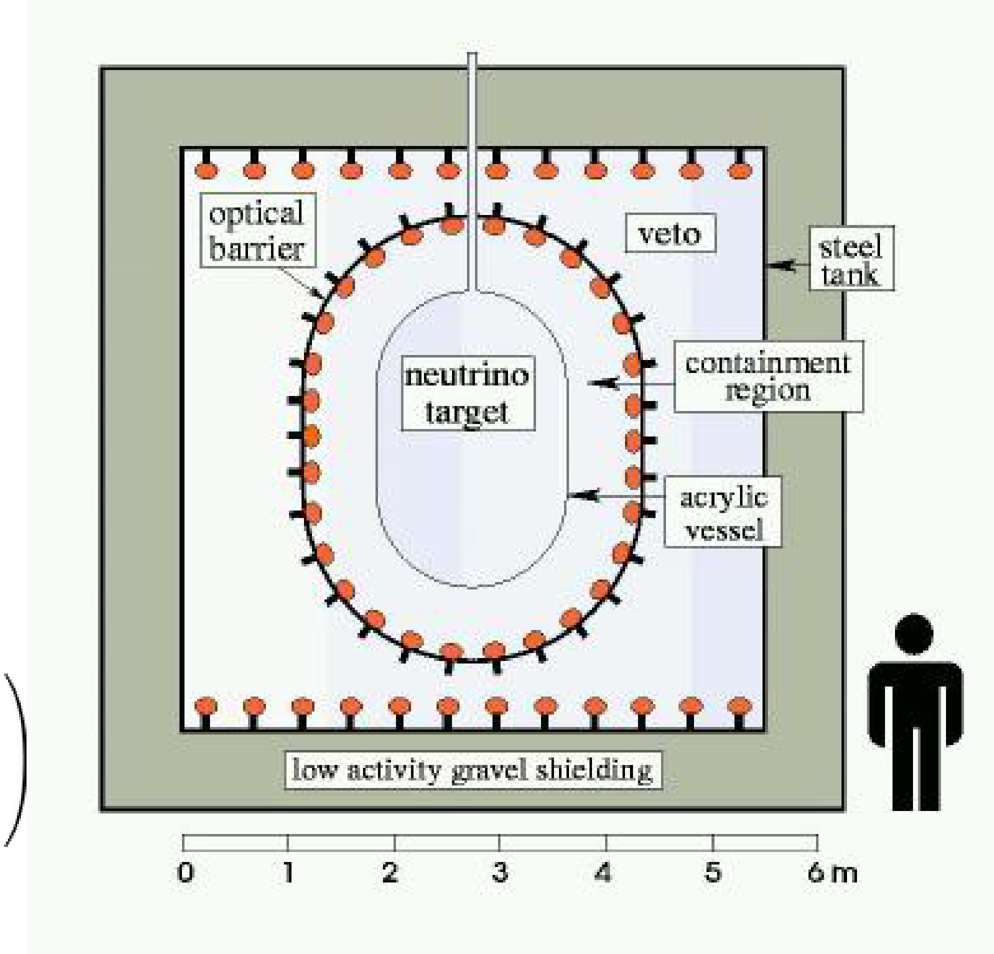

G. Signorelli et. al. , Proceedings, Durham 2002

CHOOZ: full construction and profile likelihood

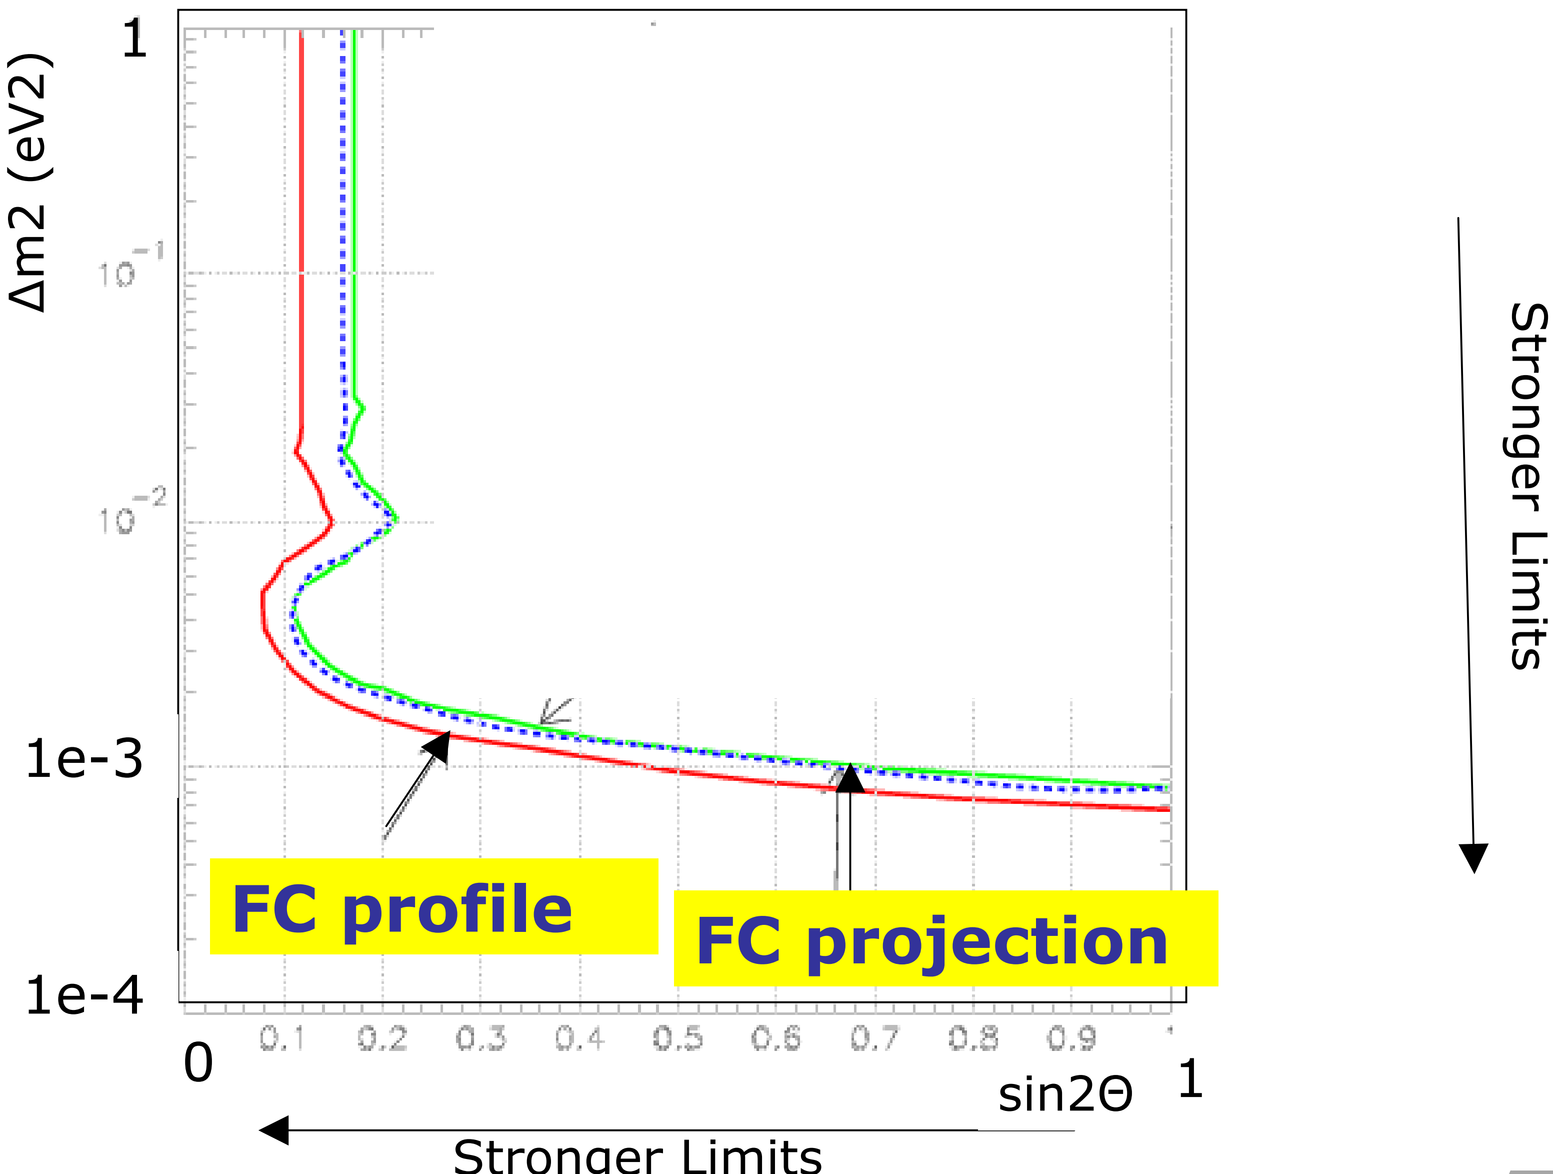

G. Signorelli et. al. , Proceedings, Durham 2002

Projection method with appropriate ordering I

rdering function:
Punzi, PhyStat05)

$$r(n, \epsilon_{meas}; s) = \int_{r_0(n')<r_0(n)} p(n'|\epsilon_{meas}; s, \epsilon_{max}(\epsilon_{meas})dn'$$

Can be any ordering in prime observable space, in this case

Upper Limit: $r_0(n) = n$

Poisson signal, Gauss eff. Unc (10 %)

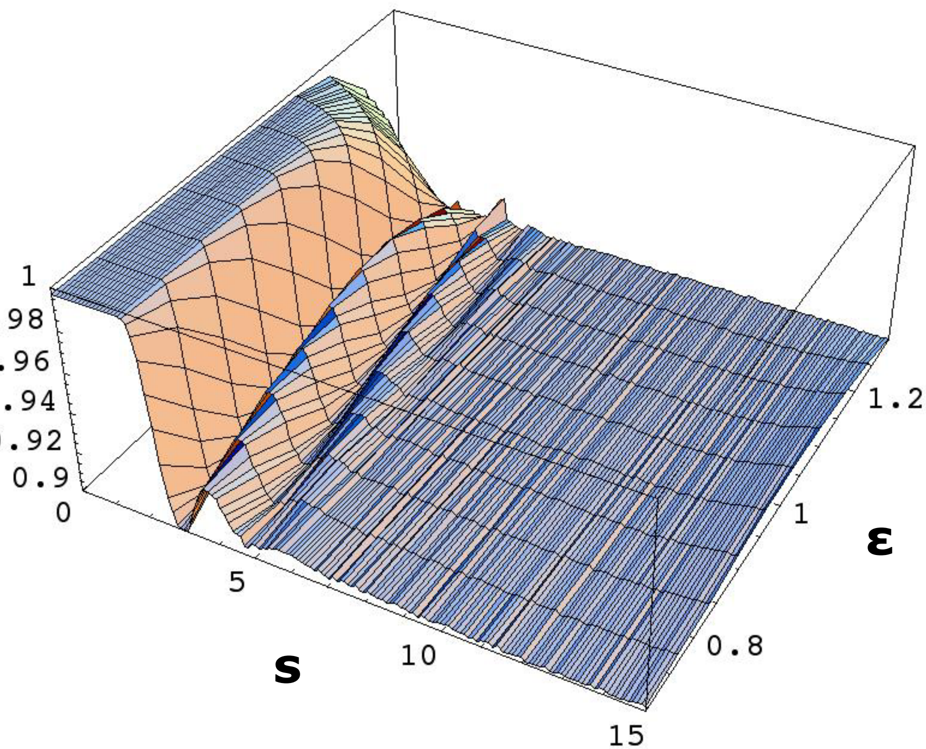

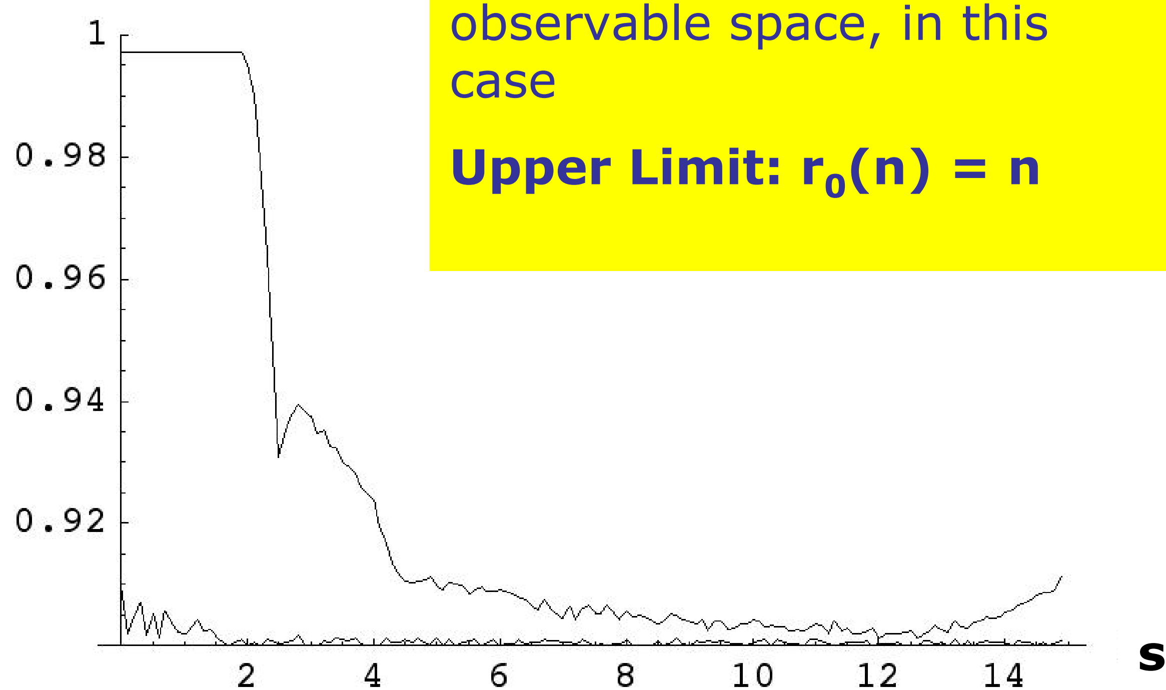

G. Punzi, Proceedings PhyStat 2005, physics/0511202

Confidence Intervals for new particle searches at LHC?

- Basic idea: calculate 5 σ confidence interval and claim discovery if s = 0 is not included.

- Straw-man model:

$$\mathcal{L}(n, n_b | s, b) = P(n|s+b)P(n_b|\tau b)$$

Poisson or Gauss

- Bayesian under-covers badly (add 16 events to get correct significance)

- Jack-knife is the only method considered here which gives coverage (exc. full construction)

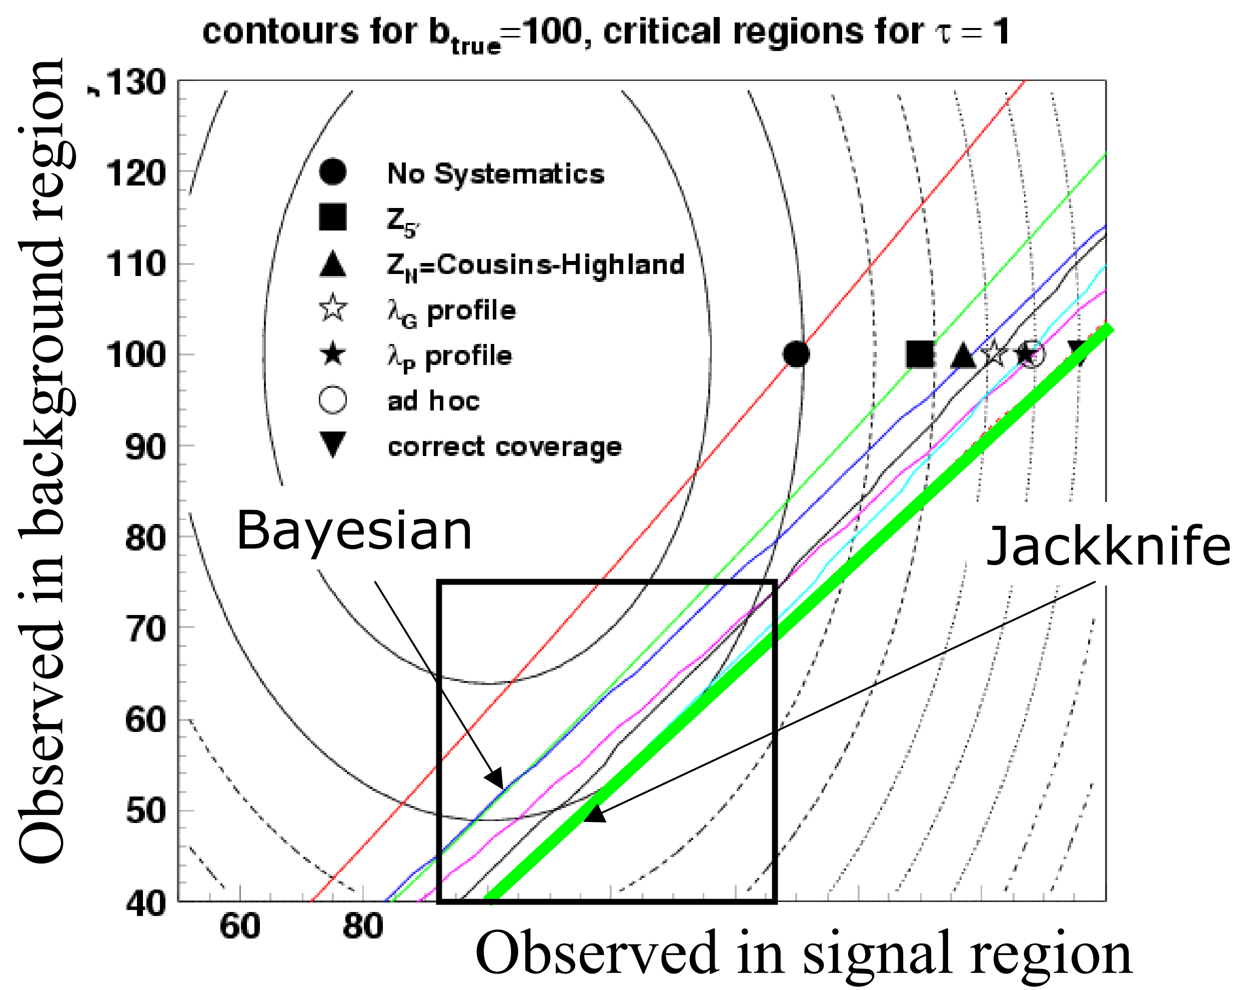

K. S. Cranmer, Proceedings PhyStat 2005

Higgs search

- Not straw-man, but still TOY (e.g. qqH→ qqττ → qqll(missEt):

$$\mathcal{L}(n, n_b, \vec{x}|s,b) = P(n|s+b)P(n_b|\tau b)\prod_{i=1}^{n_{obs}} \frac{s f_s(x|s) + b f_b(x|b)}{s+b}$$

$$f_s(x|s) = \frac{1}{N}e^{-a(x-1)} \quad f_b(x|b) = \frac{1}{N}e^{-ax}$$

ominal cov = 5σ

	b	τ	a	σ_{calc}	Pow
	20	1	-3	5.2	-
	20	1	-1	5.1	-
0	50	1	-3	5.1	0.9997

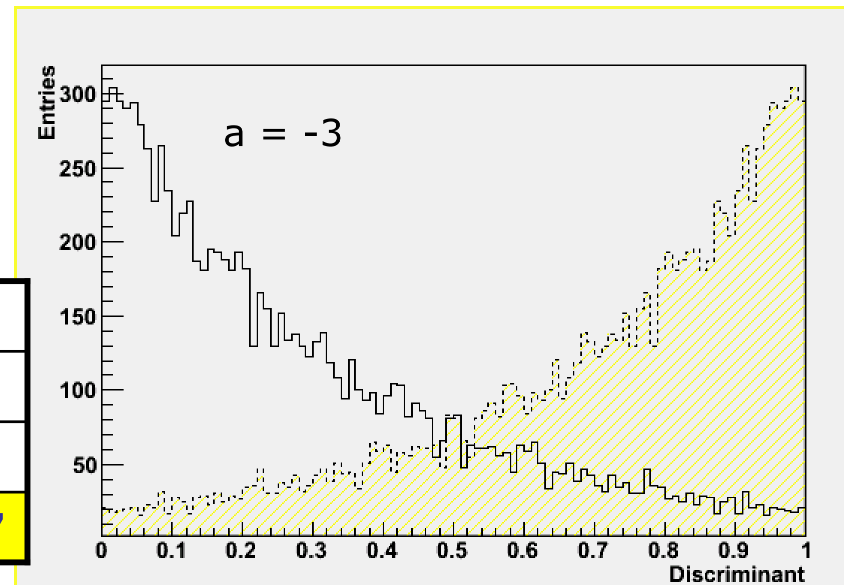

VERY PRELIMINARY

K. S. Cranmer, J. C. in preparation

Overview over presented methods
(slide added after talk)

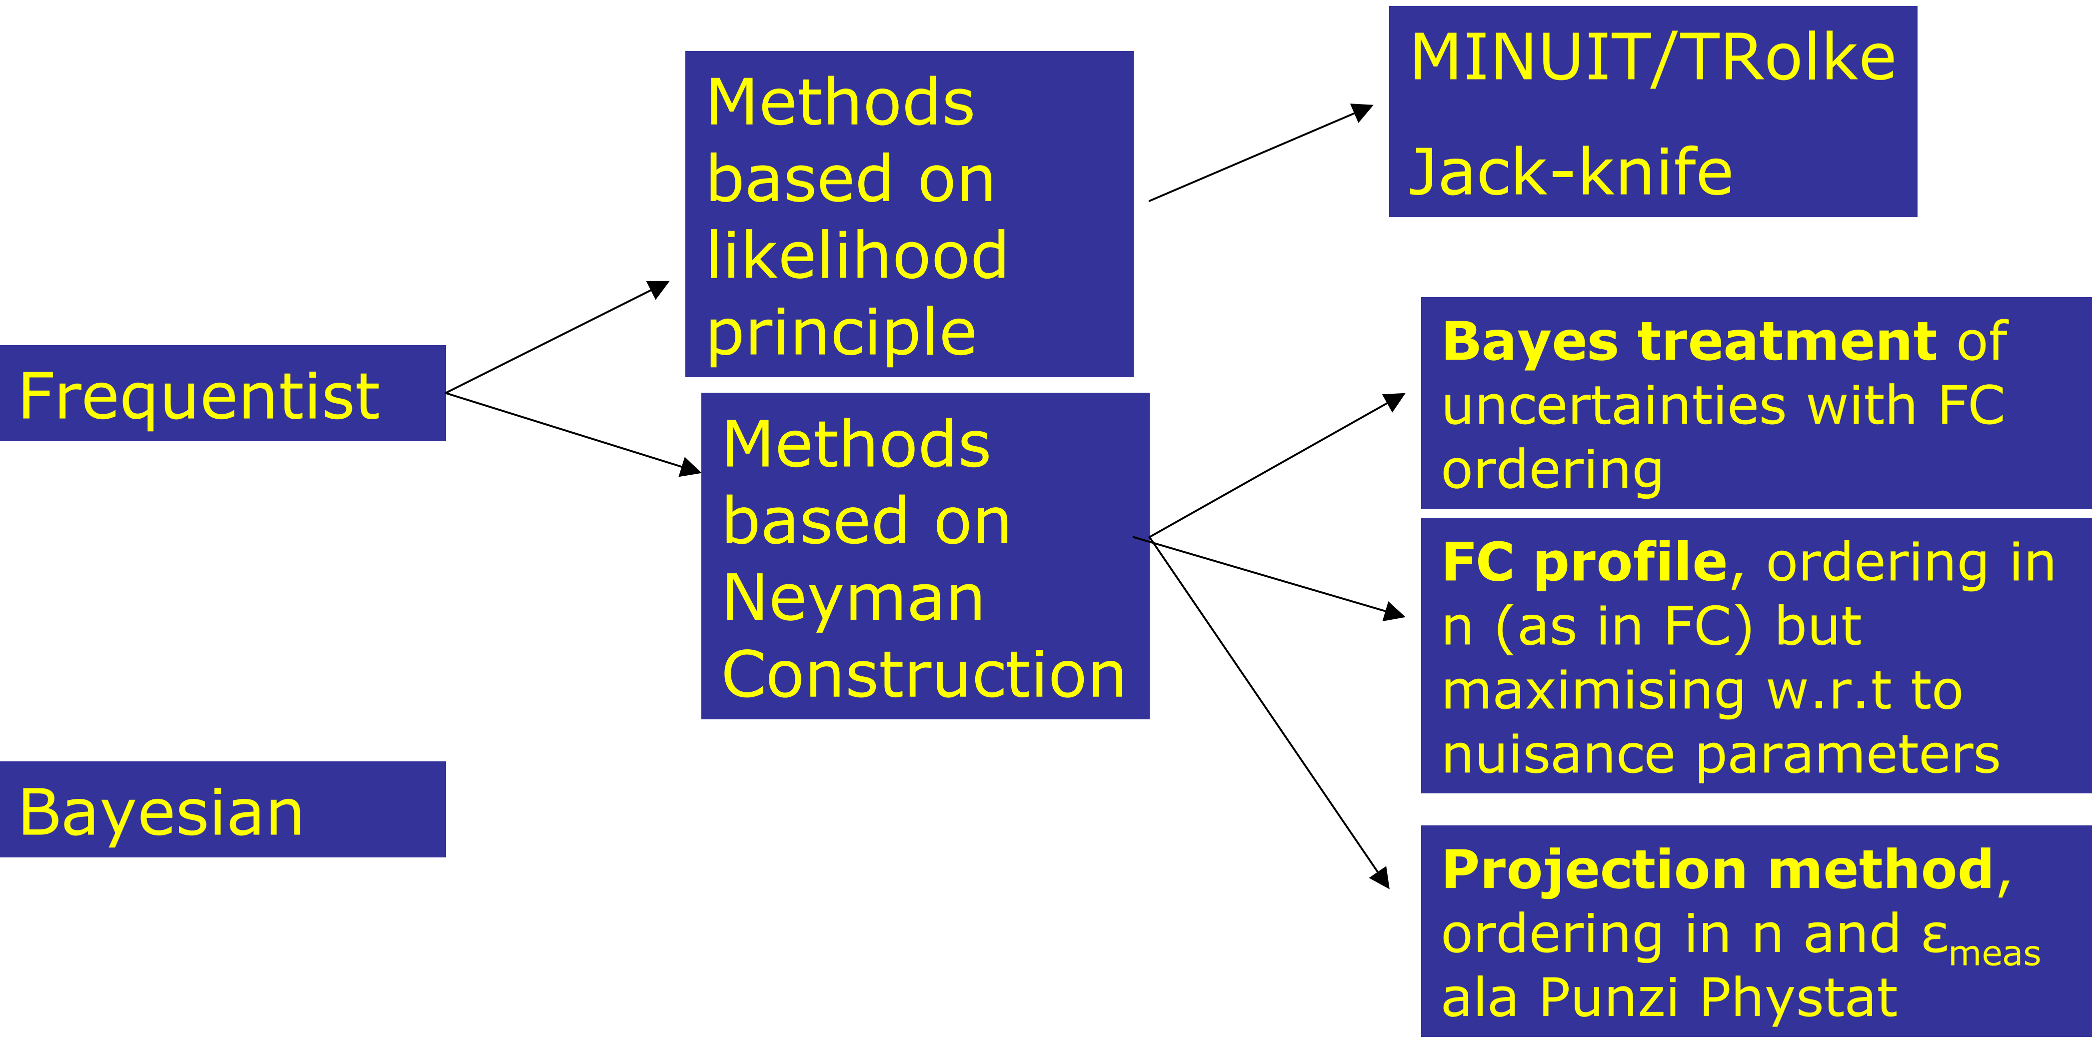

Some more details on the Punzi ordering function (slide added after the talk)

The Neyman sum condition in the discrete case is:

$\Sigma P(n|s) = 1 - a$

For the continous case this writes:

$\int P(n|s)\, dn = 1-a$

The Neyman condition for upper limit:

$$P(n < n_1(s)|s) = \alpha;\ n_2 \rightarrow \infty$$

Which can also be written as:

$\int_{r0(n) < c} P(n|s)\, dn = 1-a$, where $r0(n) = n$

is an “ordering function”, i.e. it represents a rank, where the highest rank (in this case largest n) is included first. This has of course to be done for each s, but in this case is independent of s.

Some more details on the Punzi ordering function (slide added after the talk)

In presence of uncertainty in efficiency, the Neyman condition has to be fulfilled in both n and $\varepsilon_{meas,}$ i.e.:

$\int\int_{r0(n,\varepsilon(meas)) < c} P\,(n,\varepsilon_{meas}|\varepsilon,s)\, dn\, d\varepsilon_{meas} = 1\text{-}a$

where $r0(n,\varepsilon_{neas})$ is the ordering function in now the two observed quantities (this has to be done for each s and ε_{true}).

Now, Punzi chooses an ordering function which is independent of ε_{true} :

$$r(n, \epsilon_{meas}; s) = \int_{r_0(n')<r_0(n)} p(n'|\epsilon_{meas}; s, \epsilon_{max}(\epsilon_{meas})dn'$$

…and the ordering function is itself defined via an integral which is using the r0(n) ordering function (which for example could be the Feldman-Cousins rank or again r0(n) = n).

Projection method with appropriate ordering II

Ordering function:
(Punzi, PhyStat05)

$$r(n, \epsilon_{meas}; s) = \int_{r_0(n') < r_0(n)} p(n' | \epsilon_{meas}; s, \epsilon_{max}(\epsilon_{meas}) dn'$$

Poisson signal, Gauss eff. Unc (10 %)

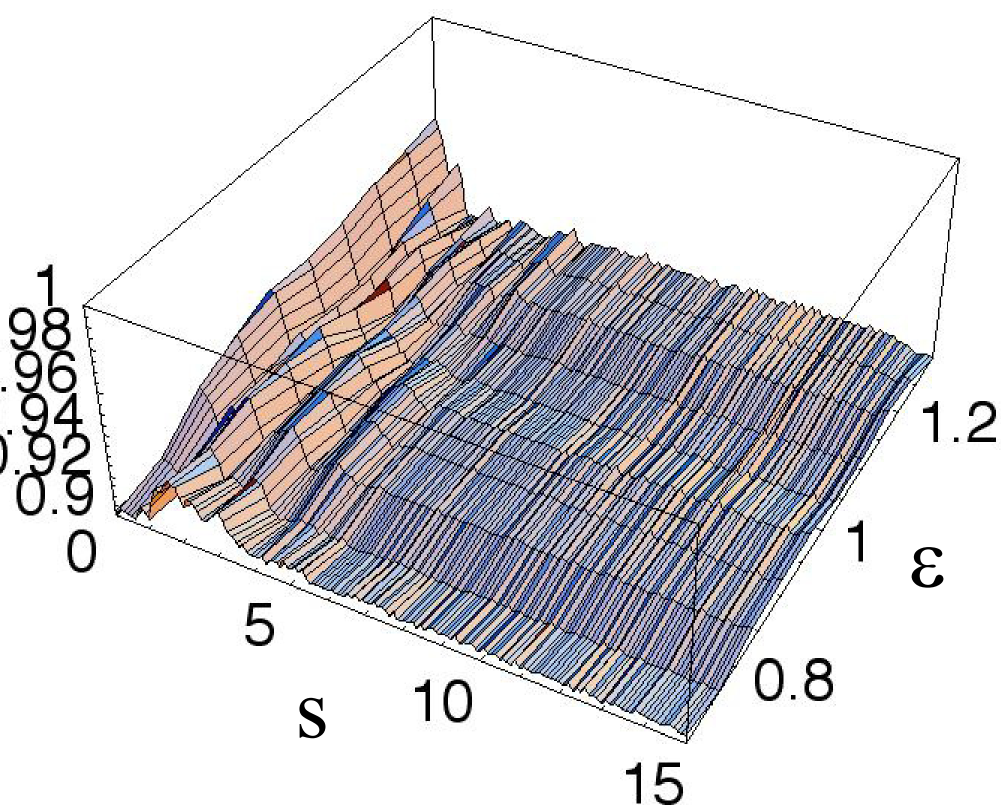

Can be any ordering in prime observable sub-space, in this case

Likelihood ratio (Feldman & Cousins)

~ FC Profile

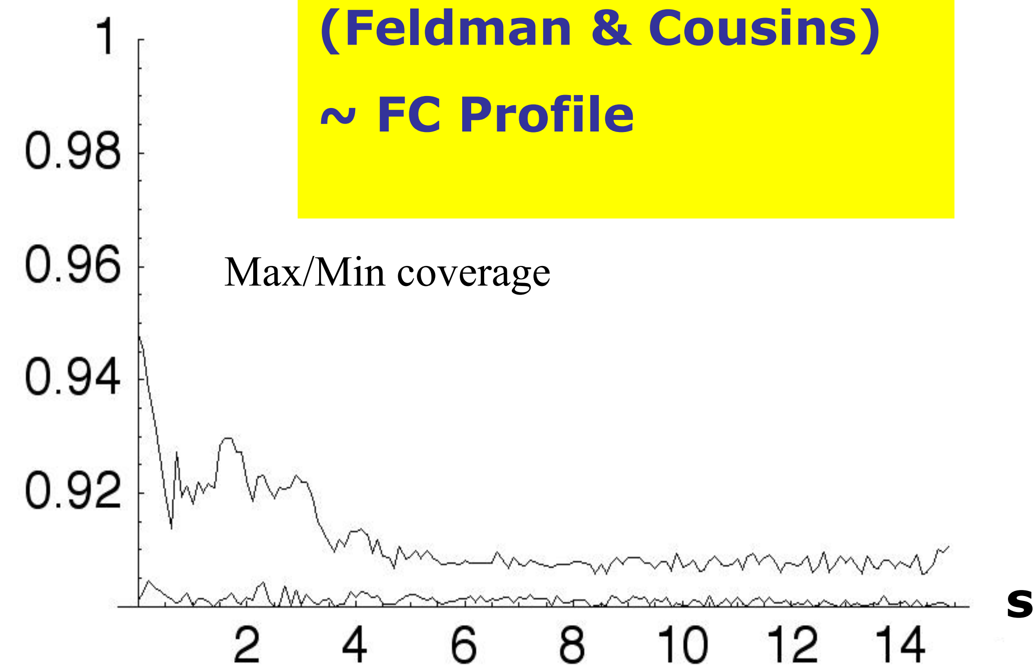

Summary : Bayesian Intervals

- Still some large experiments recommend to use Bayesian methods.
- In my personal view, there is no reason to use “objective” Bayesian methodology
- … but if you insist, spend a couple of thoughts on priors ….

Summary: Frequentist Confidence Intervals

- The (MINUIT, Rolke et. al.) jack-knife works surprisingly well for common problem with nuisance parameters and there are preliminary indications for it to have reasonable properties for extreme confidence levels (e.g. Higgs search)

- <u>The full Neyman construction:</u>
  - The most commonly used ordering scheme is: FC
  - In presence of nuisance parameters:
    - Bayesian gives some over-coverage, but is very often used
    - Projection method: gives good coverage if applied with clever ordering function, computationally quite cumbersome → not so often used
    - Profile likelihood ratio: gives good coverage and is easier to compute than projection method

      → not so often used due to lack of documentation

So, what would I do ?

NB: for combined experiments (channels): jack knife and Bayesian methods fine, constructions cumbersome

1) My preference would be to use the Neyman Construction with profile likelihood ordering (FC Profile)

2) An alternative would be the jack-knife:

- In particular I would use TRolke for problems where I can use it since here coverage has already be shown
- ….. otherwise MINUIT .....

3) If I use the jack-knife I think a check of the coverage is indispensable (**unless TRolke-problem**).
If I apply FC Profile, I would also check, if possible, but I would trust it more ….

4) If you for some reason prefer the Bayesian treatment of nuisance parameters (though giving some over-coverage), sure, use it and increase my (and twice Bob Cousin's) citation count ….. again a coverage check is recommended (unless using pole++)

5) A referee would most probably (and unfortunately) accept everything which has a refereed publication

- … even completely Bayesian methods with bad priors …

Back-Ups

What if $\sigma \rightarrow 0$?

- Bayesian treatment of nuisance parameters: approaches zero uncertainties naturally
- Frequentist methods:

Full construction (projection/profile)

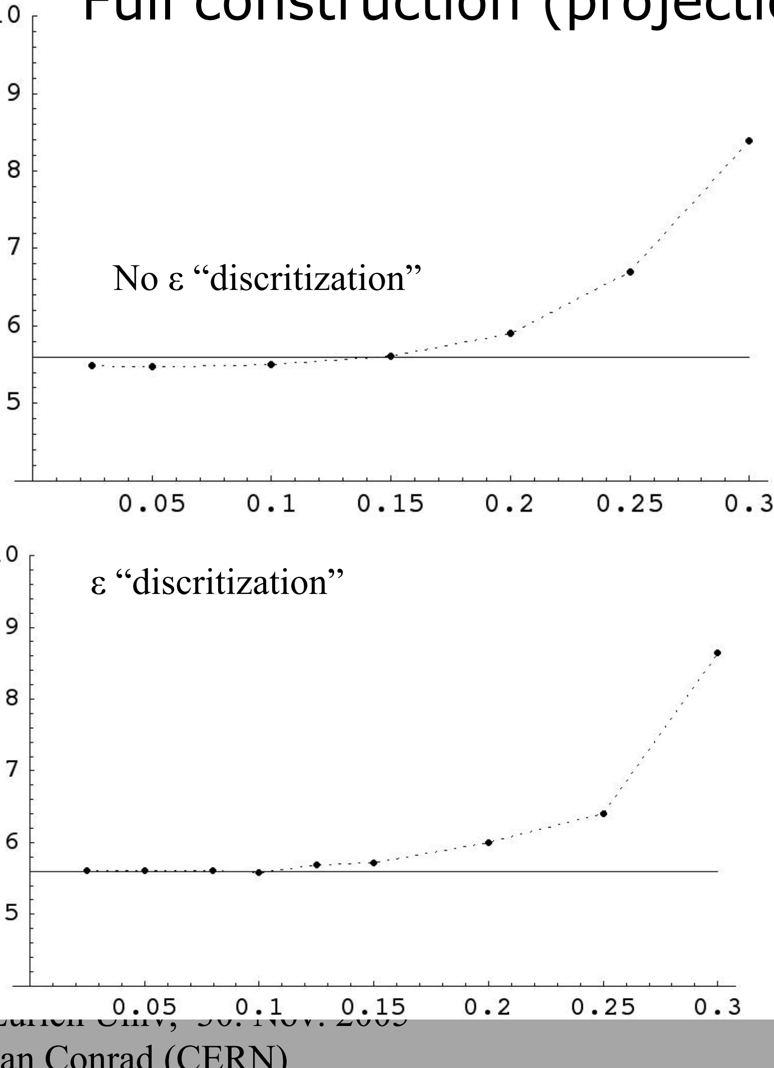

Jack-Knive

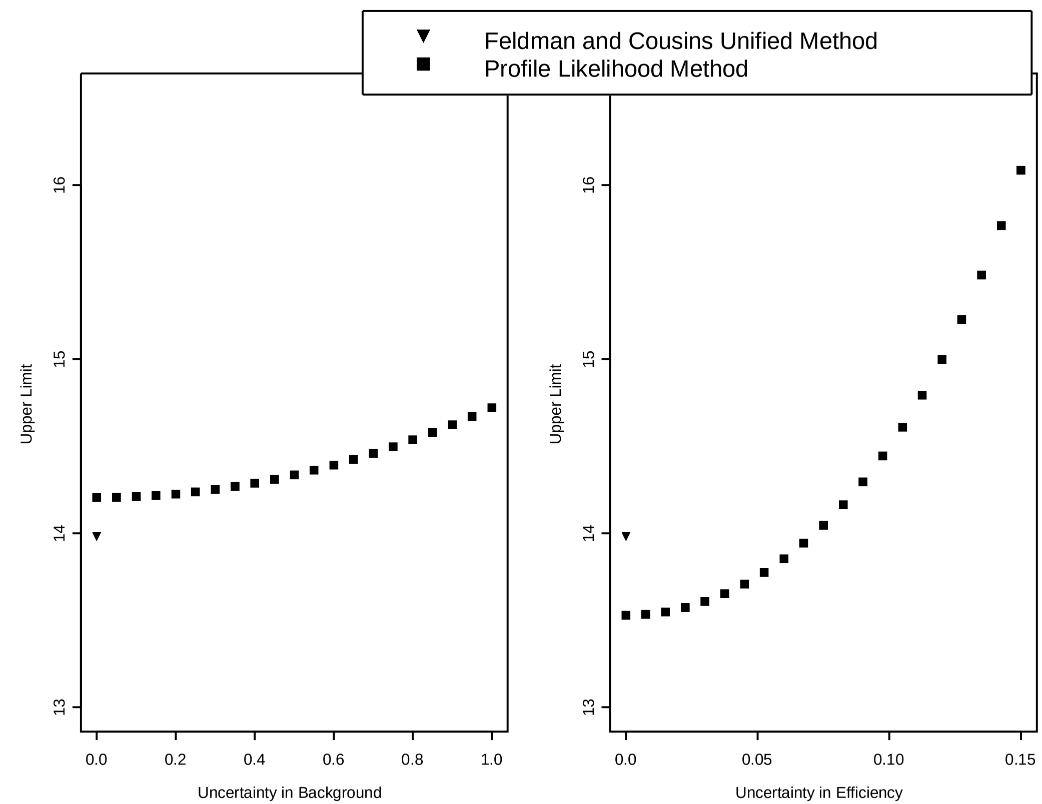

LIGO : gravitational wave bursts

- Signal: require coincidence in all three LIGO detectors
  - 2 in Hanford – 1 in Livingston (3000 km away)
- Background: noise level fluctuations, environmental and instrumental disturbances

Coincident events	6
Est. background	10.1 ± 2.3
90 % limit:	0., 2.3

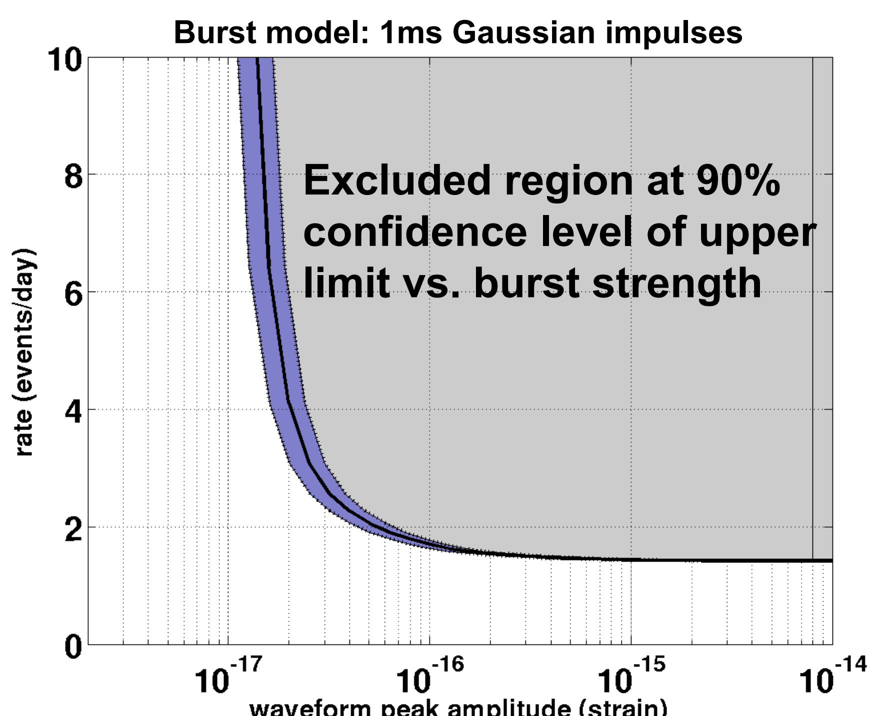

Obs:

Uncertainty in the efficiency !

Treated in Bayesian way, I will tell you later

Example: LIGO search for gravitational waves

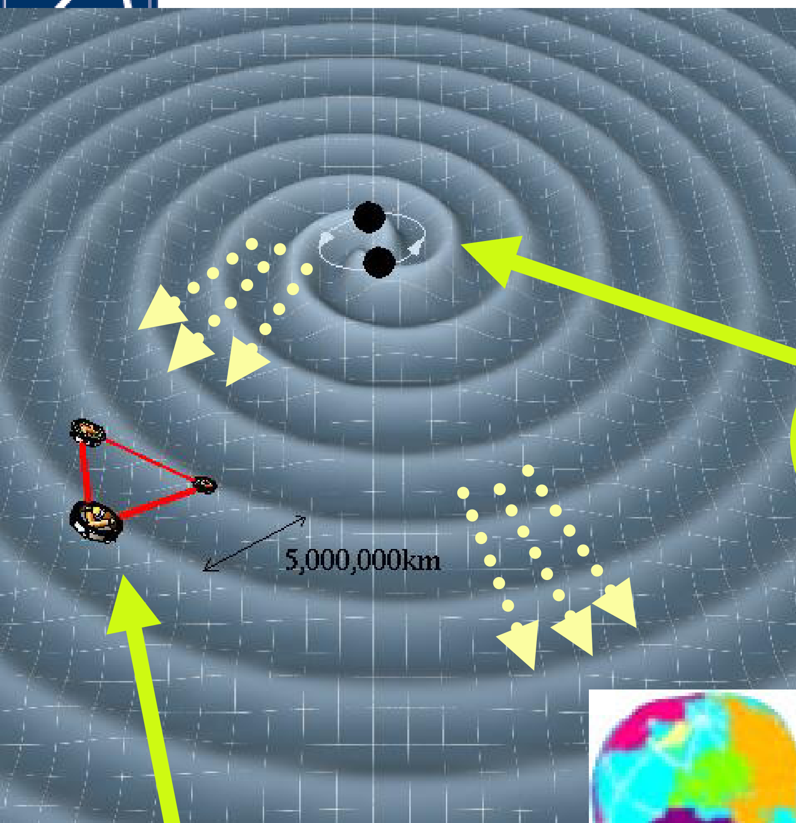

Detectors in space
LISA

Gravitational Wave Astrophysical Source

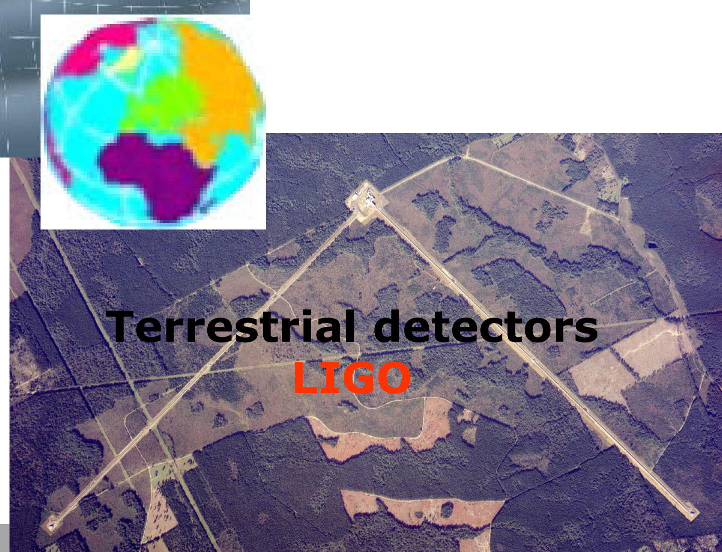

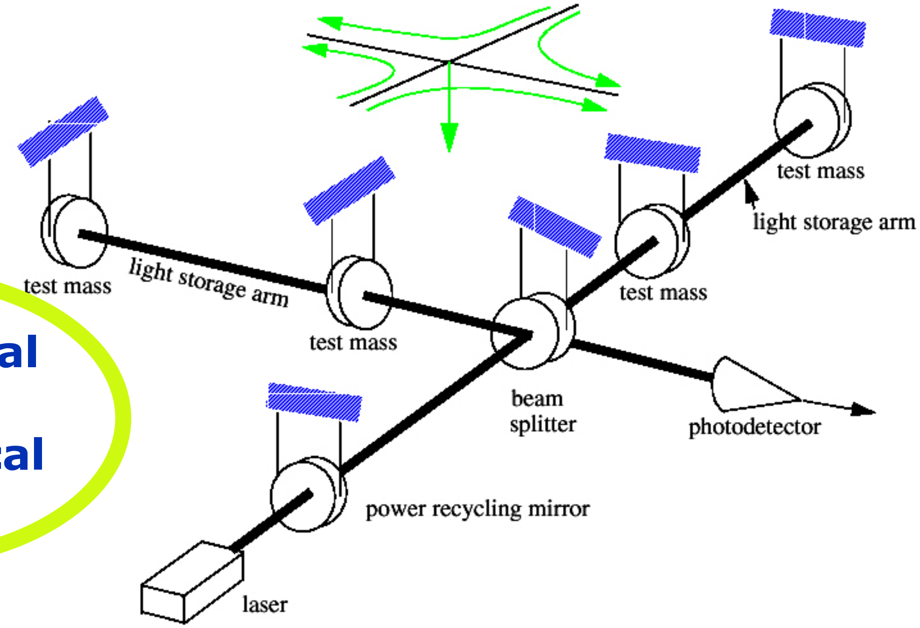

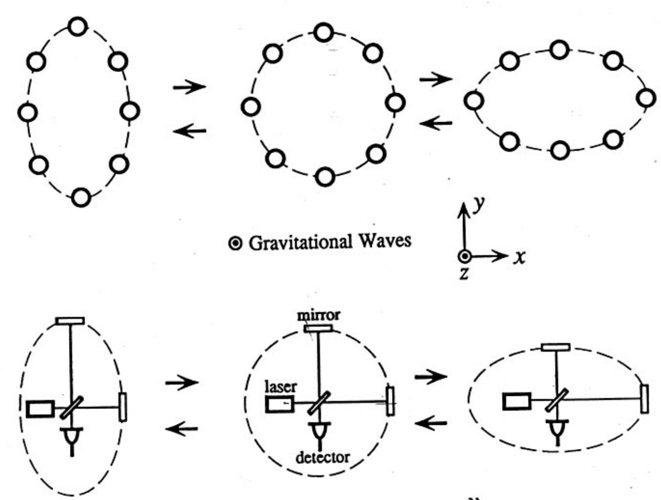

G. A Sanders

Upper limits, with appropriate 2-D ordering

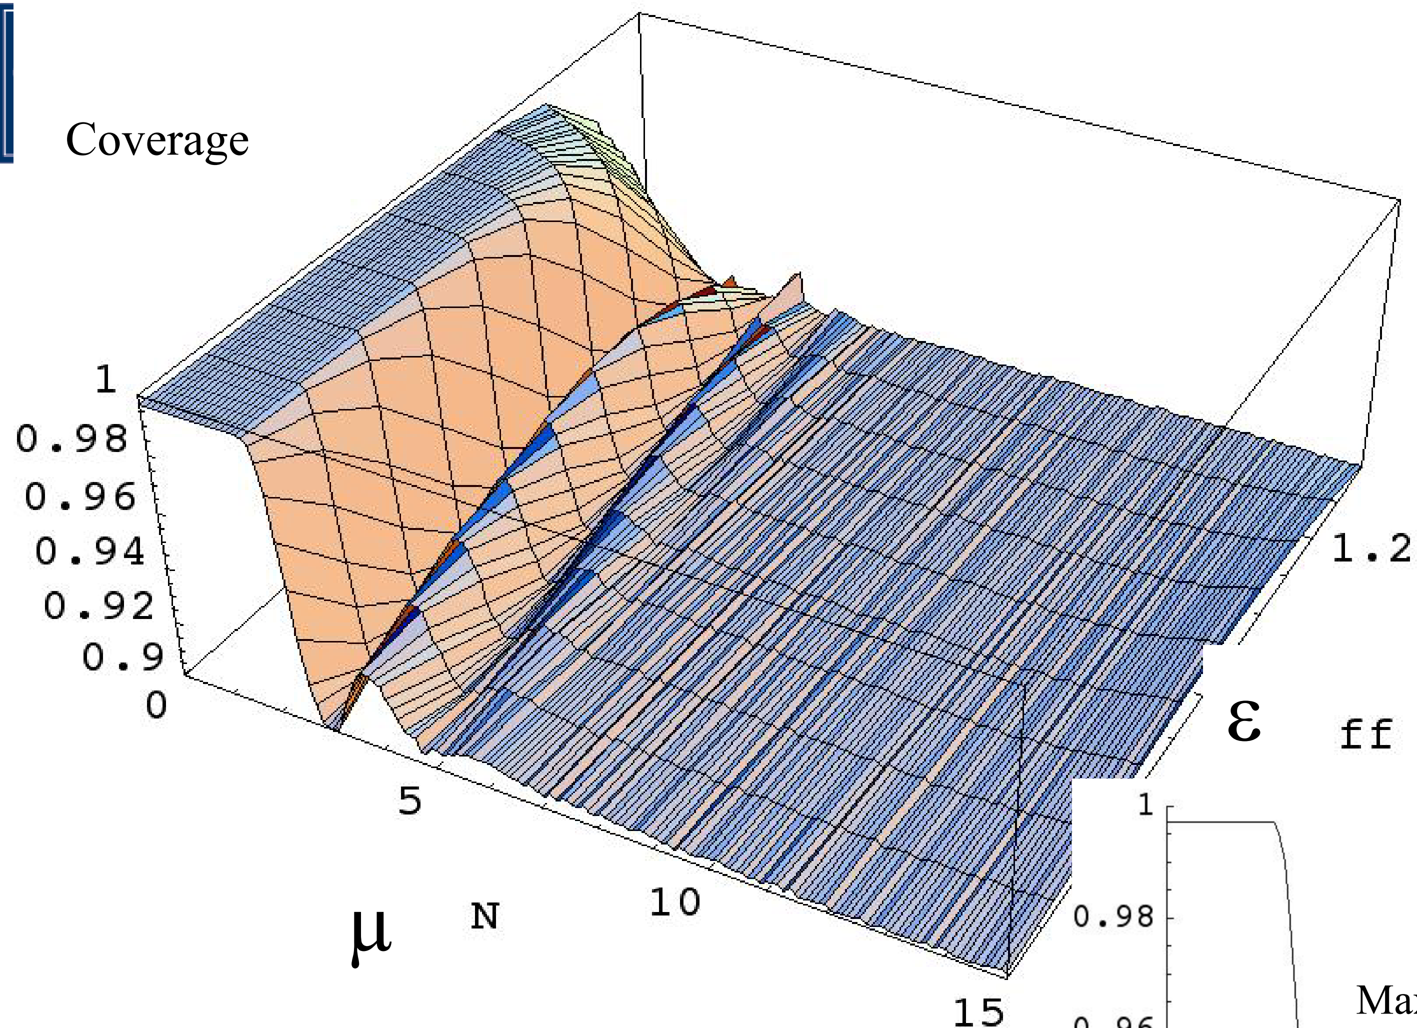

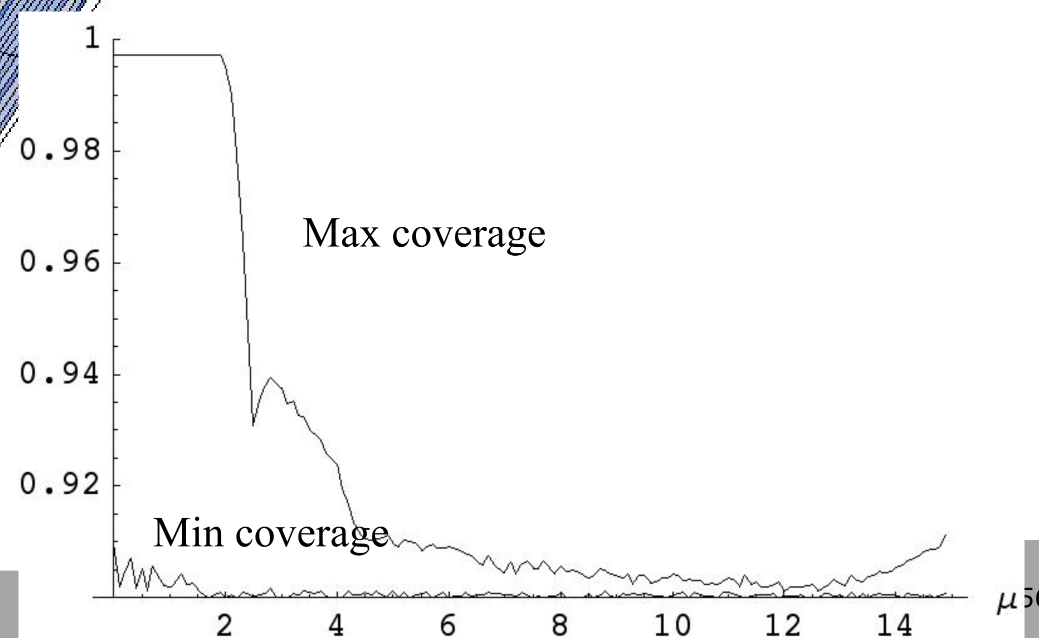

Pretty good ! No overcoverage beyond "discretization ripples"

Upper limits, naïve version (x ordering)

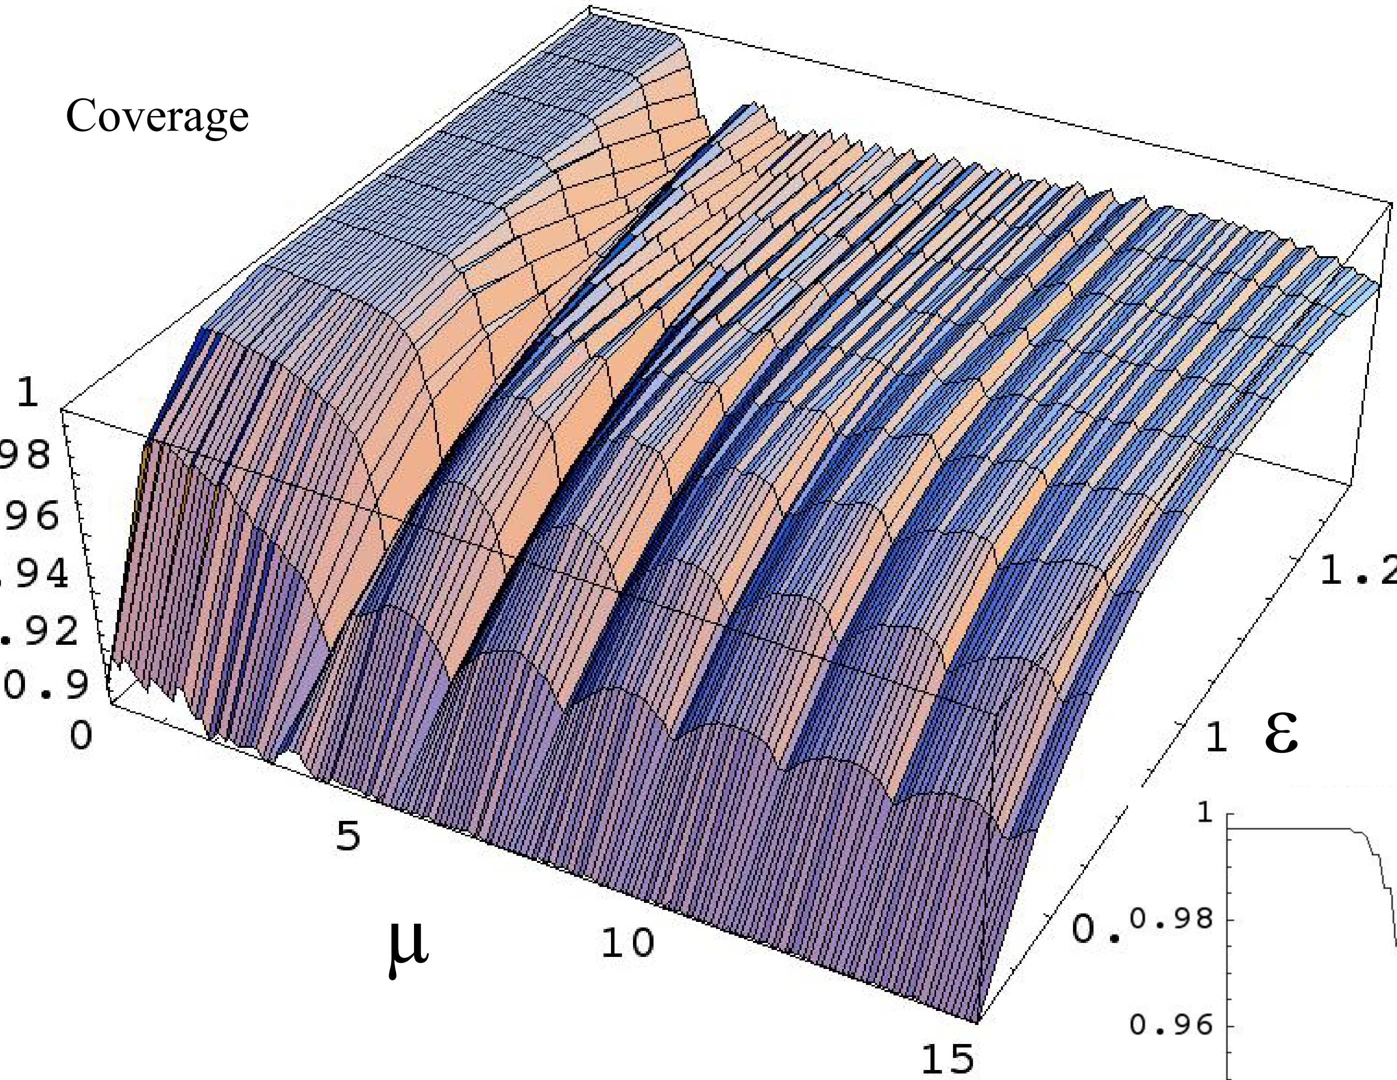

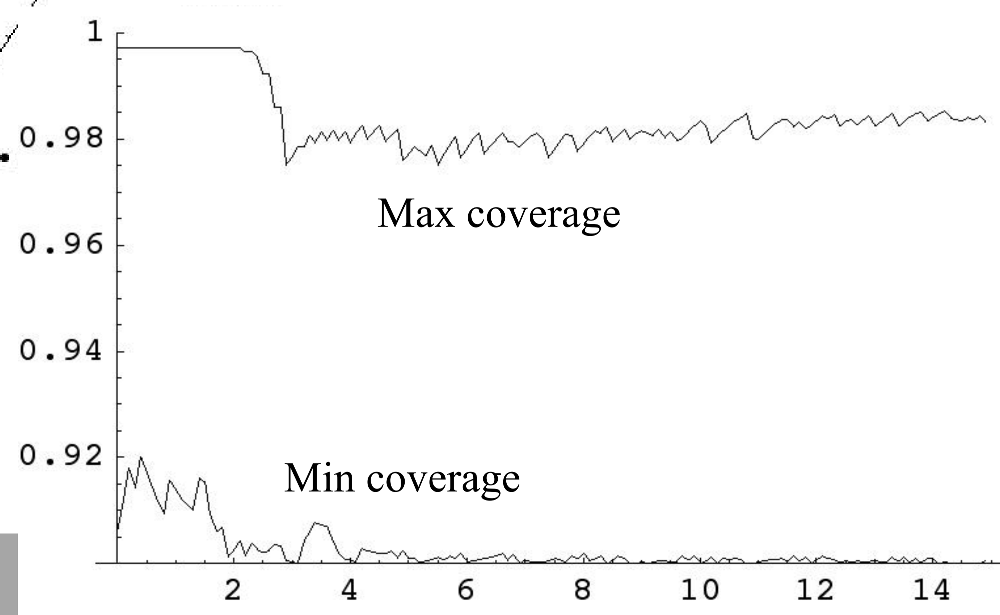

Not very exciting…

TLimit (The CL_s method)

- Class description: "Class to calculate 95 % confidence limits", violent understatement:
  - "Class for Hypothesis Testing"
  - → Caluclates p-values, power etc …..

- Can also be used to calculate upper limits at ANY given confidence → needs a loop over hypotheses

- Test statistics: likelihood ratio based on Neyman Pearson Lemma

$$-2\ln Q = -2\ln\frac{\mathcal{L}(n|s+b)}{\mathcal{L}(n|b)}$$

- Confidence Level calculated from histograms of -2 ln Q

$$CL_s = \frac{P(Q \le Q_\circ|s+b)}{P(Q \le Q_\circ|b)} = \frac{CL_{s+b}}{CL_b} = \frac{\text{p-value(s+b)}}{1-(\text{p-value(b)})}$$

- Uncertainties on signal & background can be included (Gauss), using Bayesian PDF integration (ala Highland & Cousins)

T. Junk: Nucl.Instr.Meth.A434:435-443,1999
A. Read: e.g. J.Phys.G28:2693-2704,2002

Better name:

TLevel (TCLs)

Obtained from histograms

Note: Q monotonically increasing for more signal like obs

enormalizes on
e background
nly case.

CL_s: example

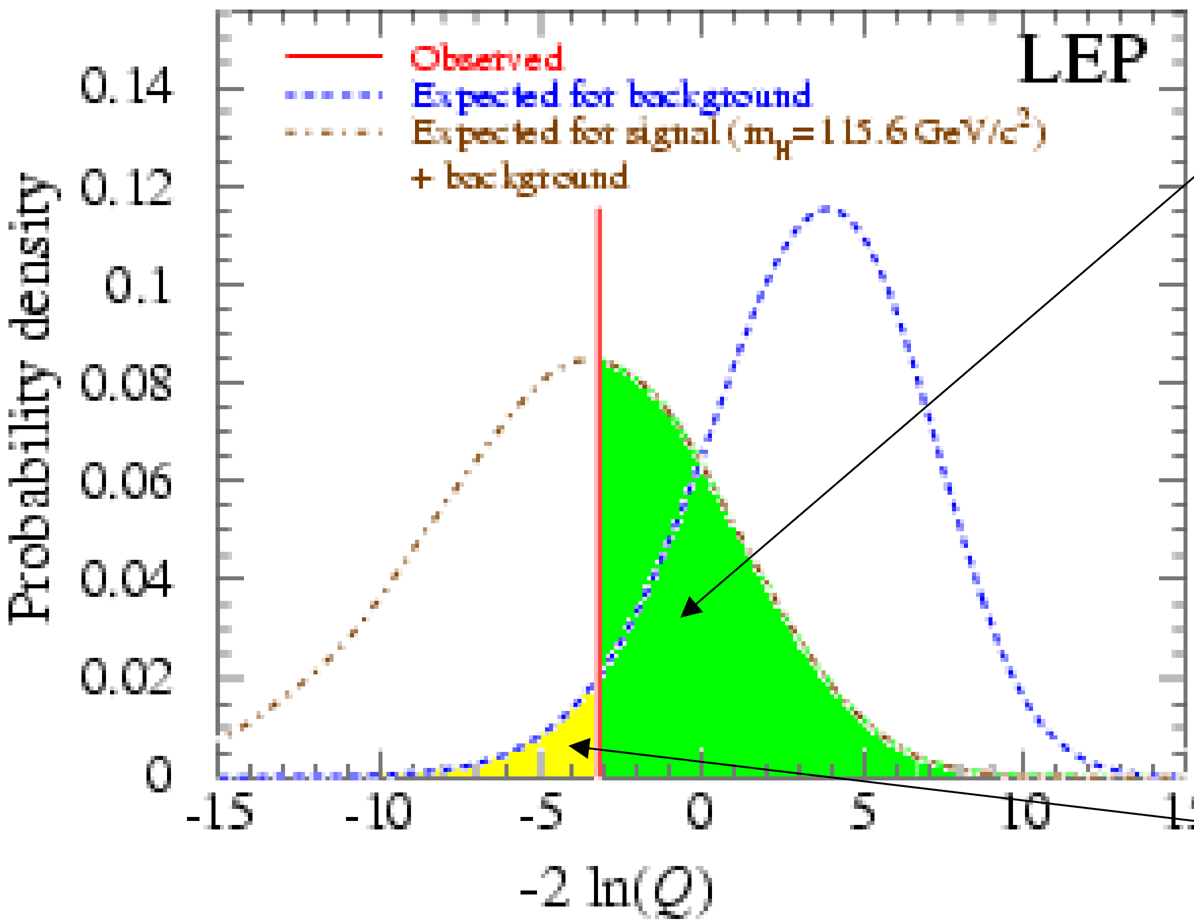

CL

CL_s : generic method

TLimit:

designed for multi-channel measurements with sufficient statistics

(you don't provide the PDF but the observed signal)

1- CL_b

A. Read: e.g. J.Phys.G28:2693-2704,2002

TLimit (CL_s): coverage

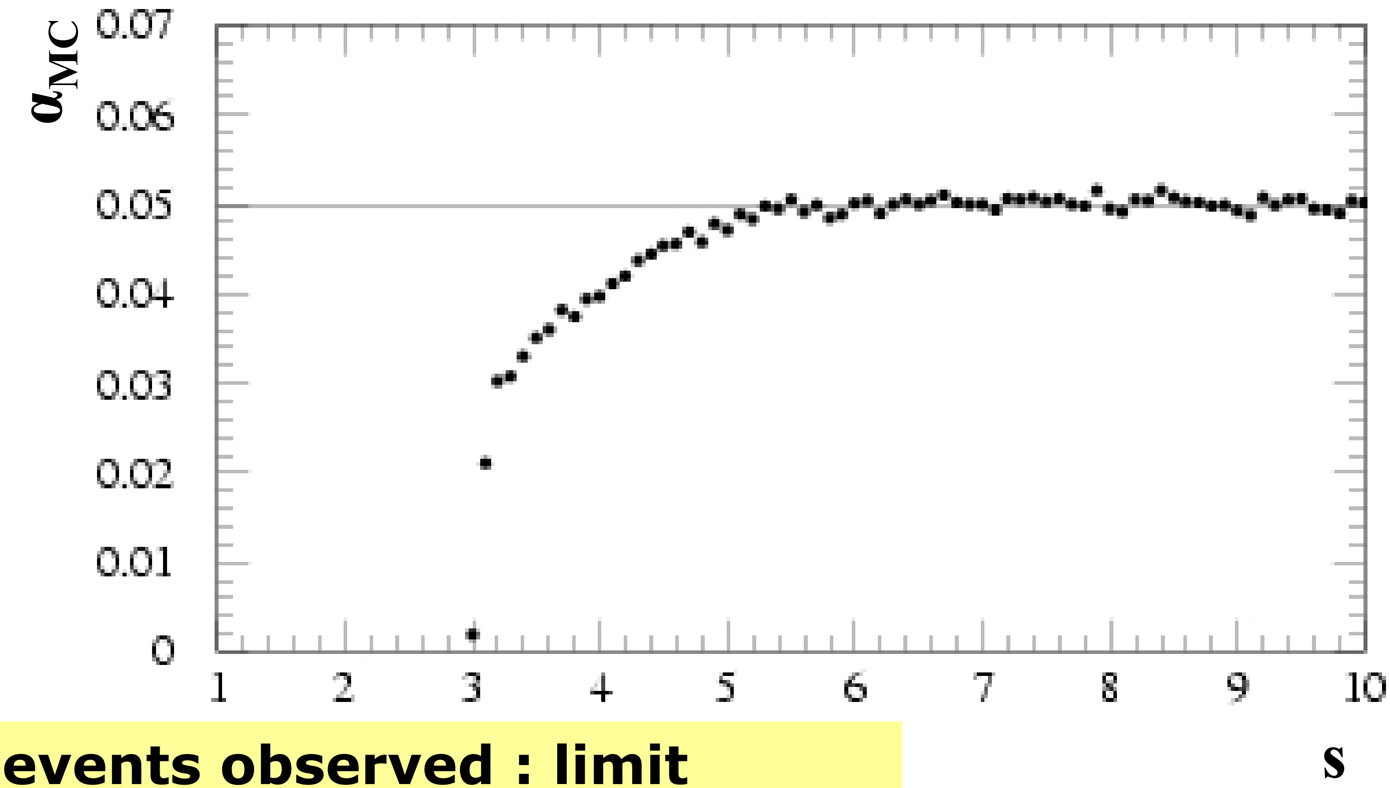

Higgs search

(distributions in rec. higgs mass)

with b =4 and a variable mass resolution

(therefore only little structure)

or zero events observed : limit ndependent of background for CL_s not so for CL_{s+b})

T. Junk: Nucl.Instr.Meth.A434:435-443,1999

Some Intervals

N_obs	b	Sys. Unc.	CHC interval
2			[illegible]3.90]
			[illegible]3.95]
			[illegible]4.10]
		0.4	[0,4.65]
6	2	0.0	[1.1,9.45]
		0.2	[1.05,10.05]
		0.3	[1.05,11.50]
		0.4	[1.05,13.35]

Rule of thumb:

$\Delta w/w \sim \sigma^2$