Limits in High Energy Physics

Stefan Schmitt, DESY

Tutorial/lecture for the
Terascale Statistics School

Hamburg, 20 March, 2013

Outline

- Part I: basic concepts, Bayes and Frequentist, simple example
- Part II: Poisson with background, expected limit, CL_s method
- Part III: systematic uncertainties and many channels, hybrid method, profile likelihood
- Exercises: Lecture is interleaved by exercises ~10-15 minutes each. Solutions are discussed in the lecture
- ROOT macros for exercises:
  `www.desy.de/~sschmitt/LimitStatSchool2013/macros`
- If available on our computer, use wget:

```
wget -N -nd www.desy.de/~sschmitt/LimitStatSchool2013/macros.list
wget -N -nd -i macros.list
```

Probability theory: selected items

- Elements of Ω: events, outcomes of an experiment
- Probability of a set A: $0 \leq P(A) \leq 1, P(\Omega) = 1$ $\quad P(\Omega \setminus A) = 1 - P(A)$

  $P(\Omega) = 1, P(\emptyset) = 0$
- Example: Poisson distribution $P(\{N\}) = \frac{e^{-\mu}\mu^N}{N!}, \Omega = \{0,1,2,...\}, A = \{N\}$
- Conditional probability of A given B: $P(A|B) = \frac{P(A \cap B)}{P(B)}$
- Bayes' law:

$$P(B|A) = \frac{P(A|B)P(B)}{P(A)}$$

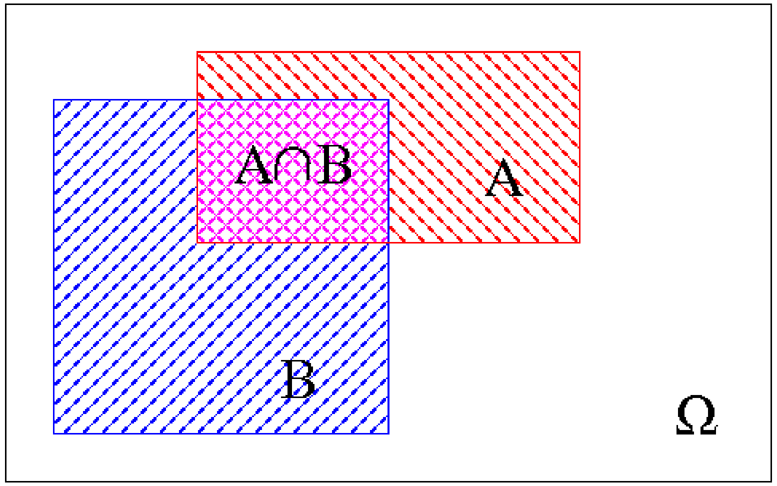

Probability densities

- Probabilites on discrete sets: each element has a finite probability

  Example: Poisson distribution

  → For event counts

$$P(\{N\}) = \frac{e^{-\mu}\mu^{N}}{N!}$$

$$\Omega = \{0, 1, 2, \ldots\}$$

- Probability densities: probabilities are defined by integrals

  Example: normal distribution

  → For systematic errors

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-(x-\mu)^2}{2\sigma^2}}$$

$$\Omega = \mathbb{R}$$

$$P(a \leq x \leq b) = \int_a^b f(x)\,dx$$

Parameters

Parameters of a probability density/distribution ⟷ The outcome of the experiments/possible observations

Examples:

- Poisson distribution:

  $$P(\{N\}) = \frac{e^{-\mu}\mu^{N}}{N!}$$

  μ is a parameter

- Normal distribution:

  $$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-(x-\mu)^2}{2\sigma^2}}$$

  μ and σ are parameters

- Symbol for (a set of) parameters: θ

Types of parameters

- During limit setting, parameters may be fixed or variable
- Types of variable parameters:
  - Parameter of interest

    → Limits are set on this parameter (e.g. Higgs coupling)
  - Nuisances

    → These are “not of interest” (e.g. background normalisation)
- Special case of “fixed” parameters:
  - Parameter scan

    → limit calculation is repeated many times (e.g. Higgs mass)

Example: limit with parameter scan

- Example: search for Rp violating SUSY at HERA (resonant single squark production)
- Limit is set on the Rp-violating coupling λ
- squark mass scanned (y-axis)
- Other SUSY parameters are also scanned (yellow area)

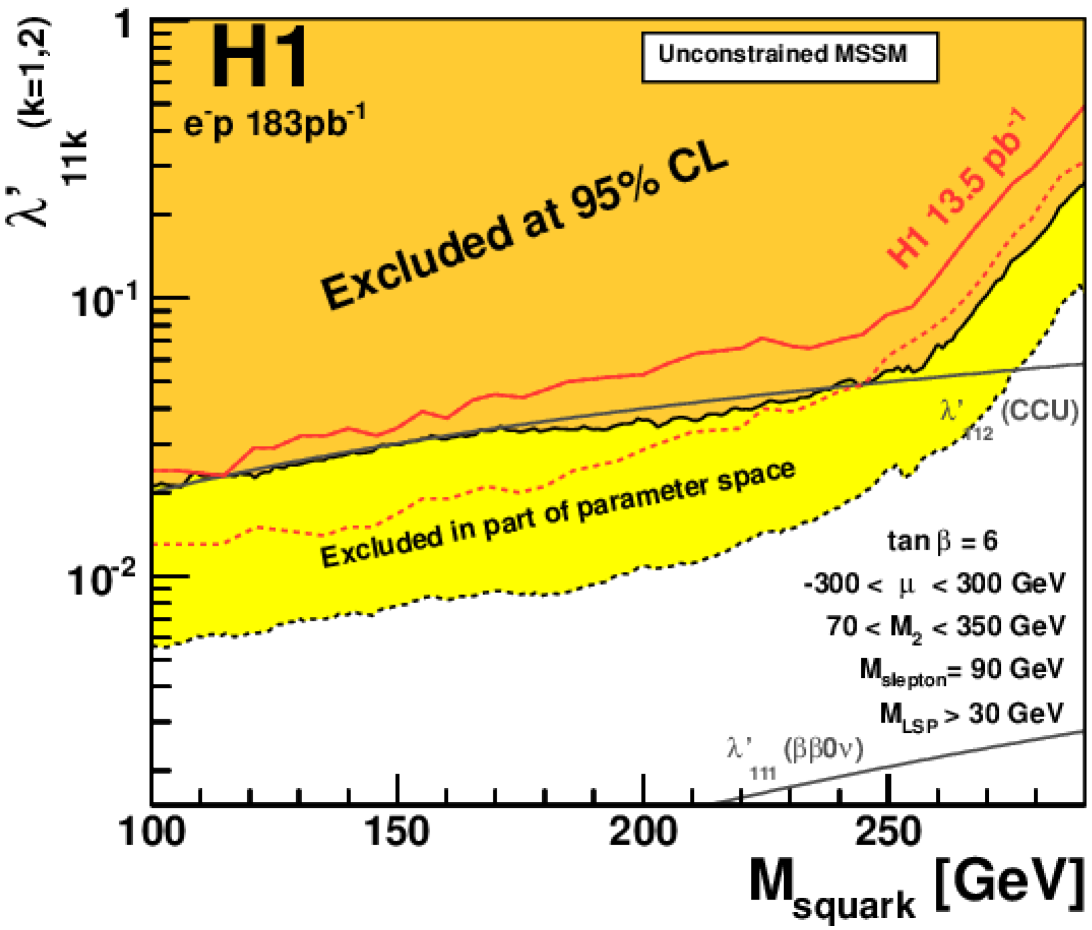

Example plot: Search for Squarks in R-parity Violating Supersymmetry in ep Collisions at HERA, Eur.Phys.J.C71 (2011) 1572

Frequentist/Bayesian probability

- Frequentist view: probabilities describe the outcomes of experiments

  Models have unknown parameters. Probabilities (to make the given observation) are quoted as a function of the parameters

- Bayesian extension: probabilities are also used to describe the “degree of belief” in parameters.

  → The parameters themselves have probabilities assigned.

Bayesian definitions

$$P(B|A) = \frac{P(A|B)\,P(B)}{P(A)}$$

- Prior: P(B) where B is the theory
- Likelihood: P(A|B) where A is the measurement
- Posterior: P(B|A) is the result of the analysis

- P(A) has no special name. The normalisation is often calculated using the relation P(B|A)+P(~B|A)=1

Exercise 1 (Bayes' law)

$$P(B|A) = \frac{P(A|B)\,P(B)}{P(A)}$$

- Disease and a test for the disease
- 0.1% of the population have the disease (prior)
- If one has the disease, the test is positive with 99% probability (likelihood)
- If one does not have the disease, the test is positive with 1% probability
- What is the posterior probability to have the disease, given a positive test?

Exercise 1 (Bayes' law)

- Disease and a test for the disease
- 0.1% of the population have the disease (prior)
- If one has the disease, the test is positive with 99% probability (likelihood)
- If one does not have the disease, the test is positive with 1% probability
- What is the posterior probability to have the disease, given a positive test?

$$P(B|A)=\frac{P(A|B)\,P(B)}{P(A)}$$

A: positive test
B: has disease

Prior: P(B)=0.001
Likelihood: P(A|B)=0.99
P(A|~B) =0.01

Normalisation:
P(A)=P(A∩B)+P(A∩~B)=
P(A|B)*P(B)+P(A|~B)*P(~B)=
0.99*0.001+0.01*0.999= 0.01098

Posterior: P(B|A)=
0.99*0.001/0.01098 = 9%

Because the disease is so rare, the probability is only 9%.
The test has to be improved, 1% of false-positive tests is too much

Probabilities in high energy physics

- Probability: predict number of events given the theory (parameter of interest) and the experimental setup (nuisances)
- Question: what does a specific observation tell about the theory
- Frequentist: give for each theory the probabillity of the observation (there is no probability for a theory)
- Bayes: assign probability (degree of belief) to theories
- High energy physics: make use of both views (preference for frequentist, in particular for discoveries)

Confidence intervals, Limits

- Confidence intervals tell about parameters of the theory
- Confidence level (CL): associated probability
  - Different meaning of CL Frequentist/Bayesian:

    Frequentist: CL=P(observation)   Bayesian: CL=P(theory)

- Double-sided interval: measurement (usually CL=68%)
- Single-sided: limit (often CL=95%)

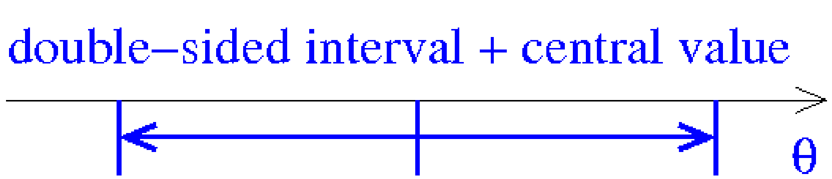

$\theta = \mu \pm \sigma \, [\text{at } 68\% \text{ CL}]$

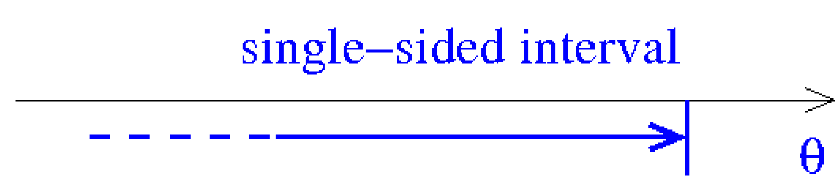

$\theta \le \theta_{max}$ at 95% CL

Frequentist limit: Neyman construction

- For each value of the parameter θ, find single-sided interval with probability≥CL (CL is fixed, e.g. CL=0.95)
- Interconnect interval edges
- For a given observation find the largest θ, where x_{obs} is just contained in the interval → limit on θ

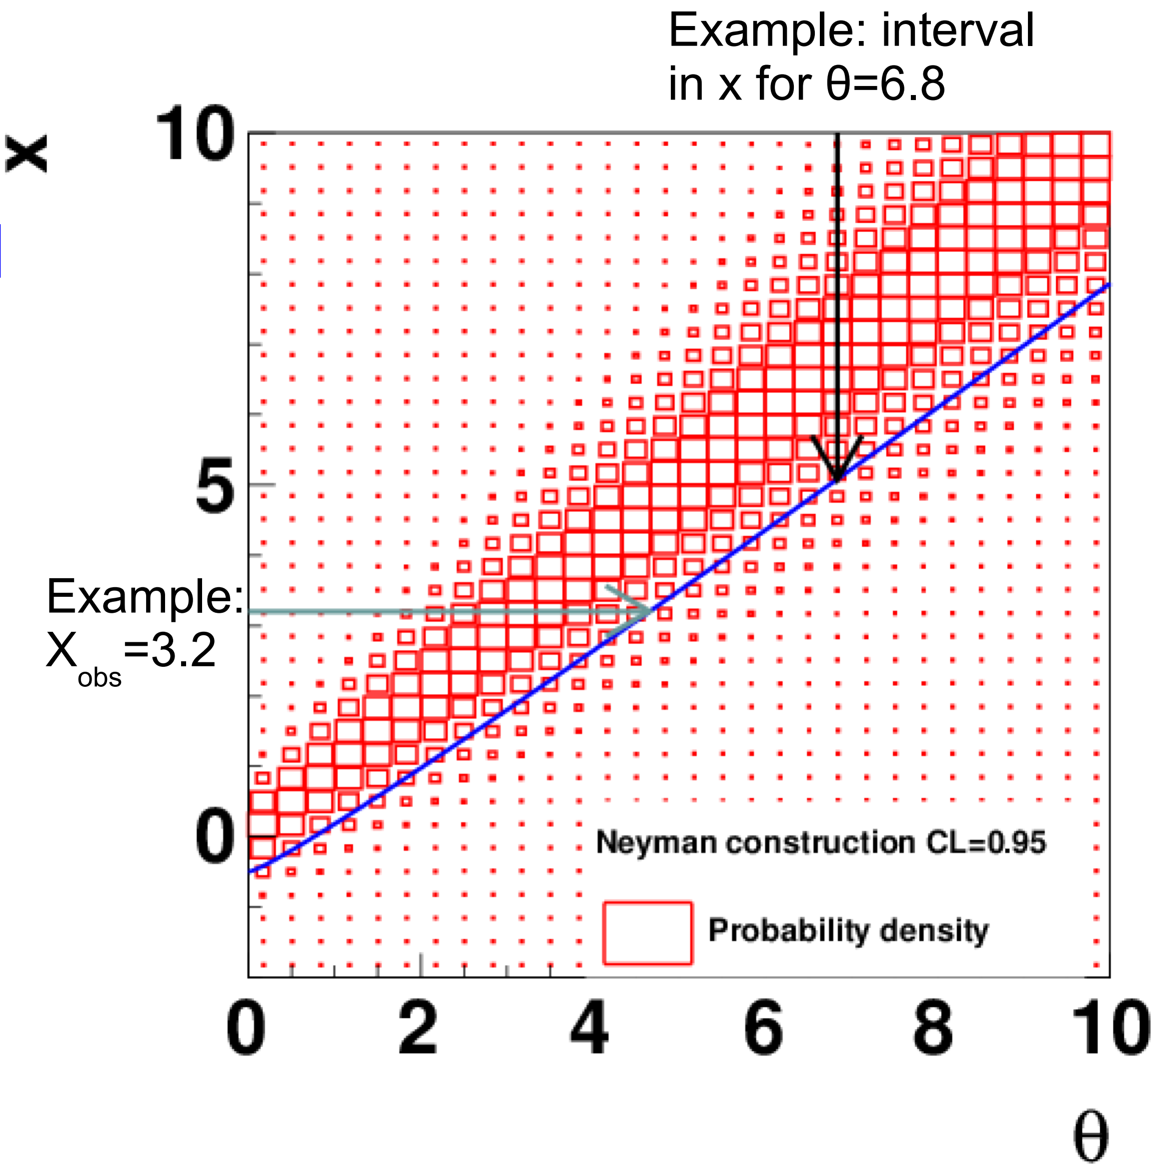

Frequentist upper limit, Poisson data

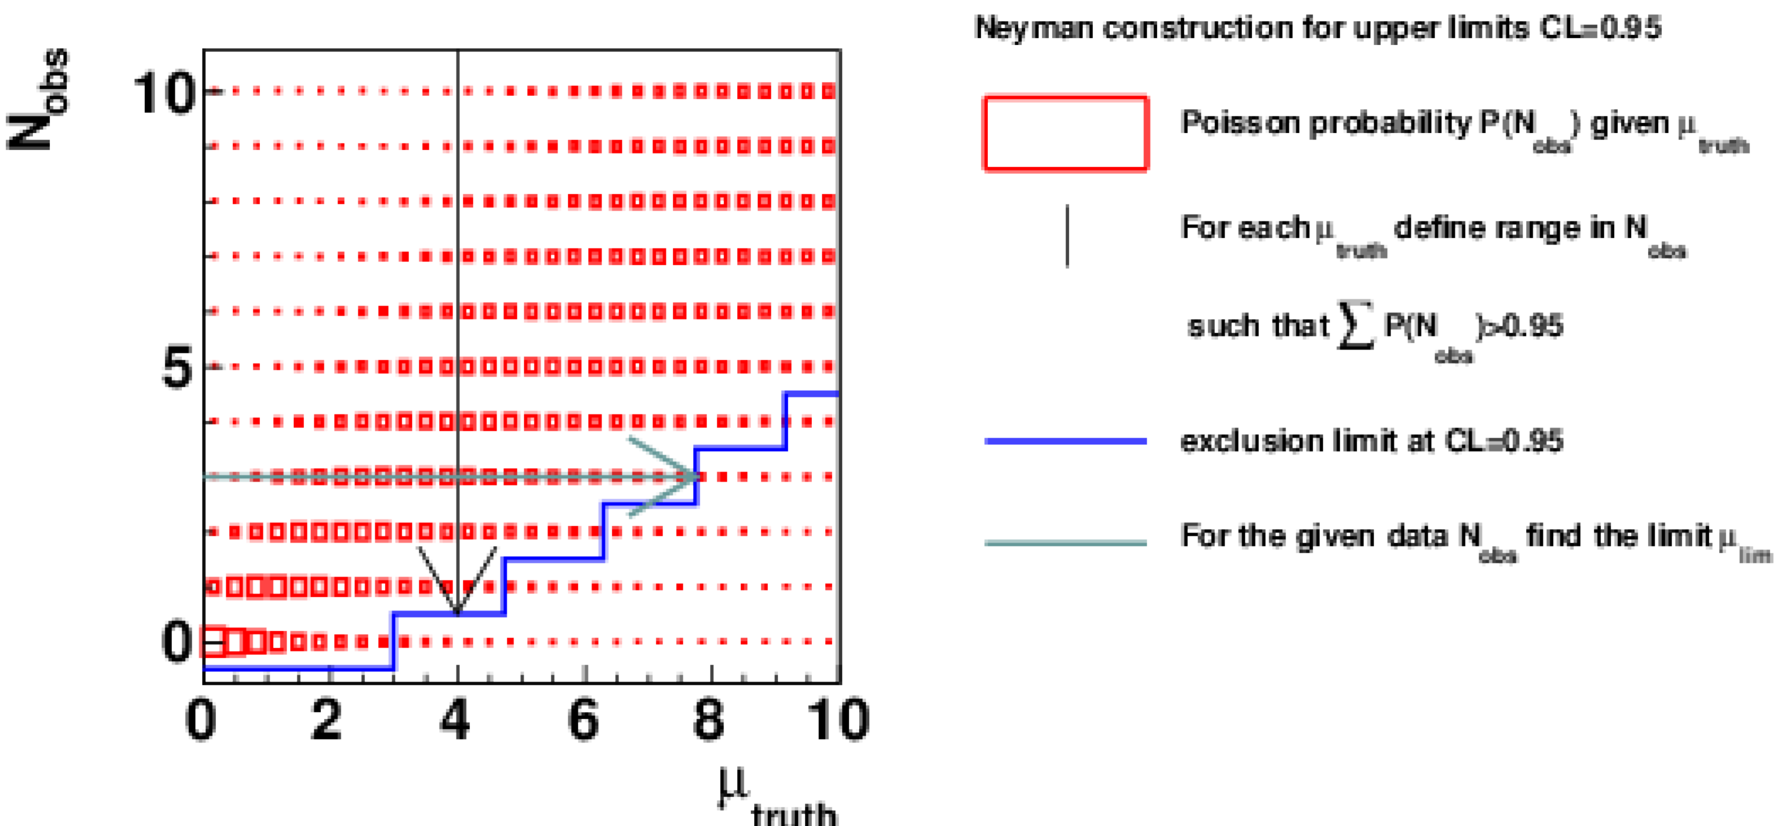

- Neyman construction, for each μ find N_{obs}-interval with P≥CL
- Then: read off μ_{limit} for a given N_{obs}
- Note: discrete N_{obs} but continuous μ → steps in the limit

Exercise 2 (Neyman construction)

- Poisson experiment, determine limits on the parameter μ, given N_{obs}

  a) determine the range $N_{obs} \leq N \leq \infty$ for CL=0.95 and μ=2,3,5,10. What is the probability to find the measurement in these ranges

  b) determine the limit on μ for N_{obs}=0,2,10,100

- Hint: the probability to find N in the interval $N_{obs} \leq N \leq \infty$ is given by:

$$\text{Probability: } \sum_{N \geq N_{obs}}^{\infty} \frac{e^{-\mu}(\mu)^N}{N!} = 1-\alpha = 1-\text{TMath::Prob}(2*\mu, 2*N_{obs})$$

$$\text{Inverse function: } 2*\mu = \text{TMath::ChisquareQuantile}(1-\alpha, 2*N_{obs})$$

(a)

μ	N_{obs}	1-α
2		
3		
5		
10		

(b)

N_{obs}	μ_{limit}
0	
2	
10	
100	

Exercise 2 (Neyman construction)

- Poisson experiment, determine limits on the parameter μ, given N_{obs}
  - a) determine the range $N_{obs} \leq N \leq \infty$ for CL=0.95 and μ=2,3,5,10. What is the probability to find the measurement in these ranges
  - b) determine the limit on μ for N_{obs}=0,2,10,100
- Hint: the probability to find N in the interval $N_{obs} \leq N \leq \infty$ is given by:

$$\text{Probability: } \sum_{N \geq N_{obs}}^{\infty} \frac{e^{-\mu}(\mu)^N}{N!} = 1-\alpha = 1-\text{TMath::Prob}(2*\mu, 2*N_{obs})$$

$$\text{Inverse function: } 2*\mu = \text{TMath::ChisquareQuantile}(1-\alpha, 2*N_{obs})$$

(a)

μ	N_{obs}	1-α
2	0	1
3	1	0.95
5	2	0.96
10	5	0.97

(b)

N_{obs}	μ_{limit}
0	3.0
2	6.3
10	17.0
100	118.1

Coverage

- Coverage: given a limit procedure, calculate for each θ the probability to exclude the theory
- Poisson example (exercise 2)

$$P_{\mathrm{excl}}(\mu_{\mathrm{truth}})=\sum_N P_{\mu,\mathrm{truth}}(N)\,\Theta(\mu_{\mathrm{truth}}\leq\mu_{\mathrm{limit}}(N))$$

$$\text{where }\Theta(\mu_{\mathrm{truth}}\leq\mu_{\mathrm{limit}})=\begin{cases}1 & \text{if } \mu_{\mathrm{truth}}\leq\mu_{\mathrm{limit}}\\ 0 & \text{otherwise}\end{cases}$$

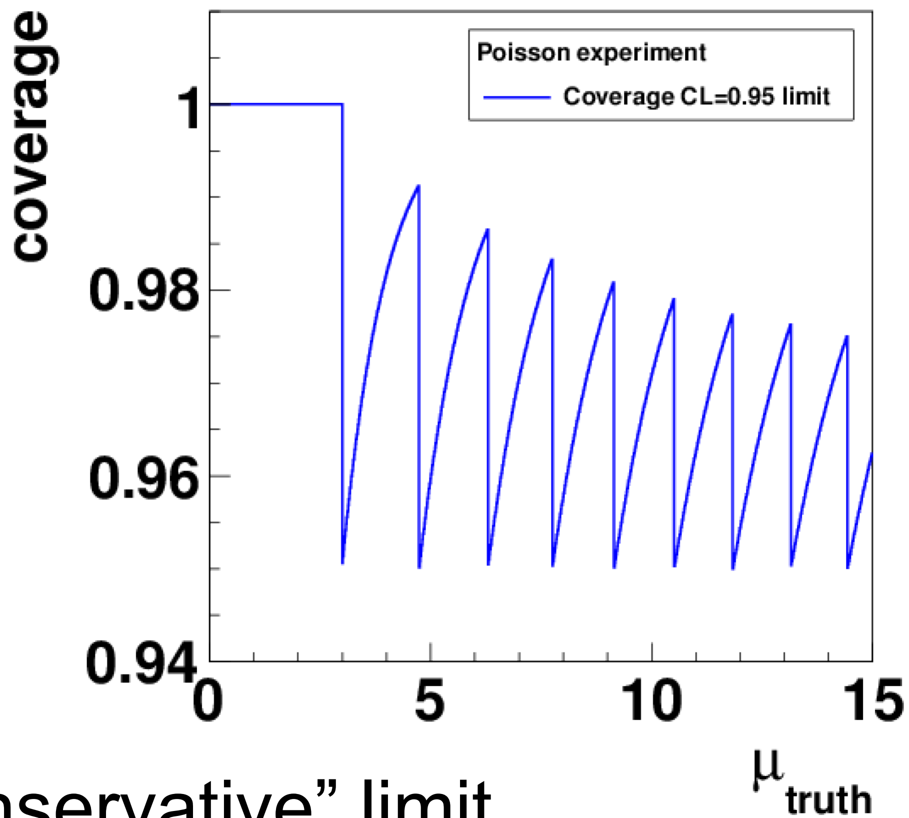

- coverage=0.95: exact coverage
- coverage<0.95: undercoverage
- coverage>0.95: overcoverage, “conservative” limit
- “Simple” Poisson case: overcoverage (discrete measurement)

Bayesian limits

- Bayesian limit: exclude a set of theories, such that the posterior probability of the excluded theories is 1-CL

Enumerator: integrate over allowed theories

$$CL = P(\mu \leq \mu_0 | N = N_{obs}) = 1 - \alpha = \frac{\int_0^{\mu_0} P(N = N_{obs}|\mu)\, P(\mu)\, d\mu}{\int_0^{\infty} P(N = N_{obs}|\mu)\, P(\mu)\, d\mu}$$

Denominator: integrate all theories (normalisation)

$P(\mu)$: prior probability of the model μ

$P(N = N_{obs}|\mu)$: Likelihood

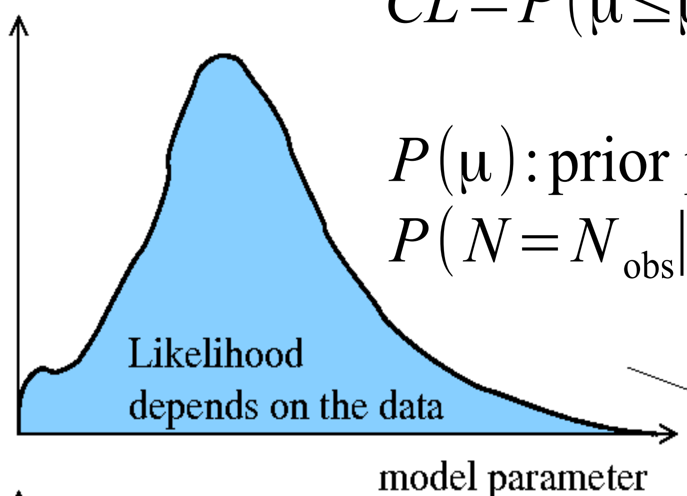

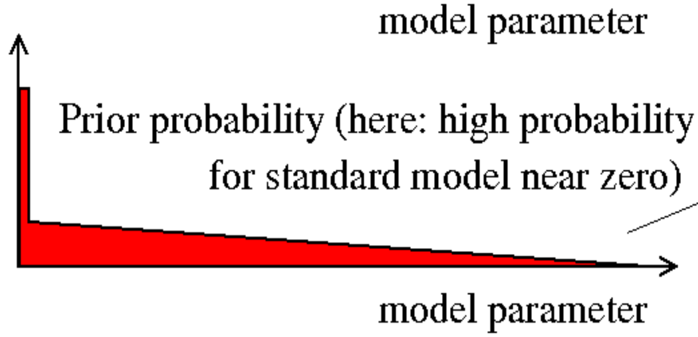

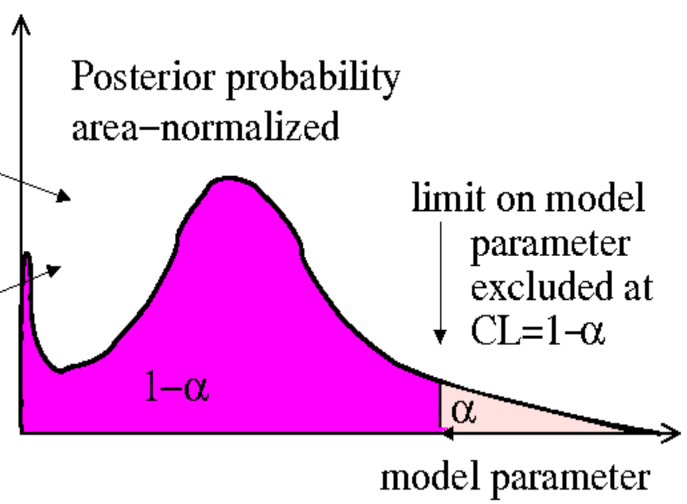

Bayesian limit: integrate over models, fixed N_{obs}

Frequentist limit: integrate over N_{obs}, test each model

Exercise 3 (Bayesian limit)

- Exercise 3a: Bayesian limit for N_{obs}=0,2,10,100 (flat prior) (use Root macro)
- Exercise 3b: use a prior P(μ)=μ, N_{obs}={0,2,10,100}
- Exercise 3c: use a flat prior up to μ_{max}=90, set prior to zero above μ_{max}
- Compare to exercise 2

- Bayesian limit with arbitrary prior → numerical integration
- `GetPosterior.C(muLimit,nObs)`

$$\text{Posterior} \sim \int_0^{\mu_0} d\mu \, \text{Prior}(\mu) \frac{\exp[-\mu]\mu^{N_{obs}}}{N_{obs}!}$$

- Vary muLimit until Posterior=0.95

	frequentist	Bayes flat	Bayes P(μ)=μ	Bayes flat μ_{max}=90
N_{obs}	μ_{limit}	μ_{limit}	μ_{limit}	μ_{limit}
0	3.0			
2	6.3			
10	17.0			
100	118.1			

Bayesian limit exercise

- Exercise 3a: Bayesian limit for N_{obs}=0,2,10,100 (flat prior) (use Root macro)
- Exercise 3b: use a prior P(μ)=μ, N_{obs}={0,2,10,100}
- Exercise 3c: use a flat prior up to μ_{max}=90, set prior to zero above μ_{max}
- Compare to exercise 2

- For this example: Bayes flat=Frequentist
- Prior P(μ)=μ gives more conservative limit
- μ_{max}=90 fails for N_{obs}=100

	frequentist	Bayes flat	Bayes P(μ)=μ	Bayes flat μ_{max}=90
N_{obs}	μ_{limit}	μ_{limit}	μ_{limit}	μ_{limit}
0	3.0	3.0	4.7	3.0
2	6.3	6.3	7.8	6.3
10	17.0	17.0	18.2	17.0
100	118.1	118.2	119.3	89.7

Lecture Part I, Summary

- Setting limits: related to parameter estimation, hypothesis tests
- Limit: special case of a confidence interval (single-sided)

- Frequentist limit: Neyman construction (sum over observations)
- Concept of “coverage”: test the validity of the limit procedure

- Bayesian limit: integral over parameter of interest
- Dependence on the choice of prior (for parameter of interest)

Limits with background

- Expected number of events: sum of a signal and background cross section, times integrated luminosity

  $\mu = s + b, \; s, b$: signal and background event yield, respectively

  - s=0: standard model
  - s>0: new physics
- Assume background known. What is the limit on the signal?
- Frequentist: set limit on μ, then subtract b
- Bayesian: use prior probability which is zero for s<0

Exercise 4 (limit with background)

- Calculate Frequentist and Bayesian limits for N_{obs}={0,2} and b={0.5,2.0,3.5}

Poisson parameter: $\mu = s + b$

	b=0.5		b=2.0		b=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
Bayesian						
Frequentist						

- Frequentist: use methods from exercise 2
- Bayes: try to modify exercise 3 macro, or use macro GetPosteriorWithBackground.C

Exercise 4 (limit with background)

- Calculate Frequentist and Bayesian limits for N_{obs}={0,2} and b={0.5,2.0,3.5}

Poisson parameter: $\mu = s + b$

	b=0.5		b=2.0		b=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
Bayesian	3.0	5.8	3.0	4.8	3.0	4.3
Frequentist	2.5	5.8	1.0	4.3	-0.5	2.8

- Problem for Frequentist limit, N_{obs}=0 and b=3.5:

  limit excludes all signal above s=-0.5.

  Even the “standard model” s=0 is excluded

Discussion Exercise 4

- Frequentist analysis can give limits where all models are "excluded" at a given CL (even the model with s=0)

  N_{obs}=0, μ =s+b, b=3.5

  → limit s<-0.5 @ 95% CL but s>=0 physical bound

- Bayesian limit uses prior knowledge s>=0

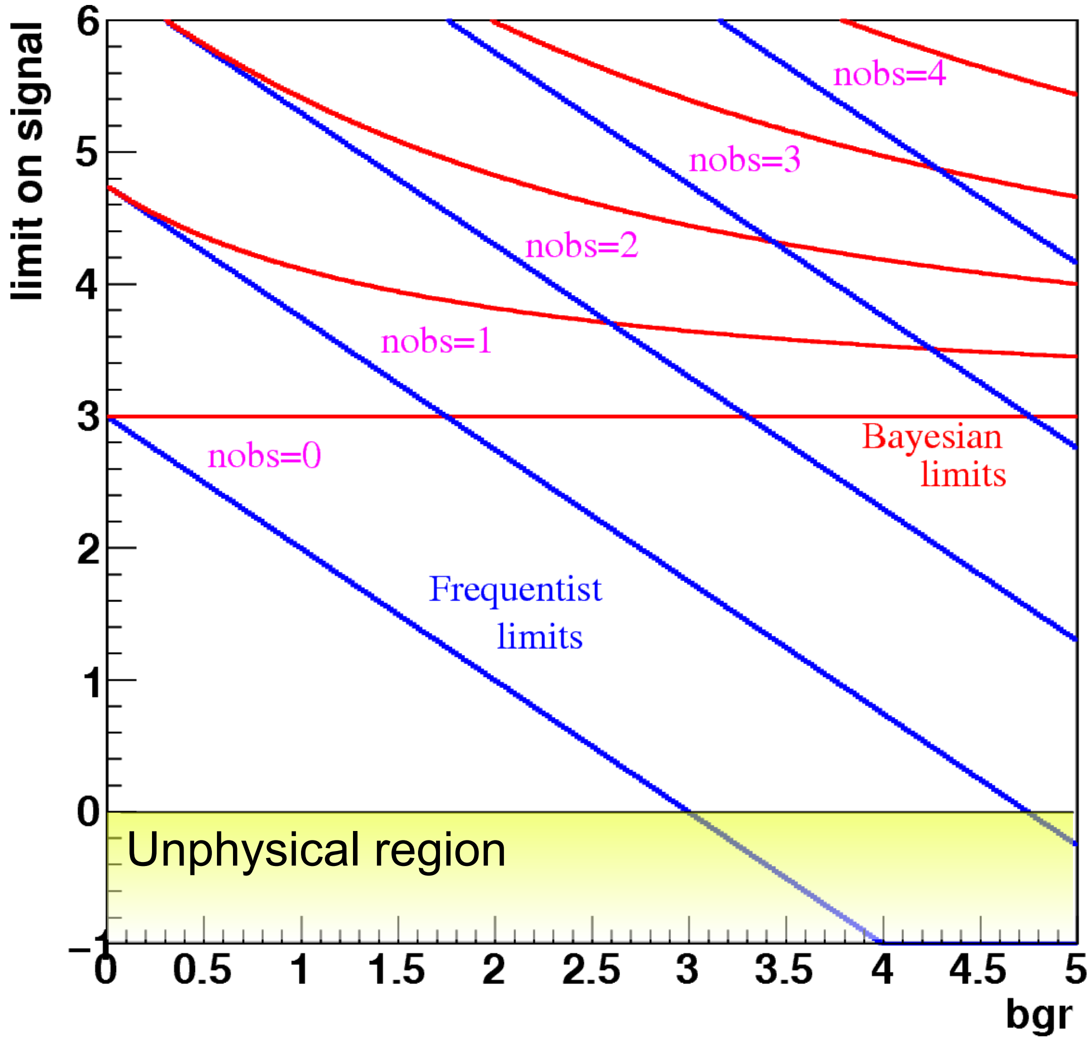

Limits near a boundary

- What to do if frequentist analysis excludes parameters beyond the sensitivity of the experiment or beyond boundaries?
- Give expected limit to show sensitivity of the experiment (exercise 5)
- CL_s method, also known as “modified frequentist” (exercise 6)
- Bayesian methods (see exercise 4)

Expected limit (exercise 5)

- Expected limit: limit weighted by background probability

$$\langle s_{\text{limit}} \rangle = \sum_{n=0}^{\infty} \frac{e^{-b} b^{n}}{n!} \text{LimitOnSignal}(b, n)$$

	b=0.5		b=2.0		b=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
Bayesian	3.0	5.8	3.0	4.8	3.0	4.3
Frequentist	2.5	5.8	1.0	4.3	-0.5	2.8
Expected						

- Calculate expected limits for b={0.5,2.0,3.5}
- Macro GetExpectedLimit.C

Expected limit (exercise 5)

- Expected limit: limit weighted by background probability

$$\langle s_{\text{limit}} \rangle = \sum_{n=0}^{\infty} \frac{e^{-b} b^n}{n!} \text{LimitOnSignal}(b, n)$$

	b=0.5		b=2.0		b=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
Bayesian	3.0	5.8	3.0	4.8	3.0	4.3
Frequentist	2.5	5.8	1.0	4.3	-0.5	2.8
Expected	3.3		4.2		4.9	

- Problematic case: expected limit differs a lot from observed limit

  → Recognize statistical fluctuation or problem with background

The CL_S (modified frequentist) method

- Frequentist limit: $1-CL \geq \alpha = \mathrm{CL}_{SB} = P(N \leq N_{\mathrm{obs}}; \mu = s+b)$
- CL_S limit: $$1-CL \geq \mathrm{CL}_S = \frac{\mathrm{CL}_{SB}}{\mathrm{CL}_B} = \frac{P(N \leq N_{\mathrm{obs}}; \mu = s+b)}{P(N \leq N_{\mathrm{obs}}; \mu = b)}$$
- Probability is normalized to background probability
- $CL_B \leq 1 \rightarrow CL_S \geq CL_{SB}$: same α requires larger signal

  Limit is “conservative”
- For zero signal: $CL_S = 1$

  → zero signal is never excluded

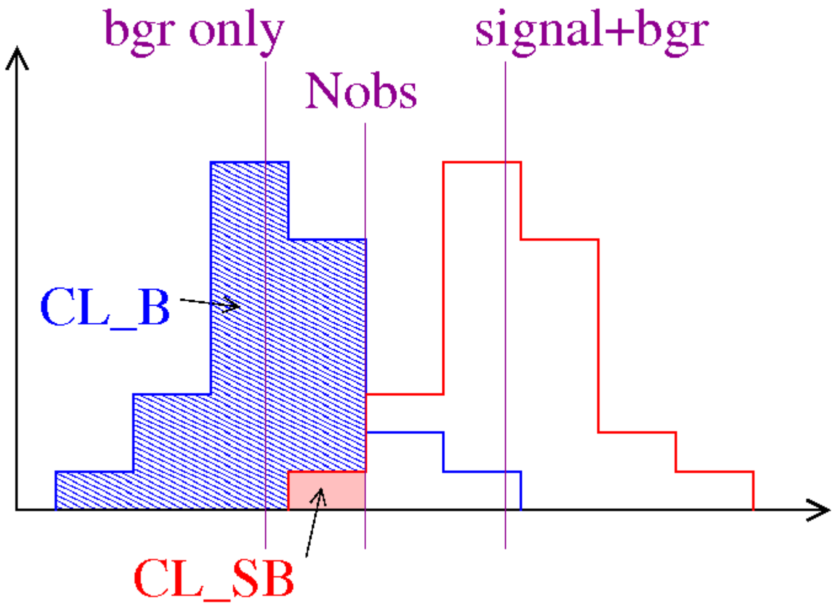

Exercise 6 (CL_S method)

- Frequentist limit: $1-CL \geq \alpha = \mathrm{CL}_{SB} = P(N \leq N_{\mathrm{obs}}; \mu = s+b)$
- CL_S limit: $1-CL \geq \mathrm{CL}_S = \frac{\mathrm{CL}_{SB}}{\mathrm{CL}_B} = \frac{P(N \leq N_{\mathrm{obs}}; \mu = s+b)}{P(N \leq N_{\mathrm{obs}}; \mu = b)}$

	b=0.5		b=2.0		b=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
Bayesian	3.0	5.8	3.0	4.8	3.0	4.3
Frequentist	2.5	5.8	1.0	4.3	-0.5	2.8
CL_S						
Expected	3.3		4.2		4.9	

- Use macro GetCLsLimit.C to calculate CL_S, iterate to get limit

Exercise 6 (CL_S method)

- Frequentist limit: $$1-CL \geq \alpha = \mathrm{CL}_{SB} = P(N \leq N_{\text{obs}}; \mu = s+b)$$
- CL_S limit: $$1-CL \geq \mathrm{CL}_S = \frac{\mathrm{CL}_{SB}}{\mathrm{CL}_B} = \frac{P(N \leq N_{\text{obs}}; \mu = s+b)}{P(N \leq N_{\text{obs}}; \mu = b)}$$

	b=0.5		b=2.0		b=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
Bayesian	3.0	5.8	3.0	4.8	3.0	4.3
Frequentist	2.5	5.8	1.0	4.3	-0.5	2.8
CL_S	3.0	5.8	3.0	4.8	3.0	4.3
Expected	3.3		4.2		4.9	

- For this example, CL_S is identical to Bayesian (with flat prior)

Limits with background, coverage

- CL_s method avoids problem with limits better than the experiments sensitivity
- Disadvantage: CL_s method is conservative, in particular for small signals

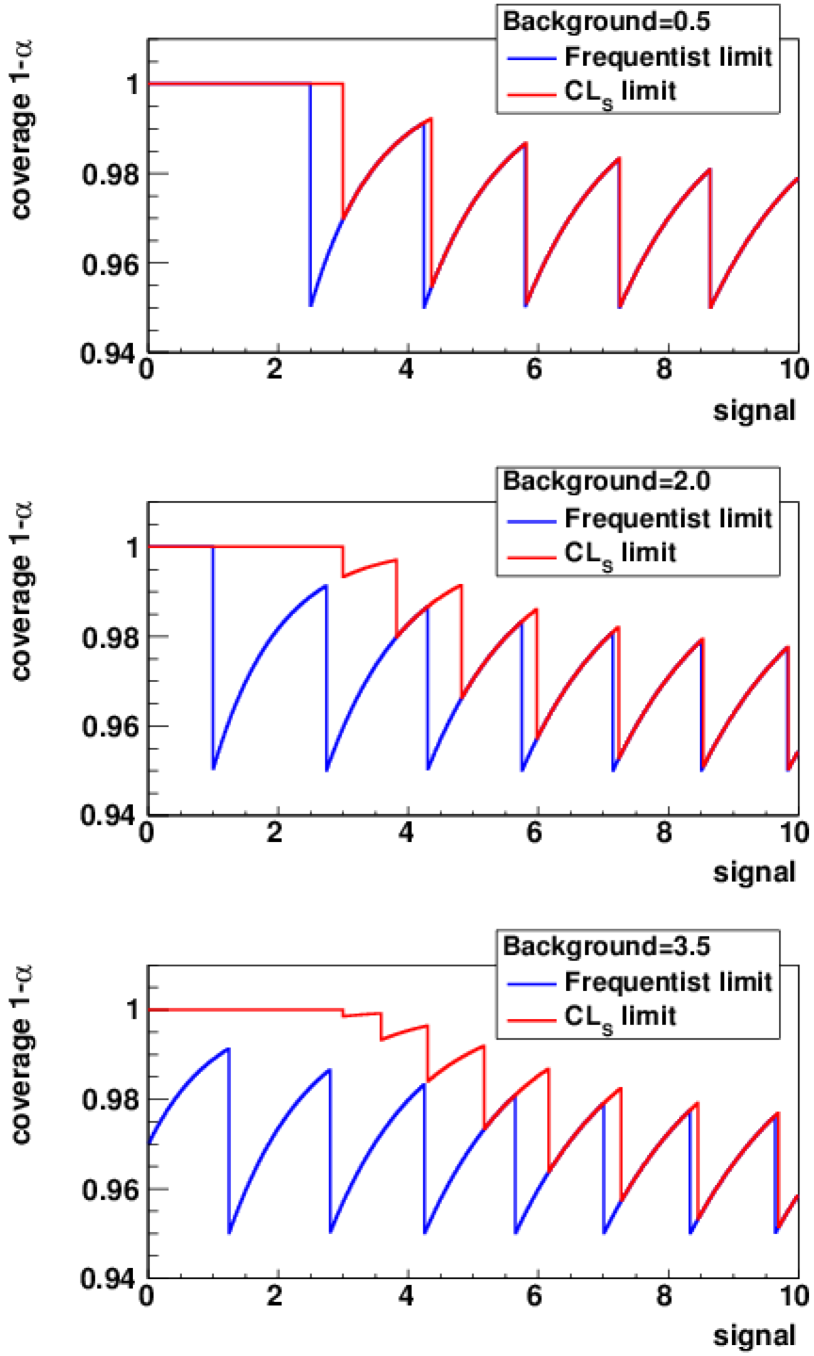

Lecture part II, summary

- Poisson experiment with background
- Unnaturally good limit if number of events is much smaller than background expectation
- “Solutions”:
  - Quote expected limit (sensitivity of the experiment)
  - CLS method (never excludes background-only)
  - Bayesian method (prior knows about boundaries)

Systematic errors, multiple bins/channels

- Examples discussed so far : events are counted in a single channel, no systematic errors
- General case: several channels (or bins) and systematic errors
- Example: mass distribution with N bins (signal and bgr shape)

  → N channels to be combined

  → Background normalisation error

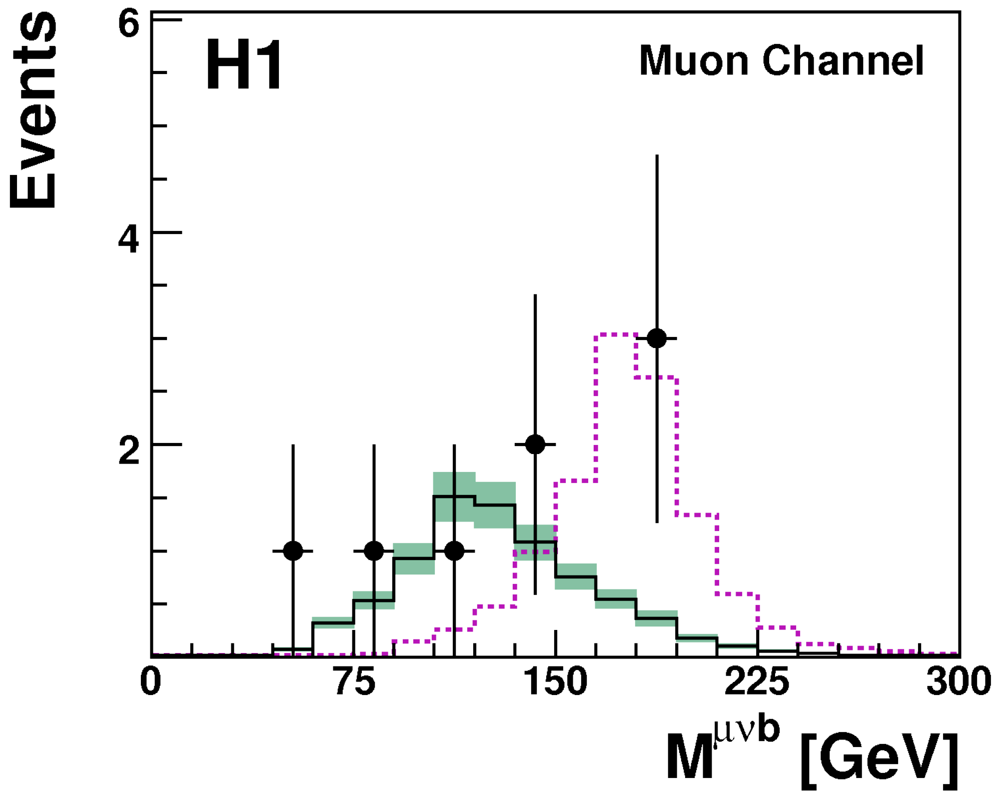

Example plot: search for single top production at HERA, Phys.Lett. B678 (2009) 450

Simple example with two syst. errrors

- Consider signal in one bin

$$\mu = l(s+b),\ \ l: \text{integrated luminosity},\ s, b: \text{signal, background cross sections}$$

with systematic errors:

$$l = l_{\text{obs}} \pm \sigma_l,\ b = b_{\text{obs}} \pm \sigma_b$$

- Full probability density has three contributions

$$P_{s,l,b}(N_{\text{obs}}, l_{\text{obs}}, b_{\text{obs}}) = \underbrace{\frac{e^{-l(s+b)}(l(s+b))^{N_{\text{obs}}}}{N_{\text{obs}}!}}_{\text{event counting}} \underbrace{\frac{1}{\sqrt{2\pi}\sigma_l} e^{\frac{-(l-l_{\text{obs}})^2}{2\sigma_l^2}}}_{\text{measurement of l}} \underbrace{\frac{1}{\sqrt{2\pi}\sigma_b} e^{\frac{-(b-b_{\text{obs}})^2}{2\sigma_b^2}}}_{\text{measurement of b}}$$

- Three channels (measurements): N_{obs}, l_{obs}, b_{obs}
- Nuisances l,b and parameter of interest s

Simple Example???

$$P_{s,l,b}(N_{\text{obs}},l_{\text{obs}},b_{\text{obs}})=\underbrace{\frac{e^{-l(s+b)}(l(s+b))^{N_{\text{obs}}}}{N_{\text{obs}}!}}_{\text{event counting}}\underbrace{\frac{1}{\sqrt{2\pi}\sigma_l}e^{\frac{-(l-l_{\text{obs}})^2}{2\sigma_l^2}}}_{\text{measurement of l}}\underbrace{\frac{1}{\sqrt{2\pi}\sigma_b}e^{\frac{-(b-b_{\text{obs}})^2}{2\sigma_b^2}}}_{\text{measurement of b}}$$

- Three observables: N_{obs}, l_{obs}, b_{obs}
- Nuisances l,b and parameter of interest s
- This looks quite complicated already
- Observed: N_{obs}, l_{obs}, b_{obs} and parameters l, b, s

  → Neyman construction in six dimensions? Perhaps not...
- How to get rid of nuisance parameters?
- How to combine channels (measurements)?

Bayesian method

$$L(s,l,b) = \underbrace{\frac{e^{-l(s+b)}\left(l(s+b)\right)^{N}_{\text{obs}}}{N_{\text{obs}}!}}_{\text{event counting}} \underbrace{\frac{1}{\sqrt{2\pi}\sigma_l} e^{\frac{-(l-l_{\text{obs}})^2}{2\sigma_l^2}}}_{\text{measurement of l}} \underbrace{\frac{1}{\sqrt{2\pi}\sigma_b} e^{\frac{-(b-b_{\text{obs}})^2}{2\sigma_b^2}}}_{\text{measurement of b}}$$

- Bayesian treatment of nuisance parameters: the measurements l_{obs} and b_{obs} correspond to priors for l,b
- Define marginalized likelihood, where the nuisances are integrated out

$$L(s) = \int dl\, db\, L(s,l,b)$$

- Only depends on s (and the observations N_{obs}, l_{obs}, b_{obs})
- Analysis (Bayesian) as for the case without systematic errors

Hybrid method

$$P_{s,l,b}(N_{obs}, l_{obs}, s_{obs}) = \underbrace{\frac{e^{-l(s+b)}(l(s+b))^{N_{obs}}}{N_{obs}!}}_{\text{event counting}} \underbrace{\frac{1}{\sqrt{2\pi}\sigma_l} e^{\frac{-(l-l_{obs})^2}{2\sigma_l^2}}}_{\text{measurement of l}} \underbrace{\frac{1}{\sqrt{2\pi}\sigma_b} e^{\frac{-(b-b_{obs})^2}{2\sigma_b^2}}}_{\text{measurement of b}}$$

- Bayesian treatment of nuisance parameters: the measurements l_{obs} and b_{obs} correspond to priors for l,b
- Use marginalized likelihood as if it were the probability density for N_{obs} (after integrating out the nuisances)

$$P_{s,marginalized}(N_{obs}) = \int dl\, db\, P_{s,l,b}(N_{obs}, \ldots)$$

- Only depends on s (and the observation N_{obs})
- Analysis (Neyman) as for the case without systematic errors

Exercise 7 (limits from hybrid method)

- CL_S limit, systematic error treated with hybrid method $\mu = l(s+b)$
- Background error: zero or σ_b=50% [b_{obs}={0.5, 3.5}]
- Luminosity error: zero or σ_l=10% [l_{obs}=1.0]

CL_S limits	b=0.5		bgr=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
No syst	3.0	5.8	3.0	4.3
σ_b/b=50%				
σ_l/l=10%				
Both syst.				

Use root macro GetClsSys.C

Exercise 7 (limits from hybrid method)

- Typical example for the use of Monte Carlo methods to calculate probabilities
- Probabilities are calculated by counting the outcomes of toy experiments

```
l=rnd->Gaus(1.0,dLumi);
b=rnd->Gaus(bgr,dBgr);
Int_t n_b=rnd->Poisson(l*b);
Int_t n_sb=rnd->Poisson(l*(signal+b));
. . .
if(n_b<=nobs) nexp_b += 1.0;
if(n_sb<=nobs) nexp_sb += 1.0;
. . .
Double_t cl_s=nexp_sb/nexp_b;
```

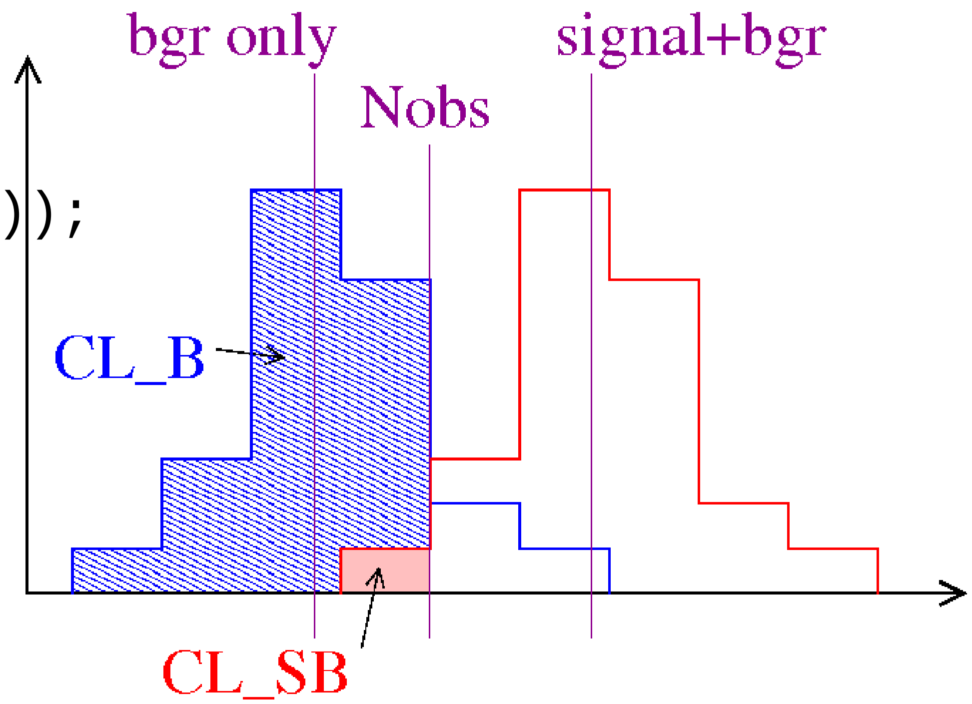

Exercise 7 (limits from hybrid method)

- Background error: zero or σ_b=50% [b_{obs}={0.5, 3.5}]
- Luminosity error: zero or σ_l=10% [l_{obs}=1.0]

$$\mu = L(s+b)$$

- Systematic errors make limits somewhat worse

CL_S limits	b=0.5		bgr=3.5	
	N_{obs}=0	N_{obs}=2	N_{obs}=0	N_{obs}=2
No syst	3.0	5.8	3.0	4.3
σ_b/b=50%	3.0	5.8	3.0	4.9
σ_l/l=10%	3.1	6.0	3.2	4.5
Both syst.	3.1	6.0	3.1	5.0

Small background: background error has little influence

Large background: background error has larger influence

Luminosity error visible in all cases

Multiple bins/channels

- Case of multiple bins/channels brings additional complication for Frequentist analysis:
- Many bins: vector of observations
- Neyman construction: not possible for a vector of observations
- Solution: define 1-dimensional random variable (test statistic X)

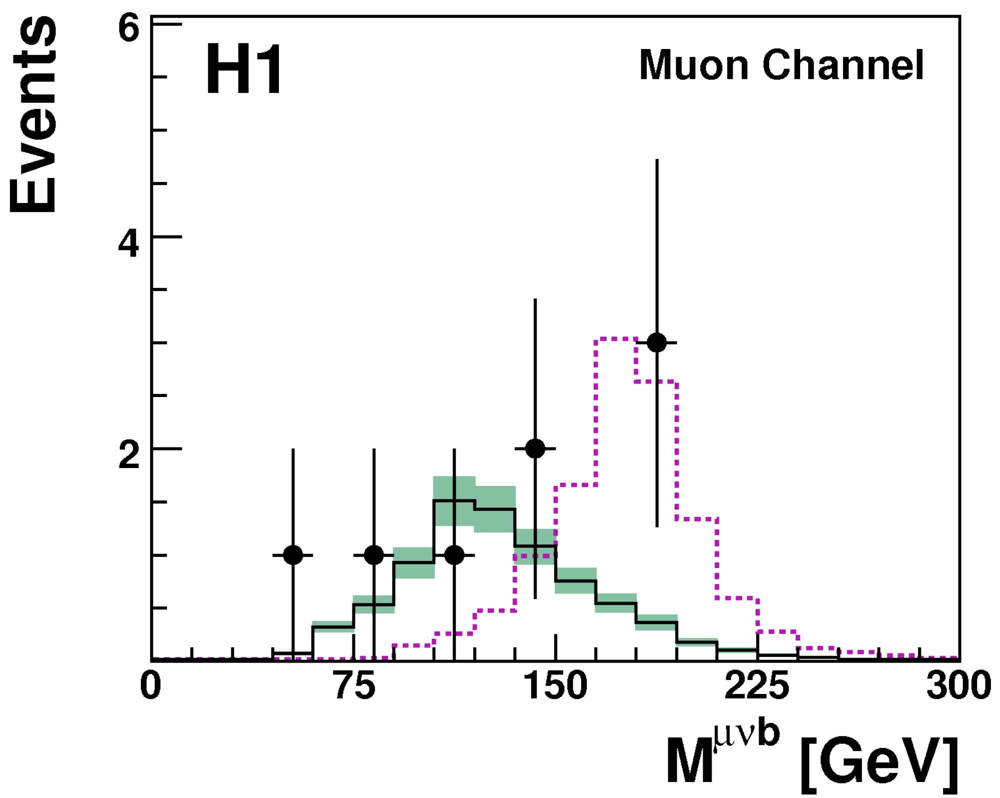

Test statistic:
X=X(N1,N2,N3,...)
where N1,N2,N3... are the event counts in bin 1,2,3,... respectively

Choice of the test statistics

- Log Likelihood ratio $X = \log\left[\frac{L(\text{signal+bgr})}{L(\text{bgr})}\right]$

- Likelihood normalized to maximum $X = \log\left[\frac{L(\text{signal+bgr})}{L_{\max}}\right]$

- Other choices: weighted sum of all channels, weight taken from signal/bgr ratio or something similar

$$X = \sum w_i N_i^{\text{obs}}, \text{ where for example } w_i = \frac{(s_i + b_i) - b_i}{(s_i + b_i) + b_i}$$

- Note: log of likelihood ratio also is a weighted sum:

$$\log L(\text{signal+bgr}) - \log L(\text{bgr}) \sim \sum_i \underbrace{\log\left(1 + \frac{s_i}{b_i}\right)}_{w_i} N_i^{\text{obs}}$$

What is a good test statistic?

- Good sensitivity to signal
- Little sensitivity to systematic effects
- Ideal case: probability density P(X) of test statistic is largely independent of the nuisances $\rightarrow$ use of hybrid method not needed $\rightarrow$ pure frequentist limit

- "Standard" choice: profile likelihood ratio
- Idea: nuisances are estimated from the data

$$X(s|\text{measurements}) = -2\log\left[\frac{L(s,\hat{\theta}(s))}{L(\hat{s},\hat{\hat{\theta}})}\right]$$

s: signal strength, θ: nuisances
$L(s,\theta)$: Likelihood function
$\hat{\theta}(s)$: value of θ which maximizes L given s
$\hat{s}$ and $\hat{\hat{\theta}}$: global maximum of L

Profile likelihood ratio

- Basic idea: nuisances are estimated from the data
- Likelihood ratio: maximum indicates signal position

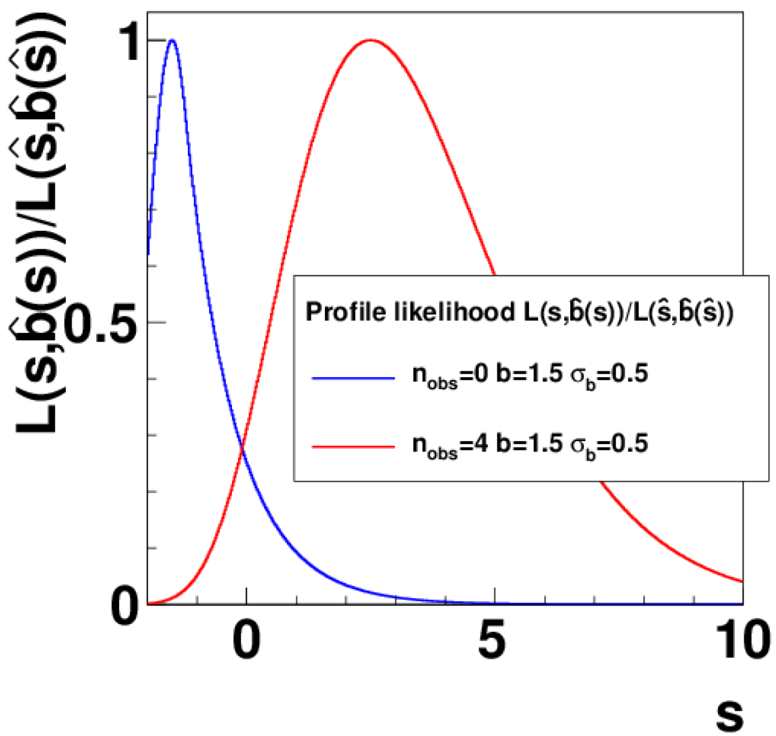

- Numerical analysis: use -2*log(likelihood ratio)
- X(s) has a minimum near the best signal

$$X(s|\text{measurements}) = -2\log\left[\frac{L(s,\hat{\theta}(s))}{L(\hat{s},\hat{\hat{\theta}})}\right]$$

s: signal strength, θ: nuisances
$L(s,\theta)$: Likelihood function
$\hat{\theta}(s)$: value of θ which maximizes L given s
$\hat{s}$ and $\hat{\hat{\theta}}$: global maximum of L

Profile likelihood analysis

- Profile likelihood is expected to have probability density

  $$P(X) \simeq \chi^2_{ndf=1}(X)$$

  in the large sample limit
- Direct access to CL_{SB}, CL_S using Tmath::Prob()
- Need to verify P(X) and dependence on nuisances with Monte Carlo methods

- Example: $\mu = s + b$
- Measurements: N_{obs} and b_{obs}
- Vary b → some influence

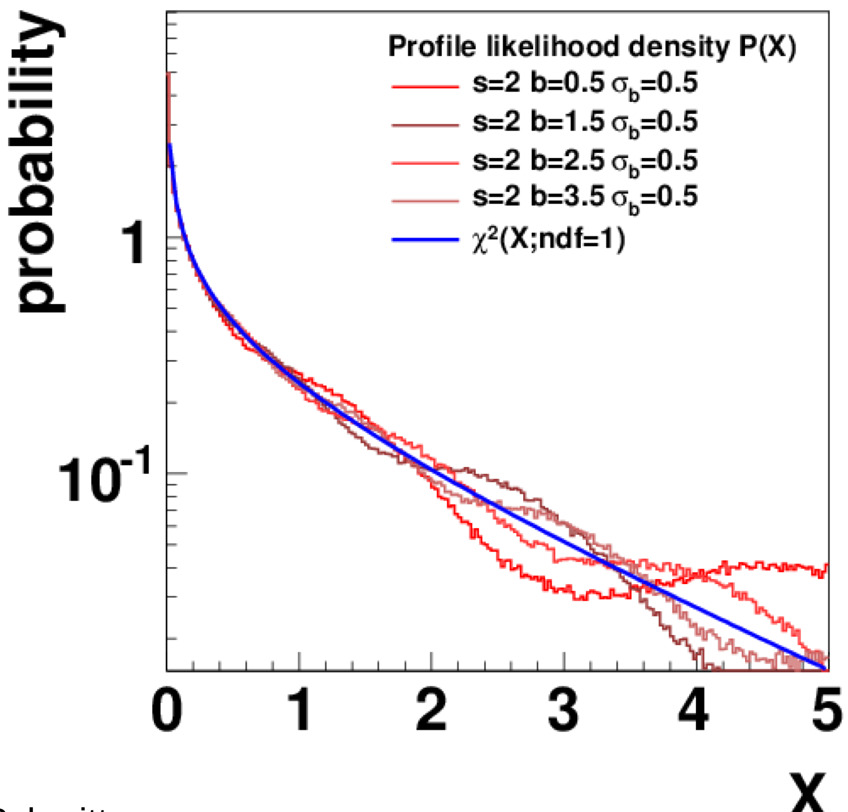

Summary

- Basic concepts of setting limits:
  - Frequentist/Bayesian methods
  - Coverage, expected limit, CL_S method
  - Systematic errors, nuisances, marginalization
  - Combining channels: test statistic, e.g. likelihood
  - Standard method: profile likelihood

Backup

Calculation of Poisson sums

- Sum over Poisson terms is related to χ^2 distribution with number-of-degrees of freedom “k”:

$$\chi^2(x;k)=\frac{x^{k/2-1}e^{-x/2}}{2^{k/2}\Gamma(k/2)} \qquad P(N;\mu)=\frac{e^{-\mu}\mu^N}{N!}$$

- Poisson sum equals integral over χ^2 distribution (partial integration)

$$\alpha(\mu,N)=\int_{2\mu}^{\infty}\chi^2(x;2(N+1))\,\mathrm{d}x=\sum_{i=0}^{N}P(i;\mu)$$

- Standard functions for χ^2 integrals:

`α(μ,N)=TMath::Prob(2*μ,2*(N+1))` and

`μ=0.5*TMath::ChisquareQuantile(1-α,2*(N+1))`

Frequentist upper limit, Gaussian case

$$CL = \int_{x_{\mathrm{obs}}}^{\infty} \exp\left[-1/2\left(\frac{x-\mu_{\mathrm{truth}}}{\sigma}\right)^2\right]\frac{1}{\sqrt{2\pi}\,\sigma}\,\mathrm{d}x$$

- Fixed σ, measurement x_{obs}, parameter of interest μ_{truth}
- Define 95% probability area under Gaussian
- If μ_{truth} is too large, it is outside the 95% → excluded

Limits with background, comparison

- Frequentist limit may become "unphysicsal" or "too good"
- Expected limit: sensitivity of the experiment
- CL_S method: normalize to "standard model", never exclude zero signal
- Disadvantage of CL_S? Study coverage