Establishing Signals, Excluding Parameters

Roger Wolf

19. November 2015

INSTITUTE OF EXPERIMENTAL PARTICLE PHYSICS (IEKP) – PHYSICS FACULTY

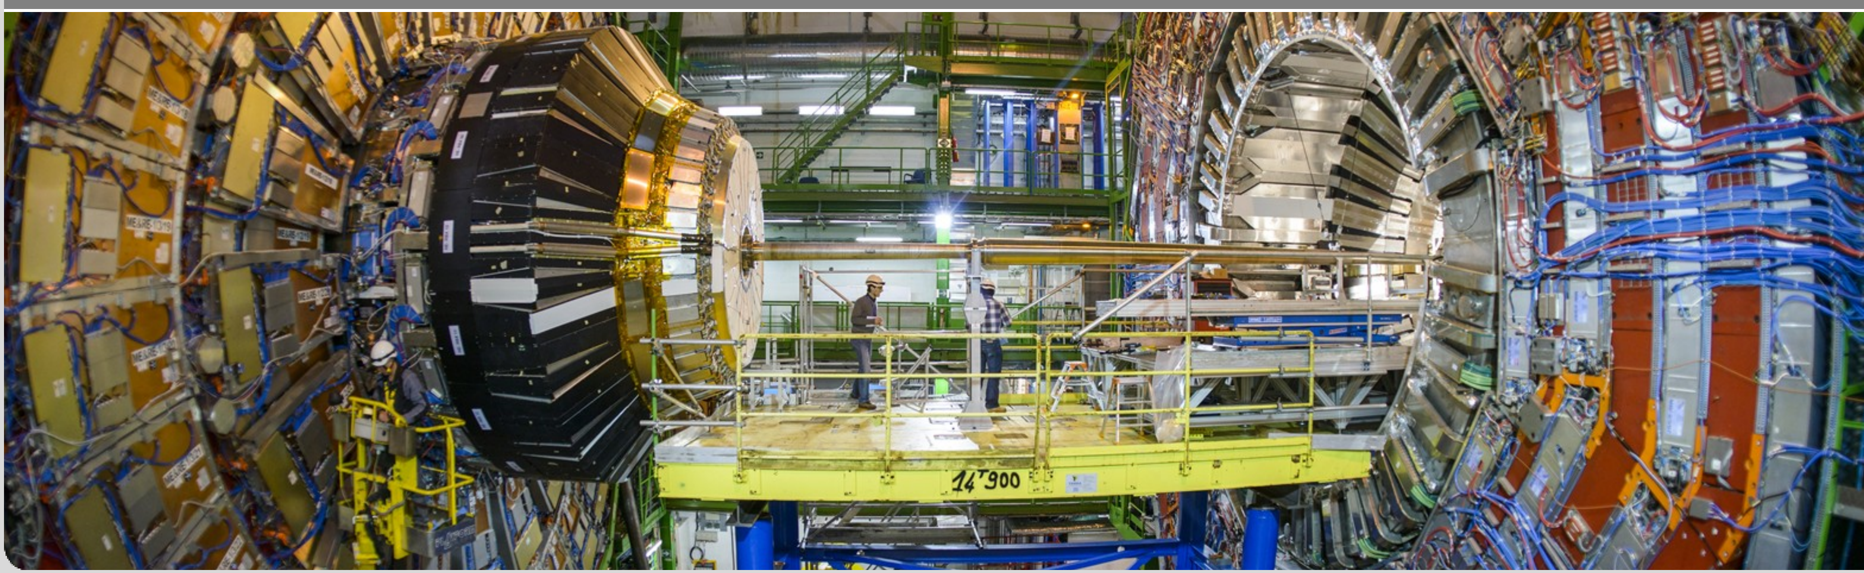

Statistics vs. particle physics

Experiment:

- All measurements we do are derived from rate measurements.
- We record millions of trillions of particle collisions.
- Each of these collisions is independent from all the others.

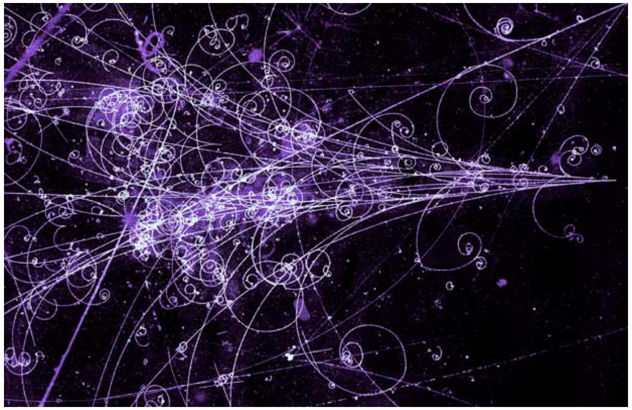

Theory:

- QM wave functions are interpreted as probability density functions.
- The Matrix Element, $\mathcal{S}_{fi}$,gives the probability to find final state f for given initial state i.
- Each of the statistical processes $pdf \rightarrow ME \rightarrow hadronization \rightarrow$ $energy\ loss\ in\ material \rightarrow digitization$ are statistically independent.
- Event by event simulation using Monte Carlo integration methods.

- Particle physics experiments are a perfect application for statistical methods.

Statistics vs. probability theory (stochastic)

Test statistic:

$\Omega^n \to \mathbb{R}: \qquad x \to f(x)$

- NLL ($q = -2\ln(\mathcal{L}_1/\mathcal{L}_0)$).
- *Boosted Decision Tree* BDT output.

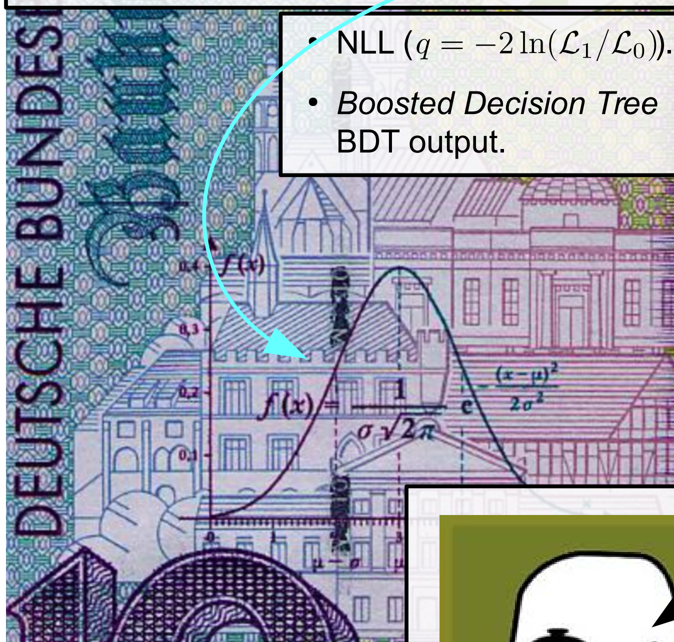

Probability (density) function:

$\Omega^n \to [0,1] \subset \mathbb{R}: \qquad x \to \mathcal{P}(x)$

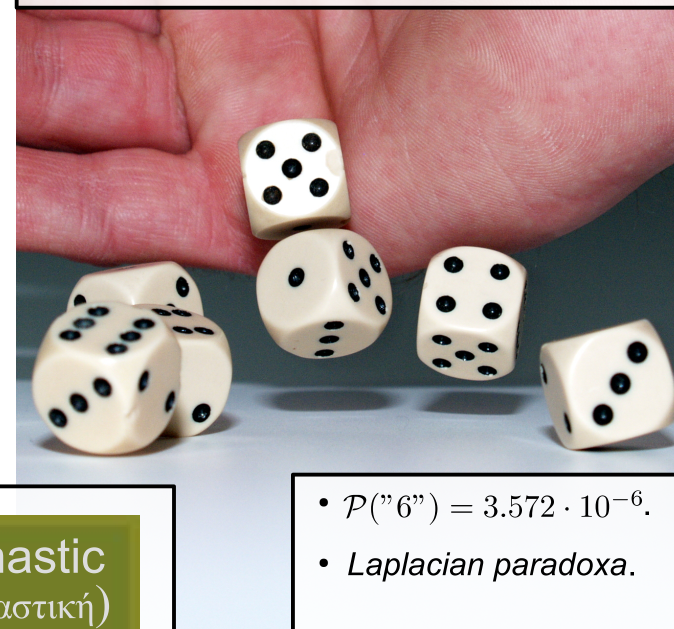

- $\mathcal{P}(\text{"6"}) = 3.572 \cdot 10^{-6}$.
- *Laplacian paradoxa*.

- Problem of statistics is usually *ill-defined*.
- Deduce *truth* from shadows in Platon's cave...

The case of “truth”

- Deduce *truth* from shadows:

Usually phrased in form of (nested) models (=*ideas* for Platon):

- Mathematically model = hypothesis.

Uncertainty model:

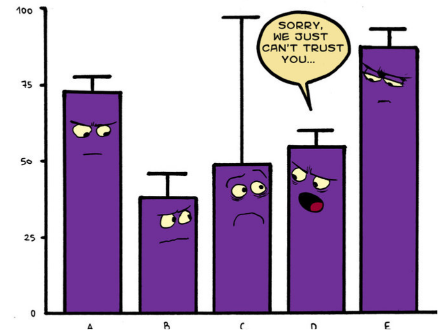

Usually determined to best knowledge (not questioned)

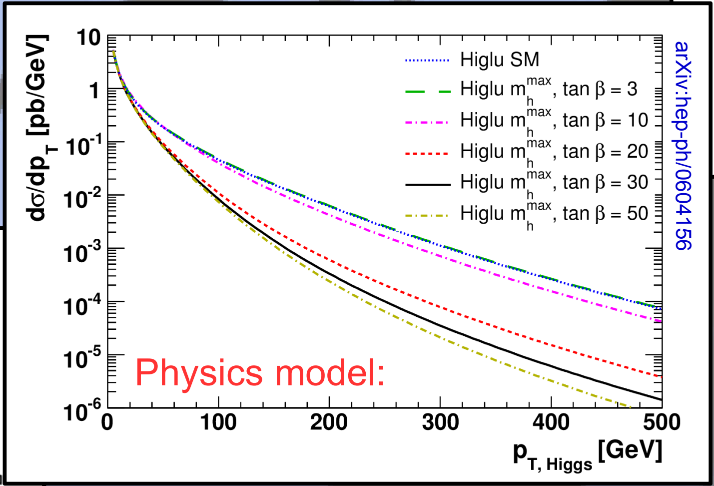

Usually competing models/ hypothe-ses will be discussed here!

Statistics model:

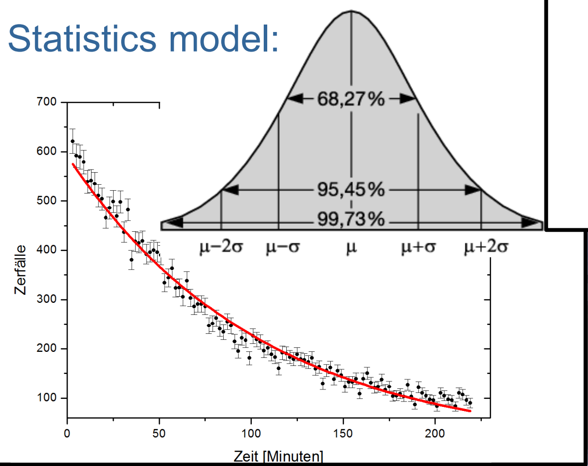

Usually not questioned

Models in counting experiments

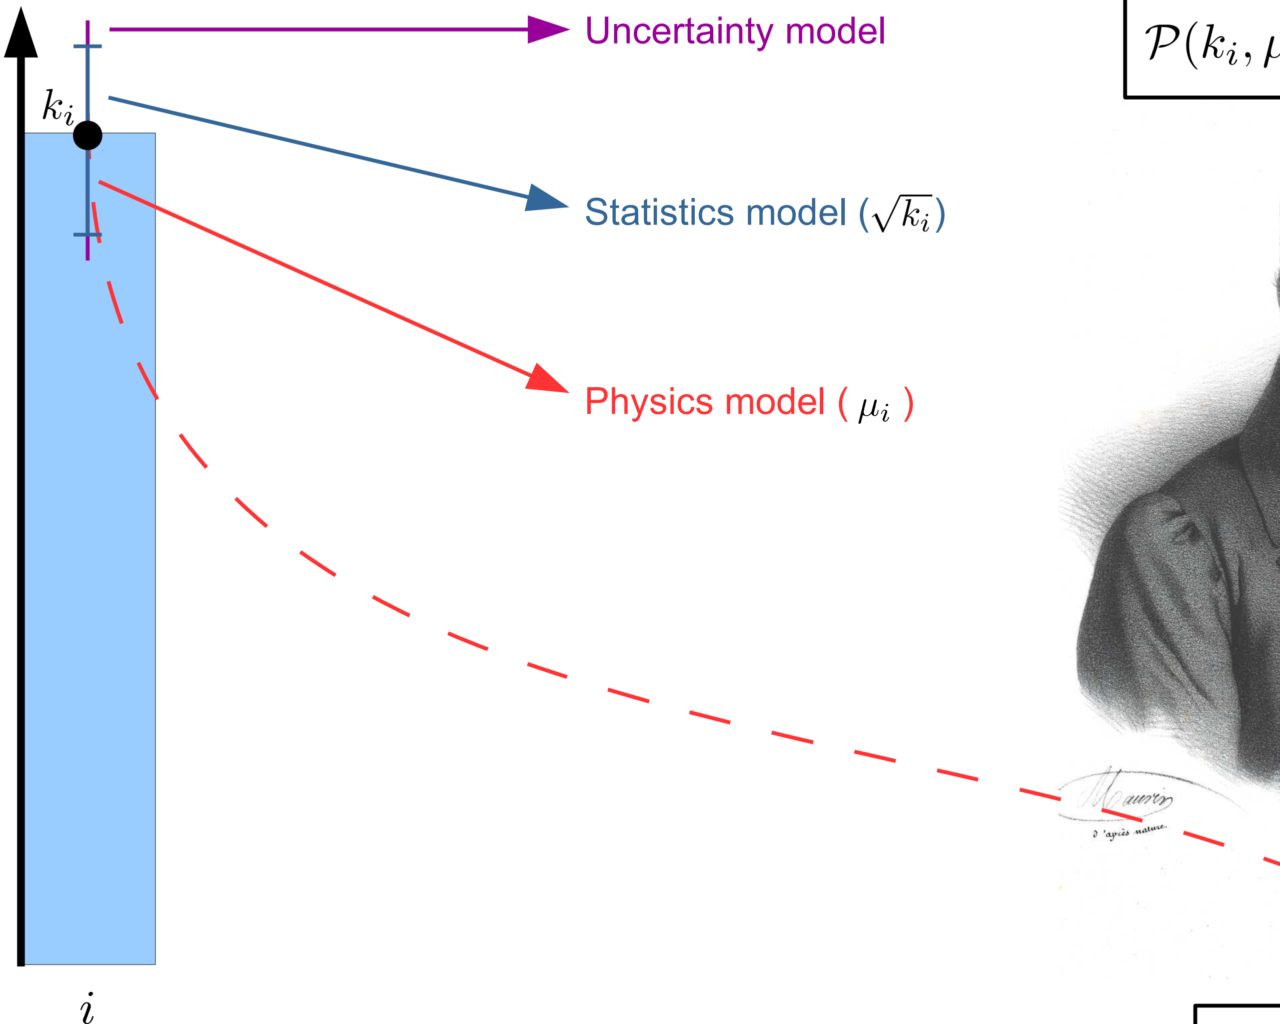

$$\mathcal{P}(k_i, \mu_i) = \frac{\mu_i^{k_i}}{k_i!} e^{-\mu_i}$$

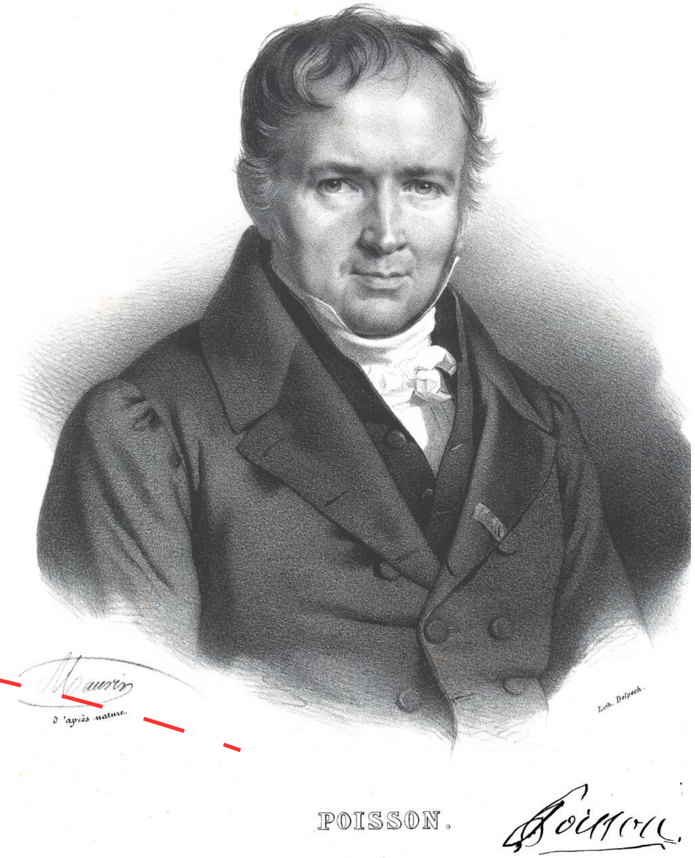

Siméon Denis Poisson
(21.07.1781 – 25.04.1840)

From one to many...

$$\prod_i \mathcal{P}(k_i, \mu_i) = \prod_i \frac{\mu_i^{k_i}}{k_i!} e^{-\mu_i}$$

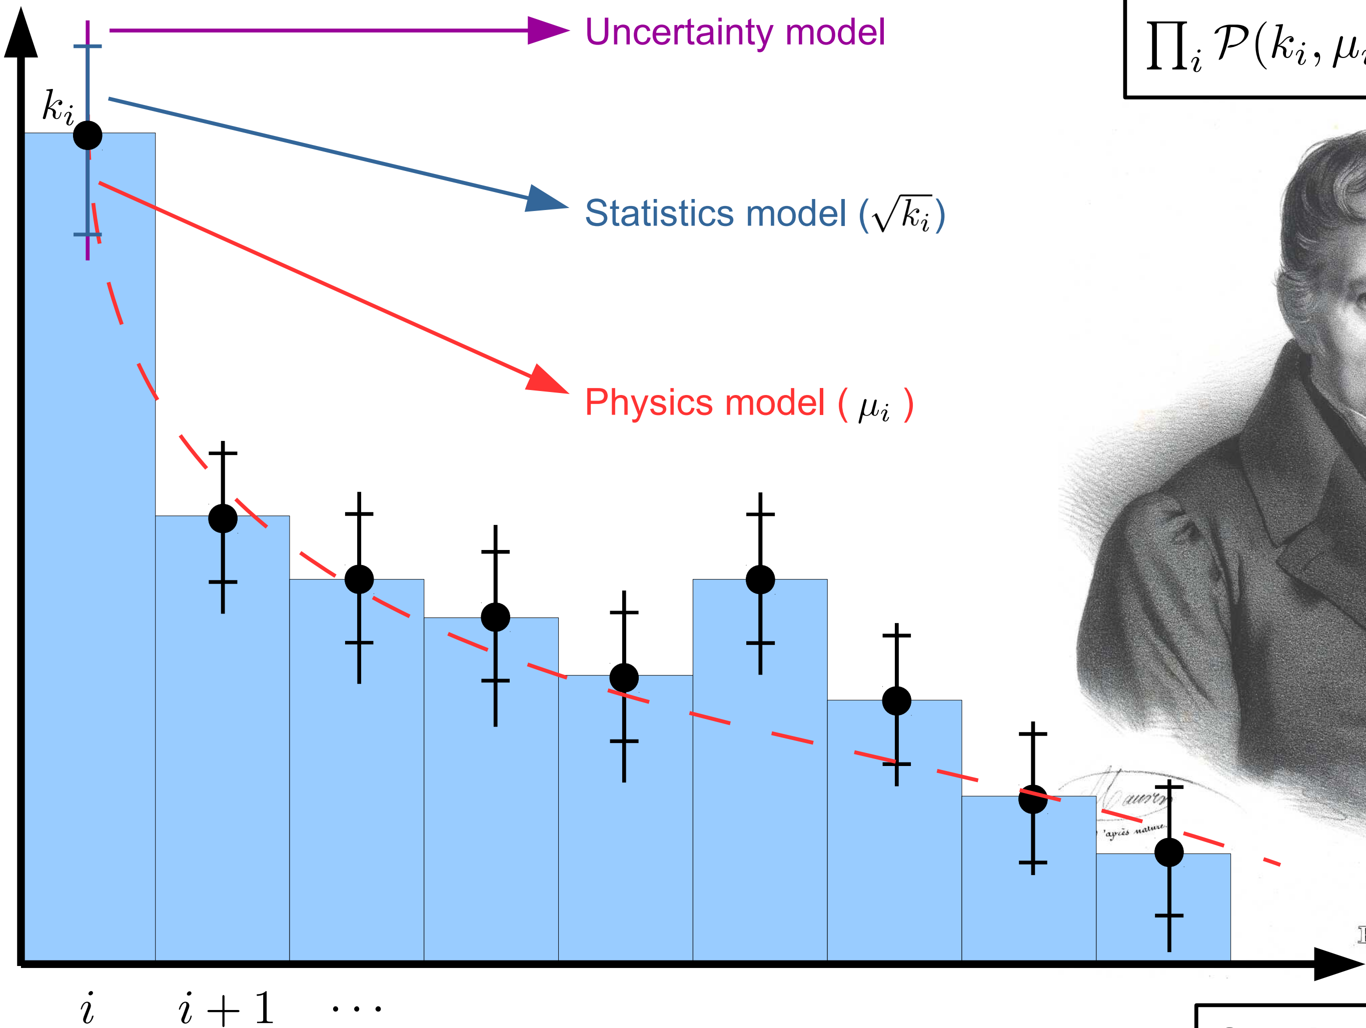

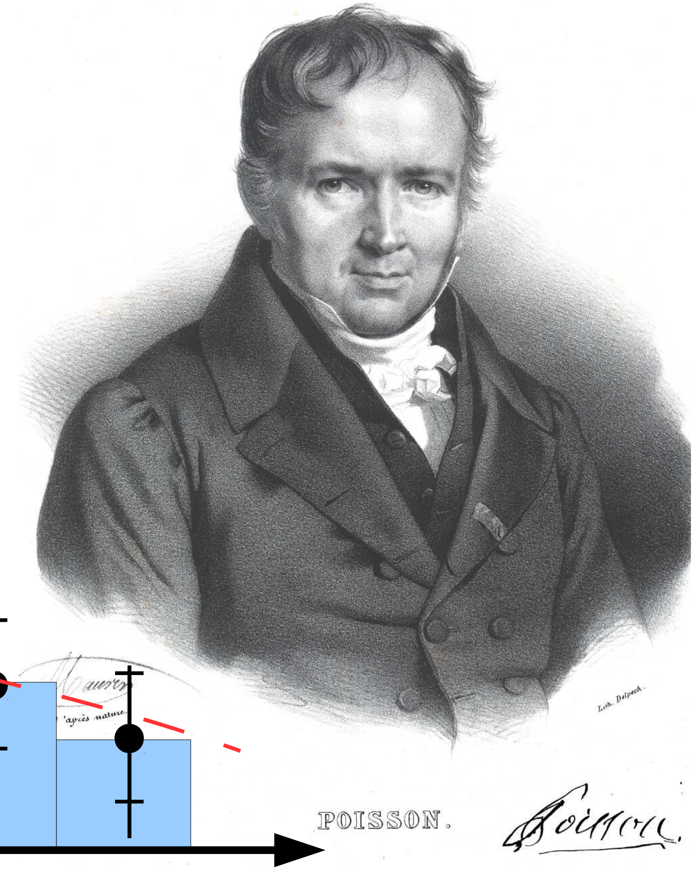

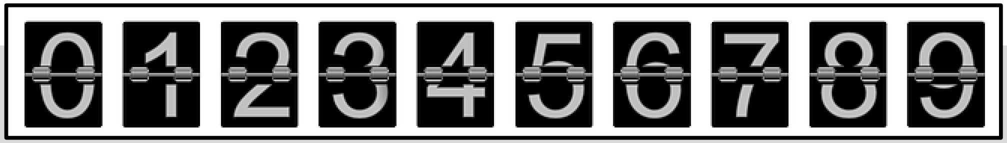

Siméon Denis Poisson
(21.07.1781 – 25.04.1840)

Model building (likelihood functions)

- Likeliness of a model to be true quantified by *likelihood function* $\mathcal{L}(\{k_i\}, \{\kappa_j\})$.

$$\prod_i \mathcal{P}(k_i, \mu_i) = \prod_i \frac{\mu_i^{k_i}}{k_i!} e^{-\mu_i}$$

$\{\kappa_j\}$ → model parameters.

$\{k_i\}$ → measured number of events (e.g. in bins i).

- Simple example:
  signal on top of known background in a binned histogram:

$$\mathcal{L}(\{k_i\}, \{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i, \mu_i(\kappa_j))}_{\text{Product of } pdf\text{s for each bin (Poisson).}}$$

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

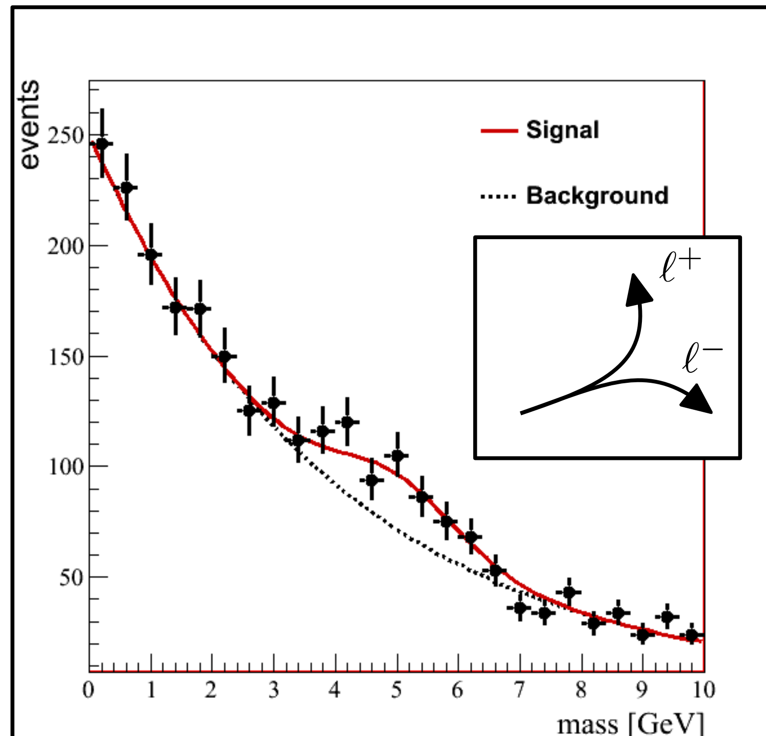

Model building (likelihood functions)

- Likeliness of a model to be true quantified by *likelihood function* $\mathcal{L}(\{k_i\},\{\kappa_j\})$.

$$\prod_i \mathcal{P}(k_i,\mu_i) = \prod_i \frac{\mu_i^{k_i}}{k_i!}e^{-\mu_i}$$

$\{\kappa_j\}$ → model parameters.

$\{k_i\}$ → measured number of events (e.g. in bins i).

- Simple example:
  signal on top of known background in a binned histogram:

$$\mathcal{L}(\{k_i\},\{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i,\mu_i(\kappa_j))}_{\text{Product of } pdf\text{s for each bin (Poisson).}}$$

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

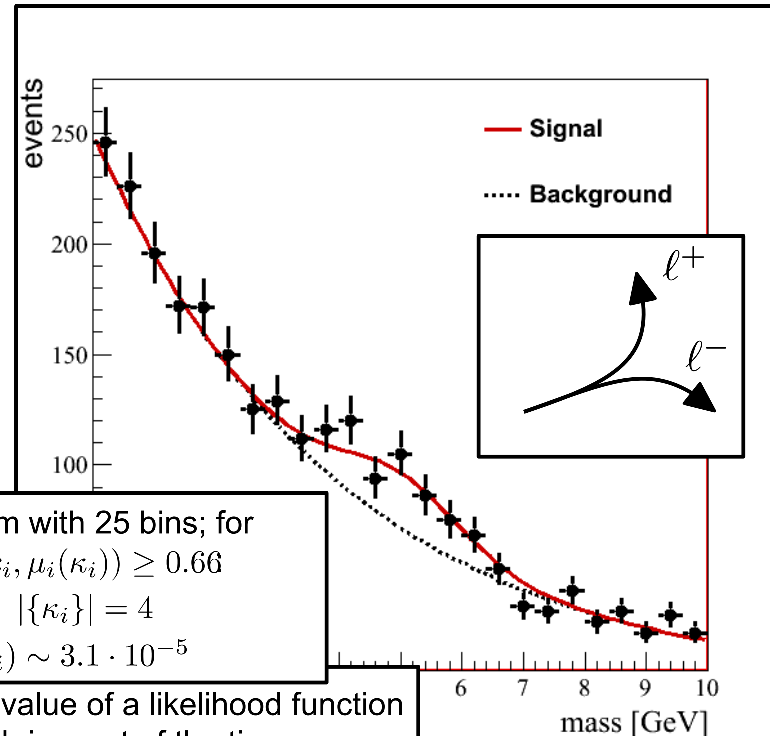

EX: histogram with 25 bins; for each bin $\mathcal{P}(k_i,\mu_i(\kappa_i)) \geq 0.66$

$|\{k_i\}| = 25 \qquad |\{\kappa_i\}| = 4$

$\prod \mathcal{P}(k_i,\mu_i(\kappa_i)) \sim 3.1 \cdot 10^{-5}$

NB: a value of a likelihood function as such is most of the time very close to zero, and w/o a reference in general w/o further meaning.

Distinguishing models (*likelihood ratio*)

Karlsruhe Institute of Technology

- Task of likelihood analyses:

  do not determine likelihood of an experimental outcome per se, but distinguish models (=hypotheses) and determine the one that explains the experimental outcome best.

- Fundamental lemma of Neyman-Pearson:

  when performing a test between two simple hypotheses H_1 and H_0 the *likelihood ratio test*, which rejects H_0 in favor of H_1 when

$$Q = \frac{\mathcal{L}_{H_1}(\{k_i\},\{\kappa_i\})}{\mathcal{L}_{H_0}(\{k_i\},\{\kappa_i\})} \leq \eta$$

$$\mathcal{P}(Q(\{k_i\},\{\kappa_i\}) \leq \eta | H_i) = \alpha$$

  is the most powerful test at significance level α for a threshold η.

- For $q = -2\ln Q$ this ratio turns into a difference ($\Delta\mathrm{NLL}$).

This is usually the *test statistic* of choice!

Maximum likelihood fit

Karlsruhe Institute of Technology

- Each likelihood ratio/function (with one or more parametric model part(s)) can be subject to a maximum likelihood fit (**NB:** negative log-likelihood finds its minimum where the log-likelihood is maximal...).

Minimization problem as known from school.

In our example e.g. four parameters κ_i.

(see next slides)

Parameters can be constraint or unconstraint.

- Simple example:
signal on top of known background in a binned histogram:

$$\mathcal{L}(\{k_i\},\{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i, \mu_i(\kappa_j))}$$

Product for each bin (Poisson).

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

The ATLAS+CMS Higgs couplings combined fit has $\mathcal{O}(4250)$ parameters and up to seven POI's.

The CMS Tracker Alignment problem has $\mathcal{O}(50'000)$ parameters and several thousand POI's.

Parameter(s) of interest (POI)

NB: this is a likelihood ratio on its own.

NB: I've also made the scan based on a likelihood ratio.

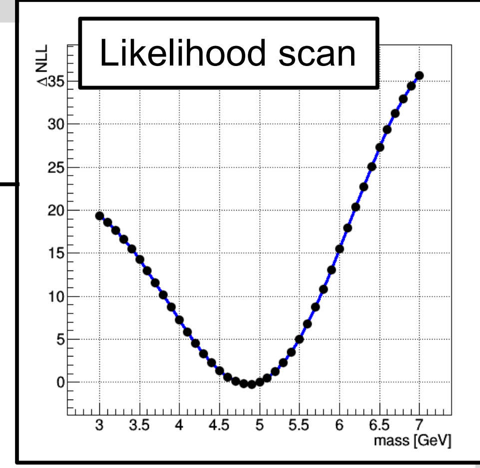

- In a maximum likelihood fit each case/problem defines its own *parameter(s) of interest* (POI's):
  - POI could be the mass (κ_3).

- Simple example:
  signal on top of known background in a binned histogram:

$$\mathcal{L}(\{k_i\},\{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i,\mu_i(\kappa_j))}_{\text{Product for each bin (Poisson).}}$$

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

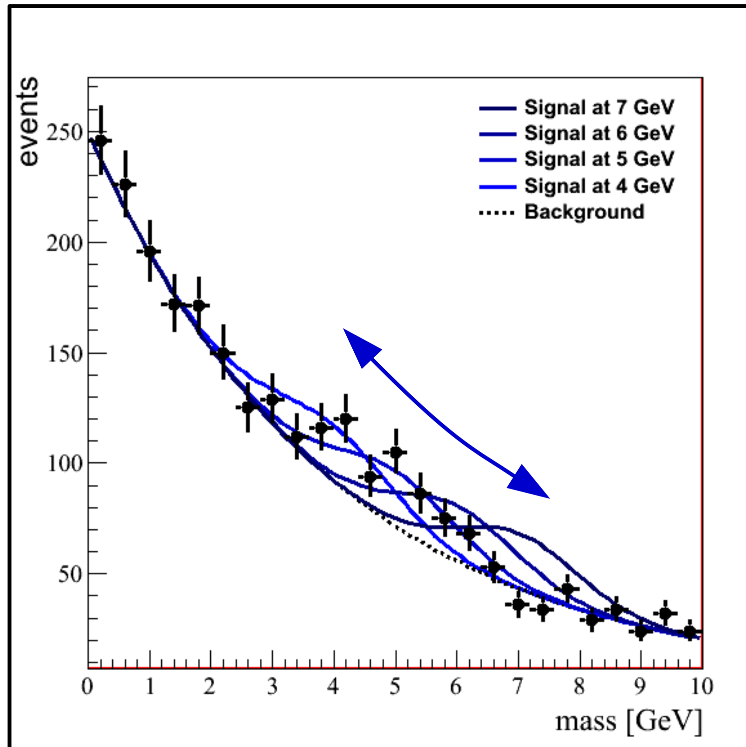

Parameter(s) of interest (POI)

NB: this is a likelihood ratio on its own

NB: I've also mad[e] a scan based on a likelihood ratio.

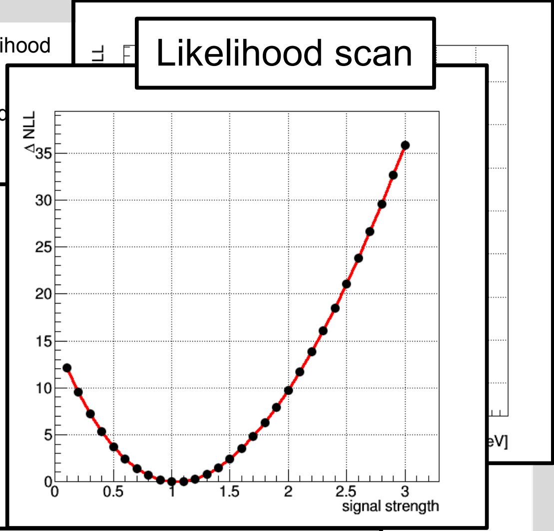

- In a maximum likelihood fit each case/problem defines its own *parameter(s) of interest* (POI's):
  - POI could be the mass (κ_3).
  - In our case POI usually is the signal strength (κ_2) (for a fixed value for κ_3).

- Simple example:
  signal on top of known background in a binned histogram:

$$\mathcal{L}(\{k_i\},\{\kappa_j\}) = \underbrace{\prod_i \mathcal{P}(k_i,\mu_i(\kappa_j))}$$

Product for each bin (Poisson).

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

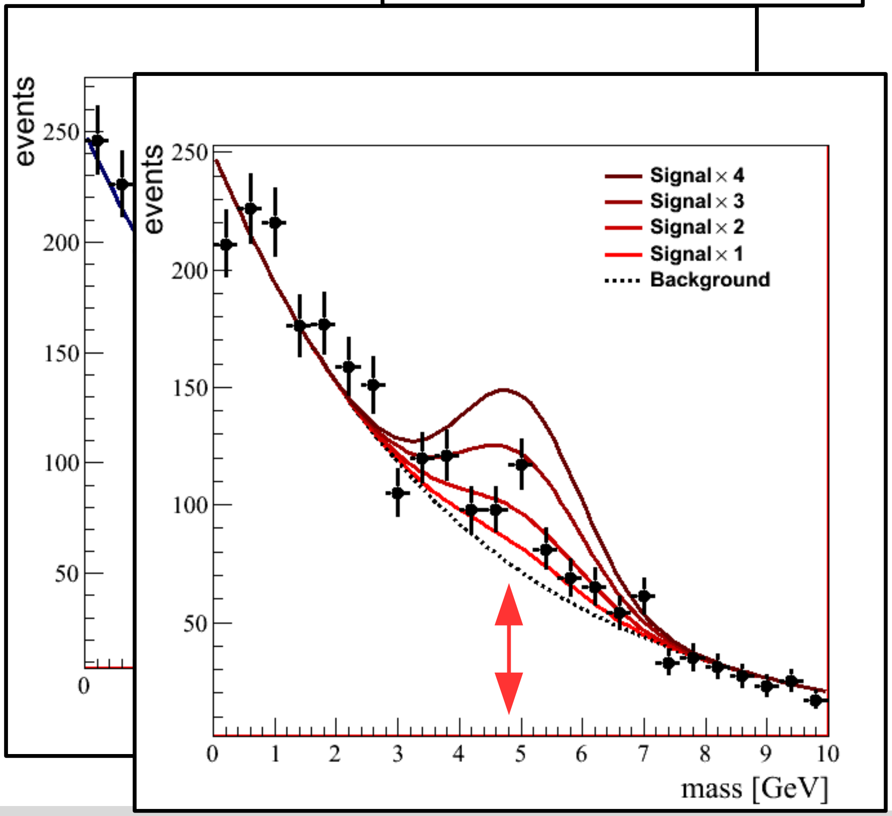

Incorporation of systematic uncertainties

- Systematic uncertainties are usually incorporated in form of *nuisance parameters*:
  - E.g. background normalization κ_0 not precisely known, but with uncertainty $\sigma(\kappa_0)$:

$$\mu_i(\kappa_j) = \kappa_0 \mathcal{P}'(x, 1, \sigma(\kappa_0)) \cdot e^{-\kappa_1 x_i} + \kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}$$

- $\sigma(\kappa_0)$ → uncertainty
- x → possible values of single "measurements" (integrated out)
- κ_0 → expected value/best knowledge

- Simple example:
  signal on top of known background in a binned histogram:

$$\mathcal{L}(\{k_i\}, \{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i, \mu_i(\kappa_j))}_{\text{Product for each bin (Poisson).}}$$

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

Incorporation of systematic uncertainties

- Systematic uncertainties are usually incorporated in form of *nuisance parameters*:
  - E.g. background normalization κ_0 not precisely known, but with uncertainty $\sigma(\kappa_0)$:

$$\mu_i(\kappa_j) = \kappa_0 \mathcal{P}'(x, 1, \sigma(\kappa_0)) \cdot e^{-\kappa_1 x_i} + \kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}$$

$\sigma(\kappa_0)$ → uncertainty

x, 1 → possible values of single "measurements" (integrated out)

κ_0 → expected value/best knowledge

Probability density function ($\mathcal{P}$)

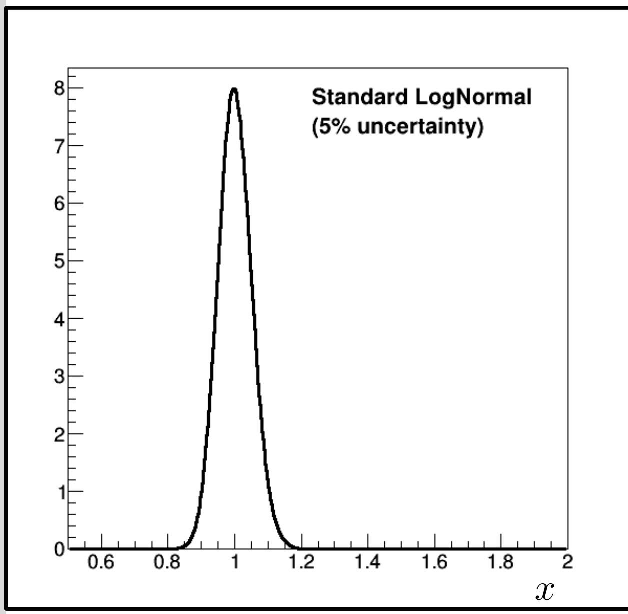

Effect on BG normalization

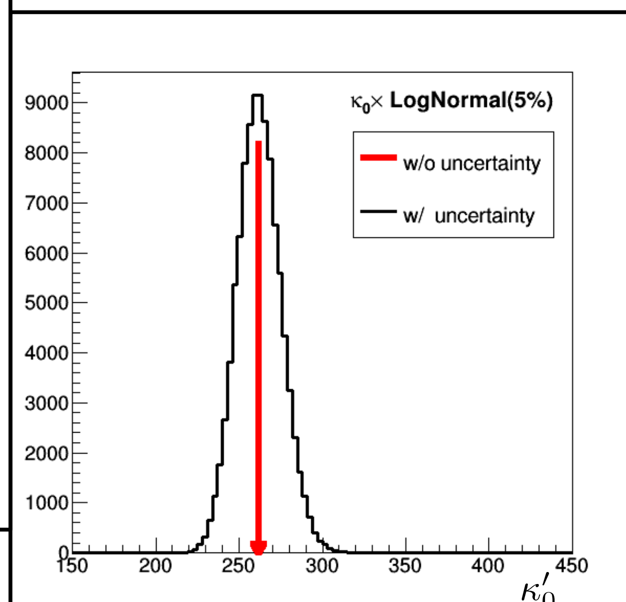

$$-\ln Q = -\ln\left(\frac{\mathcal{L}_{H_1}(\{k_i\},\{\kappa_i\})}{\mathcal{L}_{H_0}(\{k_i\},\{\kappa_i\})}\right)$$

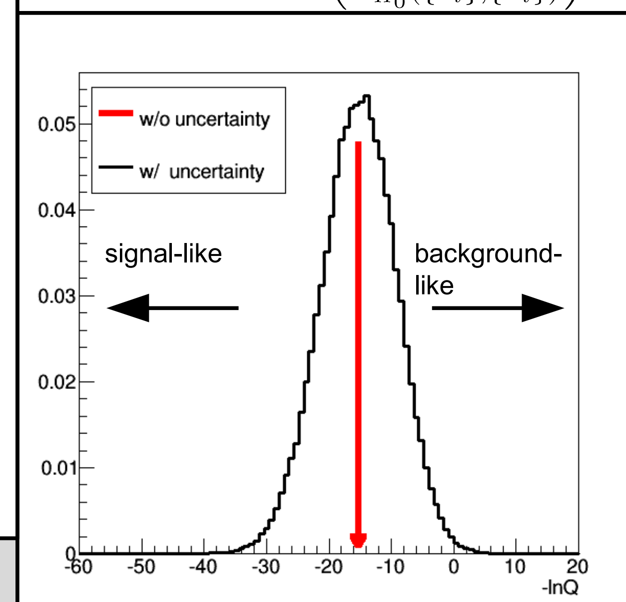

Example: test statistics (LEP ~2000)

- Test signal (H_1 , for fixed mass, m , and fixed signal strength, μ) vs. background-only (H_0).

pdf's for nuisance parameters modified according to Bayes theorem.

$$\mathcal{L}(n|\,b(\kappa_j)) = \prod_i \mathcal{P}(n_i|b_i(\kappa_j)) \times \prod_j \mathcal{C}(\kappa_j|\tilde{\kappa}_j)$$

$$\mathcal{L}(n|\,\mu s(\kappa_j) + b(\kappa_j)) = \prod_i \mathcal{P}(n_i|\mu s_i(\kappa_j) + b_i(\kappa_j)) \times \prod_j \mathcal{C}(\kappa_j|\tilde{\kappa}_j)$$

$$q_\mu = -2\ln\left(\frac{\mathcal{L}(\,n|\mu s+b)}{\mathcal{L}(\,n|b)}\right), \quad 0 \leq \mu$$

nuisance parameters $\tilde{\kappa}_j$ integrated out before evaluation of q_μ (→marginalization).

Example: test statistics (Tevatron ~2005)

- Test signal (H_1, for fixed mass, m, and fixed signal strength, μ) vs. background-only (H_0).

pdf's for nuisance parameters modified according to Bayes theorem.

$$\mathcal{L}(n|\,b(\kappa_j)) = \prod_i \mathcal{P}(n_i|b_i(\kappa_j)) \times \prod_j \mathcal{C}(\kappa_j|\tilde{\kappa}_j)$$

$$\mathcal{L}(n|\,\mu s(\kappa_j) + b(\kappa_j)) = \prod_i \mathcal{P}(n_i|\mu s_i(\kappa_j) + b_i(\kappa_j)) \times \prod_j \mathcal{C}(\kappa_j|\tilde{\kappa}_j)$$

$$q_\mu = -2\ln\left(\frac{\mathcal{L}(n|\mu s(\hat{\kappa}_\mu)+b(\hat{\kappa}_\mu))}{\mathcal{L}(n|b(\hat{\kappa}_{\mu=0}))}\right), \quad 0 \leq \mu$$

→ profiling.

nominator maximized for given μ before marginalization. Denominator for $\mu = 0$. Better estimates of nuisance parameters w/ reduced uncertainties.

Example: test statistics (LHC ~2010)

- Test signal (H_1, for fixed mass, m, and fixed signal strength, μ) vs. background-only (H_0).

profile likelihood (→ *Feldman-Cousins* test statistic).

$$\mathcal{L}(n|\,b(\kappa_j)) = \prod_i \mathcal{P}(n_i|b_i(\kappa_j))$$

$$\mathcal{L}(n|\,\mu s(\kappa_j) + b(\kappa_j)) = \prod_i \mathcal{P}(n_i|\mu s_i(\kappa_j) + b_i(\kappa_j))$$

$$q_\mu = -2\ln\left(\frac{\mathcal{L}(n|\mu s(\hat{\kappa}_\mu)+b(\hat{\kappa}_\mu))}{\mathcal{L}(n|\hat{\mu} s(\hat{\kappa}_{\hat{\mu}})+b(\hat{\kappa}_{\hat{\mu}}))}\right),\quad 0 \leq \hat{\mu} \leq \mu$$

→ profiling.

nominator maximized for given μ before marginalization. For the denominator a global maximum is searched for at $\hat{\mu}$. In addition allows use of asymptotic formulas (→ no more toys needed!(*)).

(*) will not be discussed further here.

Meaning and interpretation of the test statistic

Karlsruhe Institute of Technology

- From the evaluation of the test statistic on data always obtain a plain value q_{obs} (in our discussion: $q_{obs} < 0$ – signal-like; $q_{obs} > 0$ – background-like).
- → True outcome of the experiment (nuisance parameters estimated to best knowledge, no uncertainties involved here)!

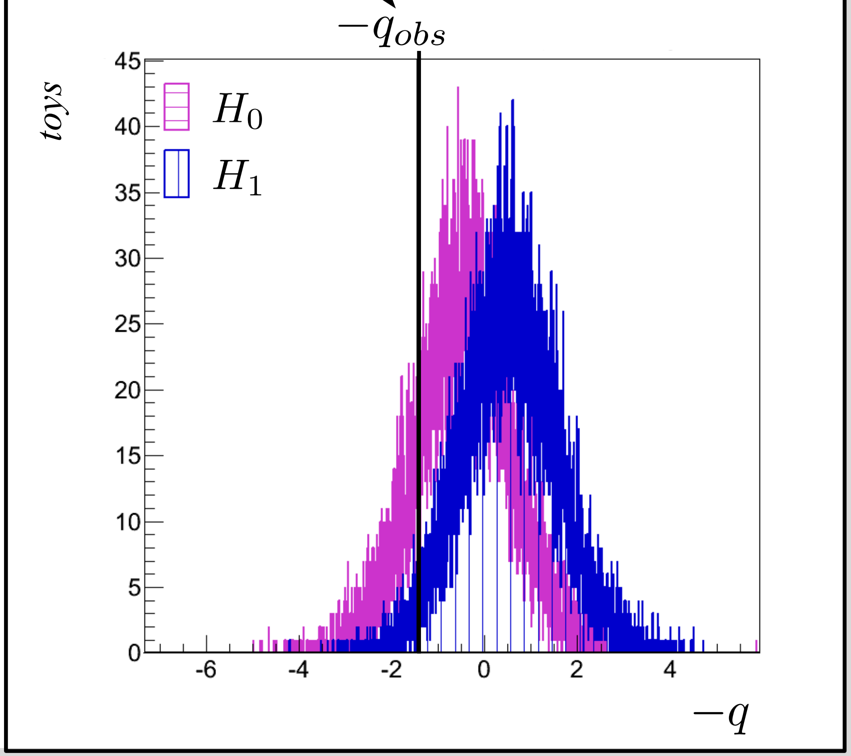

NB: real life examples (→ HIG-13-021).

How to produce a toy experiment:

$$\mathcal{L}(\{k_i\},\{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i, \mu_i(\kappa_j))}_{\text{Product for each bin (Poisson).}}$$

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

Meaning and interpretation of the test statistic

- How compatible is q_{obs} with H_0 or H_1? For this evaluate the test statistic on large number of toy experiments based on H_0 or H_1 .

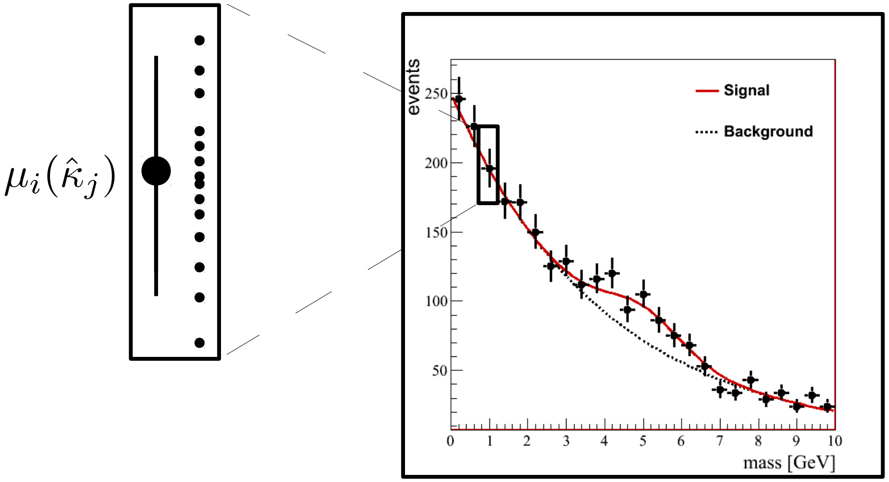

$$\mathcal{L}(\{k_i\},\{\kappa_j\}) = \prod_i \underbrace{\mathcal{P}(k_i, \mu_i(\kappa_j))}$$

Product for each bin (Poisson).

$$\mu_i(\kappa_j) = \underbrace{\kappa_0 \cdot e^{-\kappa_1 x_i}}_{\text{background}} + \underbrace{\kappa_2 \cdot e^{-(\kappa_3 - x_i)^2}}_{\text{signal}}$$

- Determine *toy dataset*.
- Determine *toy values* for all uncertainties.
- Determine value of $-q$ for toy.
- Proceed as often as possible; do the same for H_0 & H_1.

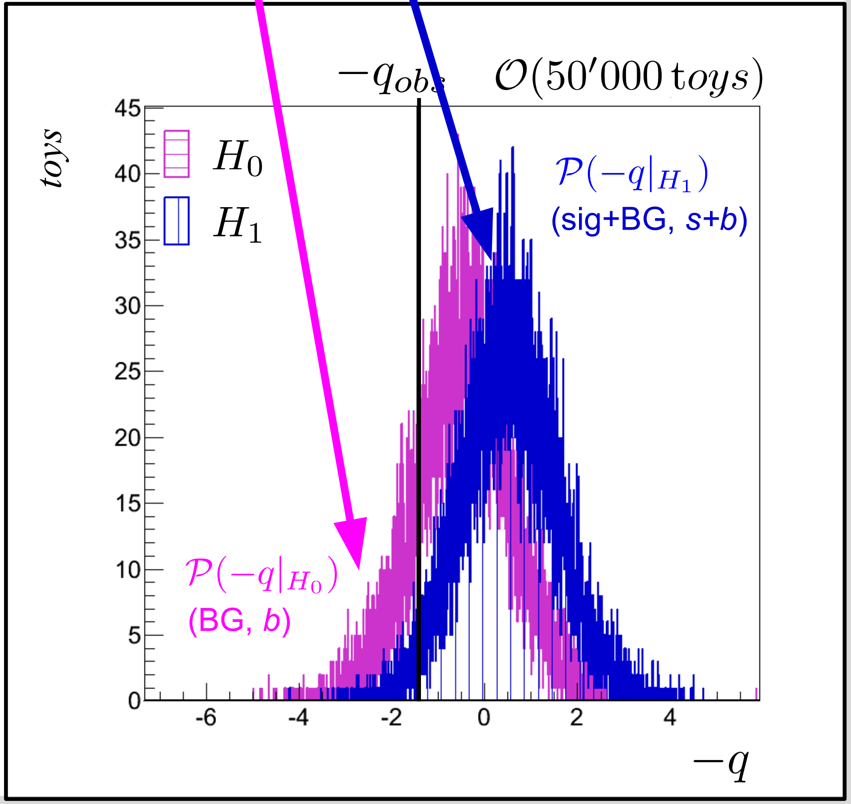

Significance levels/confidence levels (*CL*)

- The association to one or the other hypothesis can now be performed up to a given confidence level α.

Questioning H_0:

Probability to obtain values of q, which are at least as signal-like as q_{obs}. If p-value is small H_0 can be excluded.

$$(1-CL_b) = \int_{-q_{\rm obs}}^{+\infty} \mathcal{P}_b \qquad (p\text{-value})$$

$$CL_{s+b} = \int_{-\infty}^{-q_{\rm obs}} \mathcal{P}_{s+b} \qquad (CL_{s+b}\ \text{confidence})$$

$$CL_b = \int_{-\infty}^{-q_{\rm obs}} \mathcal{P}_b \qquad (CL_b\ \text{confidence})$$

$$CL_s = \frac{CL_{s+b}}{CL_b} \qquad (CL_s\ \text{confidence})$$

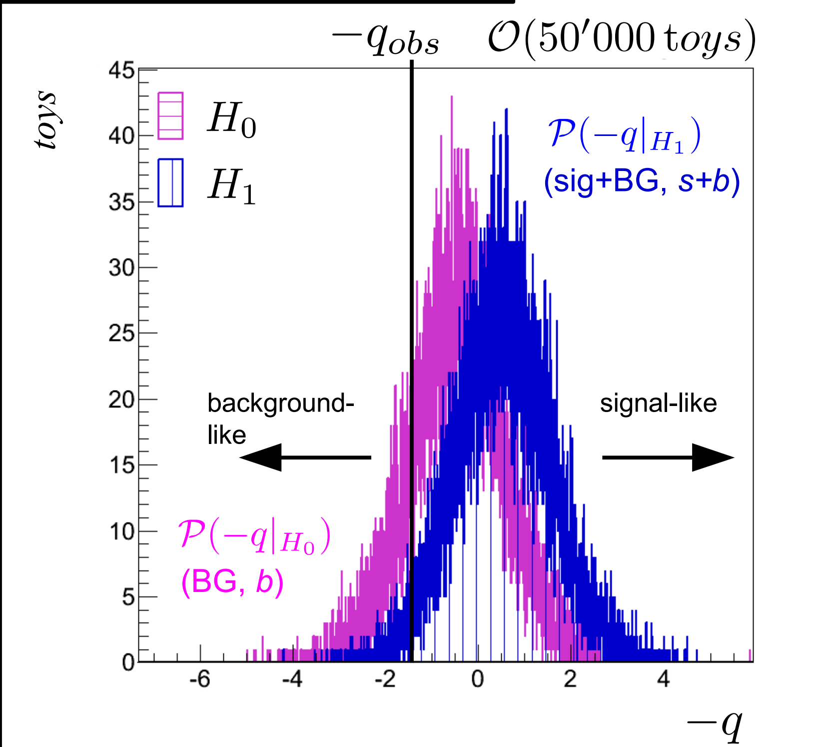

Example: test statistics (LEP)

$$q_\mu = -2\ln\left(\frac{\mathcal{L}(n|\mu s+b)}{\mathcal{L}(n|b)}\right),\quad 0 \le \mu$$

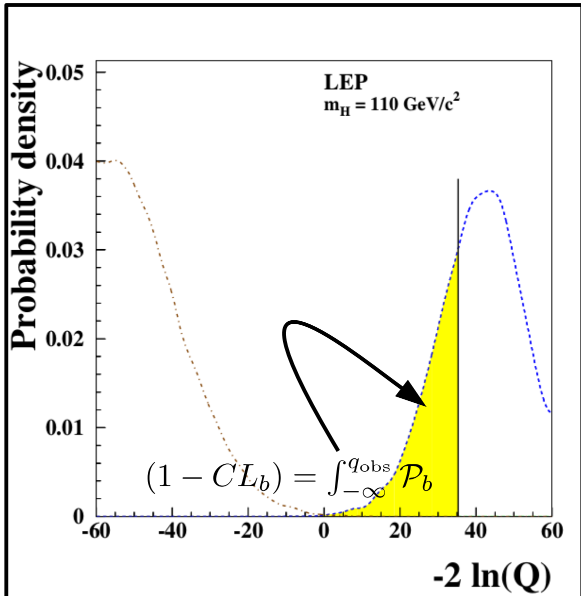

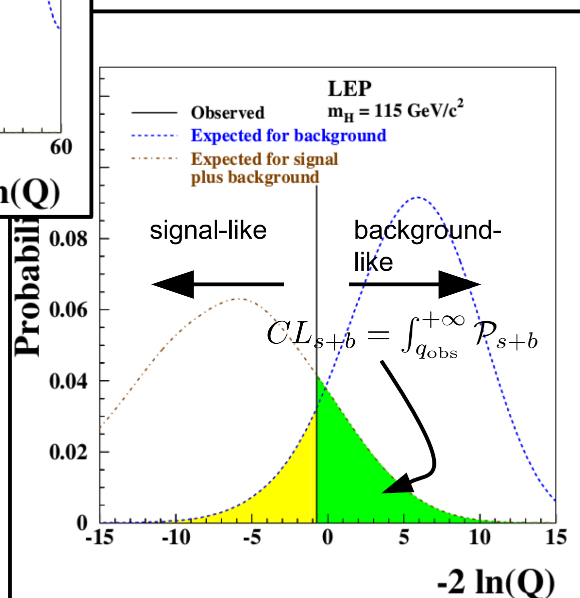

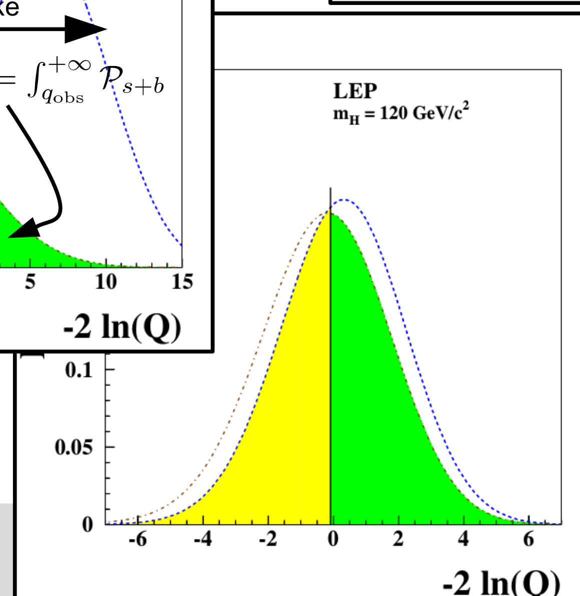

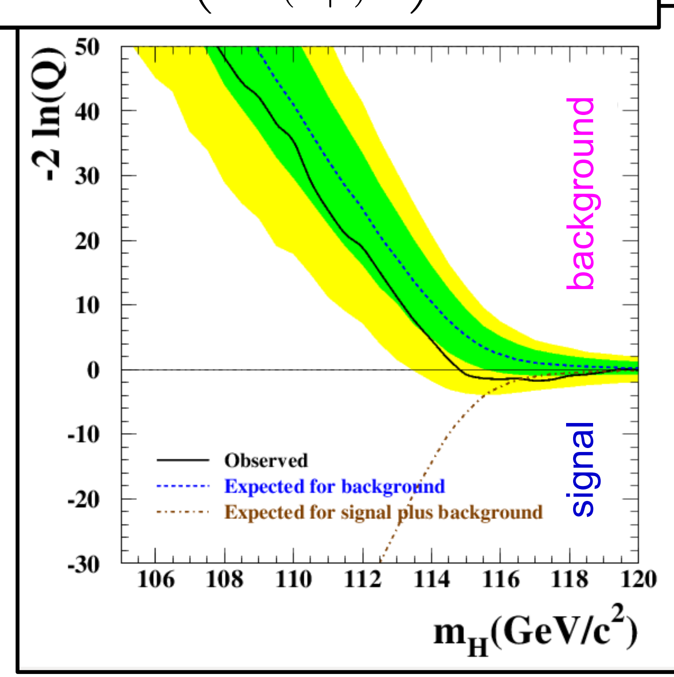

Example: *p*-value (LHC)

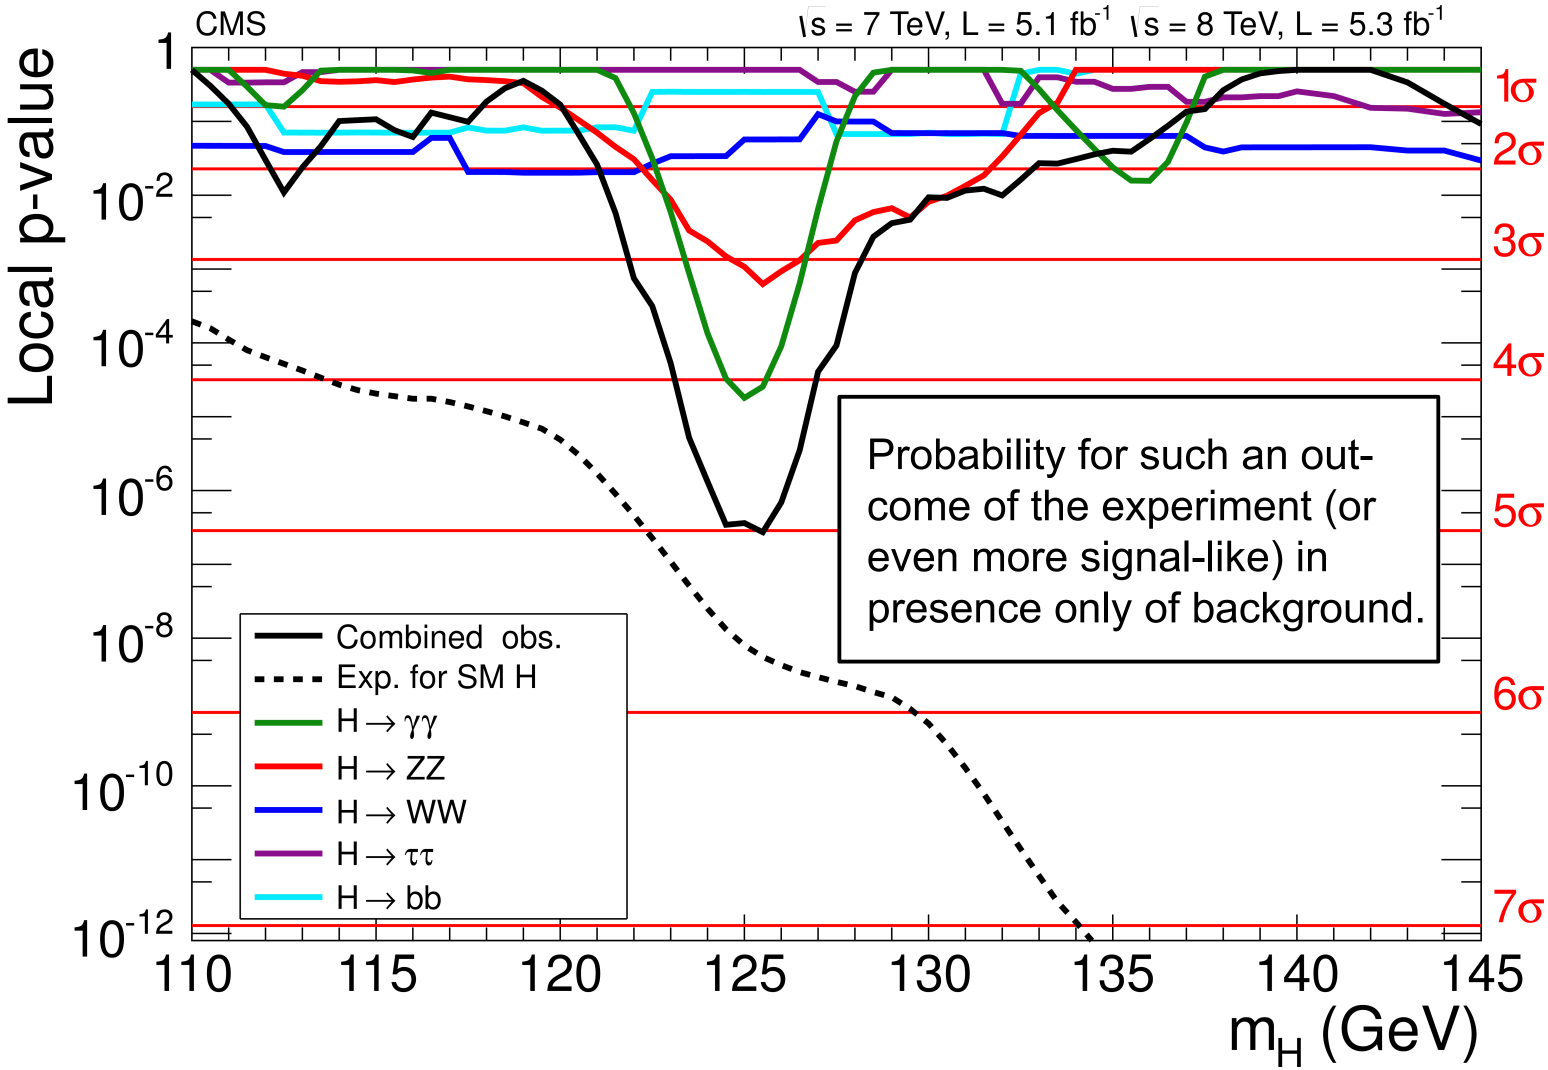

Example: saturated model

- Example of a likelihood ratio:

$$q_\lambda = -2\ln\left(\frac{\mathcal{L}(\mathrm{data}|_{\mathrm{test}})}{\mathcal{L}(\mathrm{data}|_{\mathrm{saturated}})}\right)$$

Model to be tested.

Model w/ as many parameters, λ_j, as measurements.

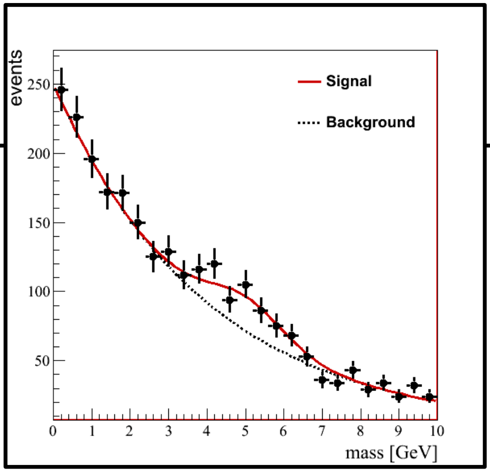

e.g. one shape for each bin.

- Special case: (i) histogram; (ii) no further nuisance parameters; (iii) uncertainties normal distributed:

$$\mathcal{L}(\mathrm{data}|_{\mathrm{test}}) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i}} e^{-(d_i-\lambda_i)^2/2\sigma_i}$$

$$\mathcal{L}(\mathrm{data}|_{\mathrm{saturated}}) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i}}$$

$$q_\lambda = -2\ln\left(\frac{\mathcal{L}(\mathrm{data}|_{\mathrm{test}})}{\mathcal{L}(\mathrm{data}|_{\mathrm{saturated}})}\right) = \sum_i \frac{(d_i-\lambda_i)^2}{\sigma_i}$$

Generalization of the χ^2 test.

- General case: (i) many histograms; (ii) many nuisance parameters:

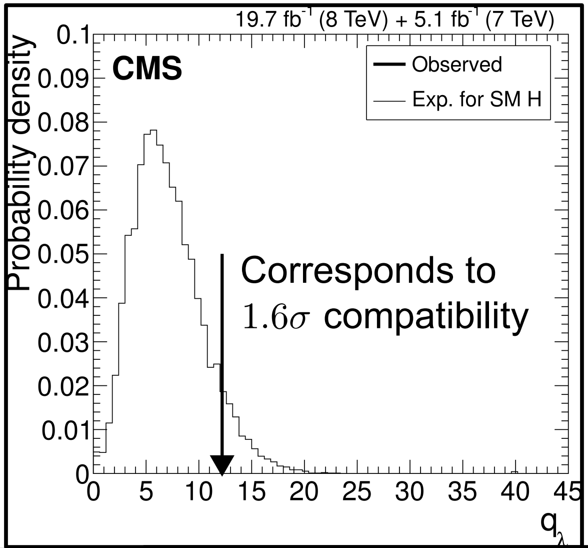

CL of interest: $\int_{q_{\mathrm{obs}}}^{+\infty} \mathcal{P}_{\mathrm{test}}$

Excluding parameters

- Questioning H_1: to be conservative choose probability α that q is more BG-like than q_{obs} low.
- $\mathcal{P}(-q|_{H_1})$ usually depends on POI: $q = -2\ln\left(\frac{\mathcal{L}(\mathrm{obs})|H_1}{\mathcal{L}(\mathrm{obs})|H_0}\right)$

varies

fixed

fixed

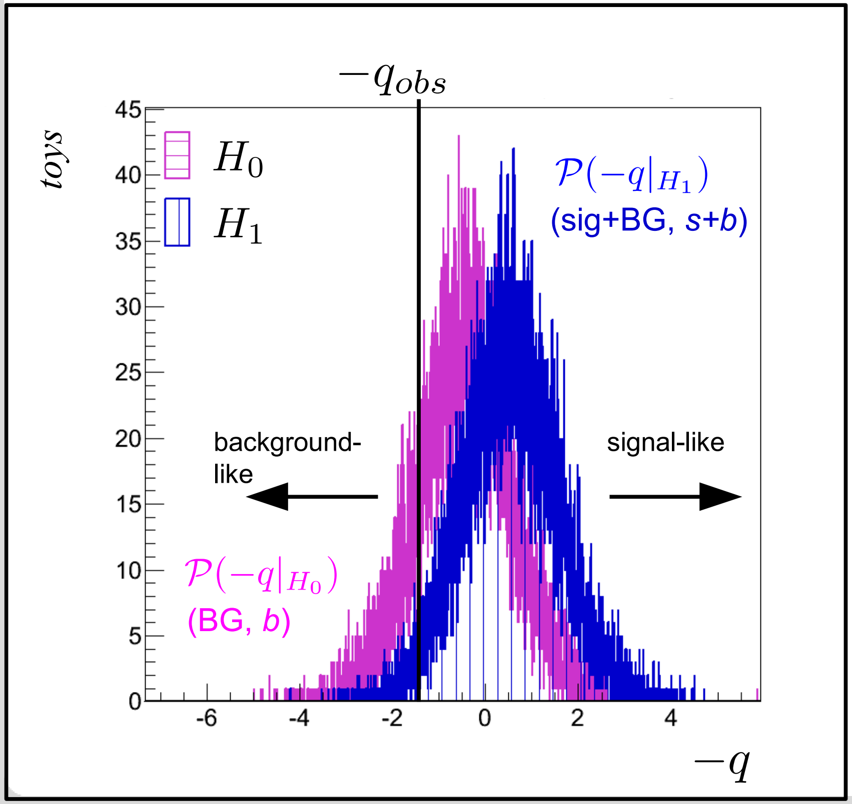

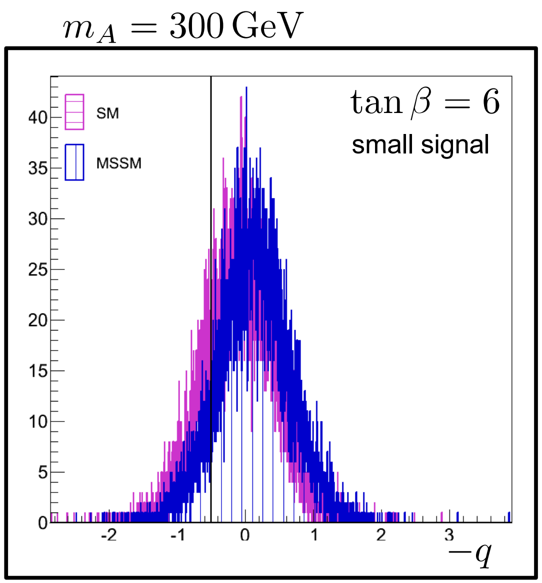

Excluding parameters

- Questioning H_1: to be conservative choose probability α that q is more BG-like than q_{obs} low.
- $\mathcal{P}(-q|_{H_1})$ usually depends on POI: $q = -2\ln\left(\frac{\mathcal{L}(\text{obs})|H_1}{\mathcal{L}(\text{obs})|H_0}\right)$

varies

fixed

fixed

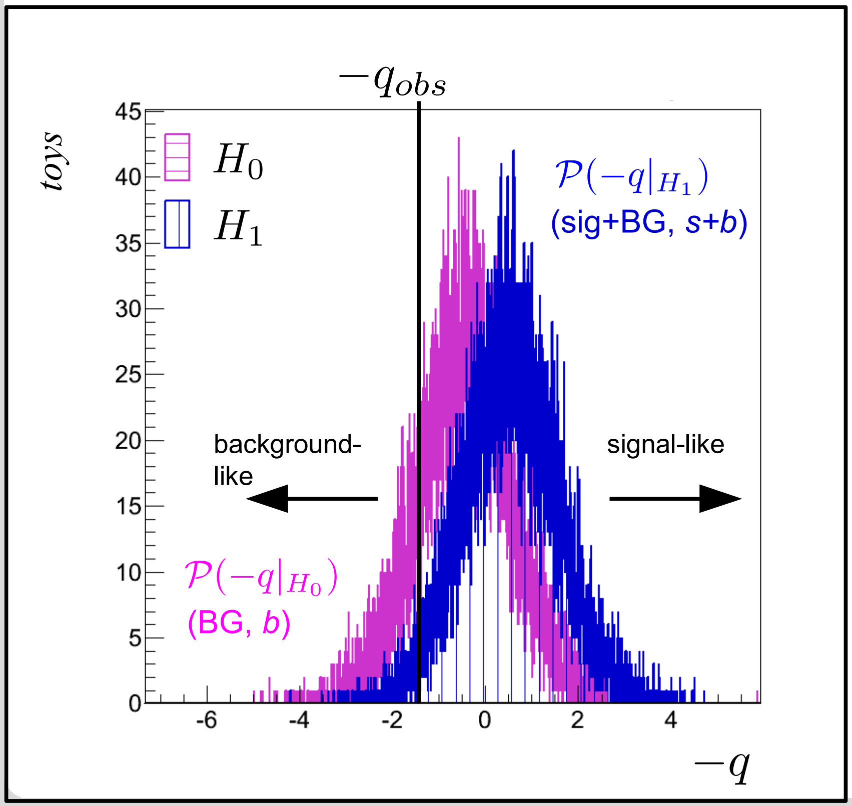

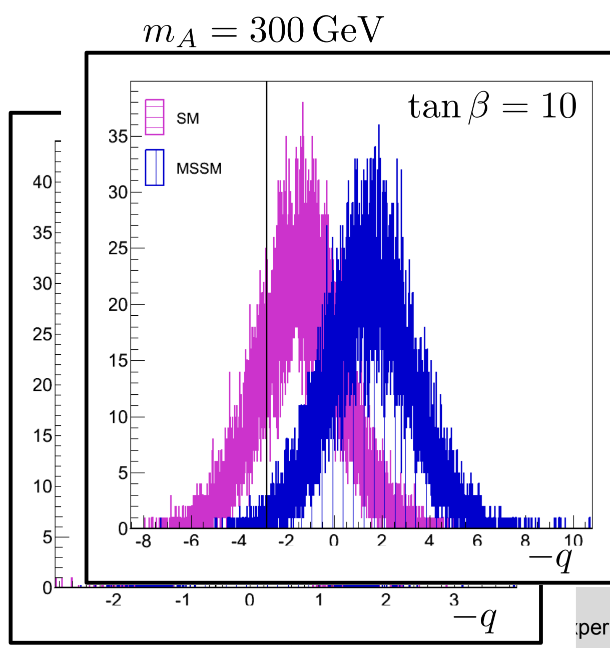

Excluding parameters

- Questioning H_1: to be conservative choose probability α that q is more BG-like than q_{obs} low.
- $\mathcal{P}(-q|_{H_1})$ usually depends on POI: $q = -2\ln\left(\frac{\mathcal{L}(\text{obs})|H_1}{\mathcal{L}(\text{obs})|H_0}\right)$

H_1 → varies; obs → fixed; H_0 → fixed

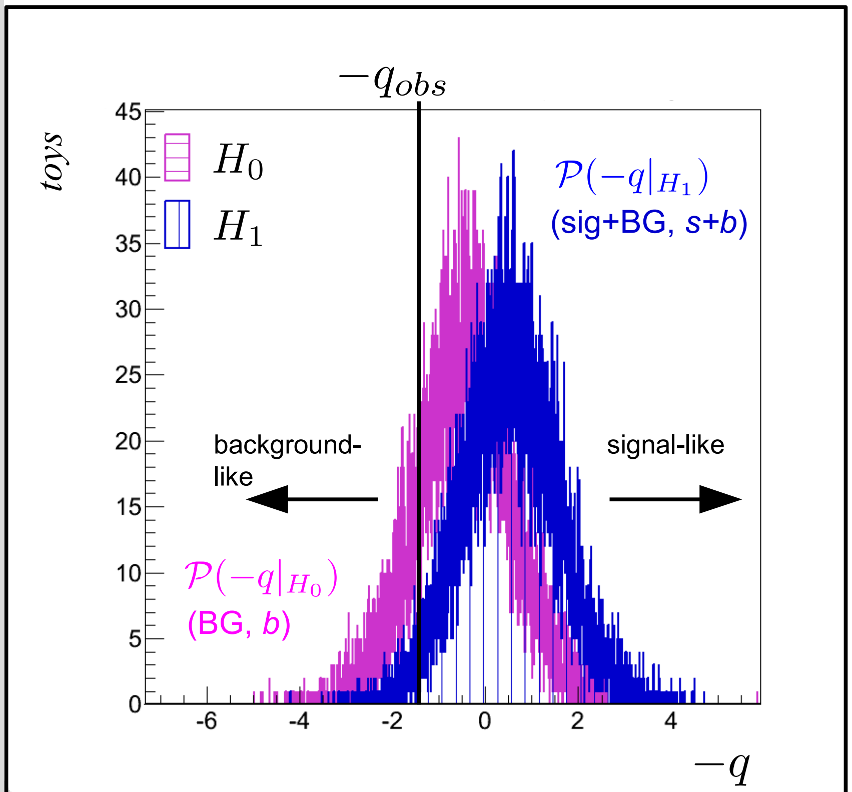

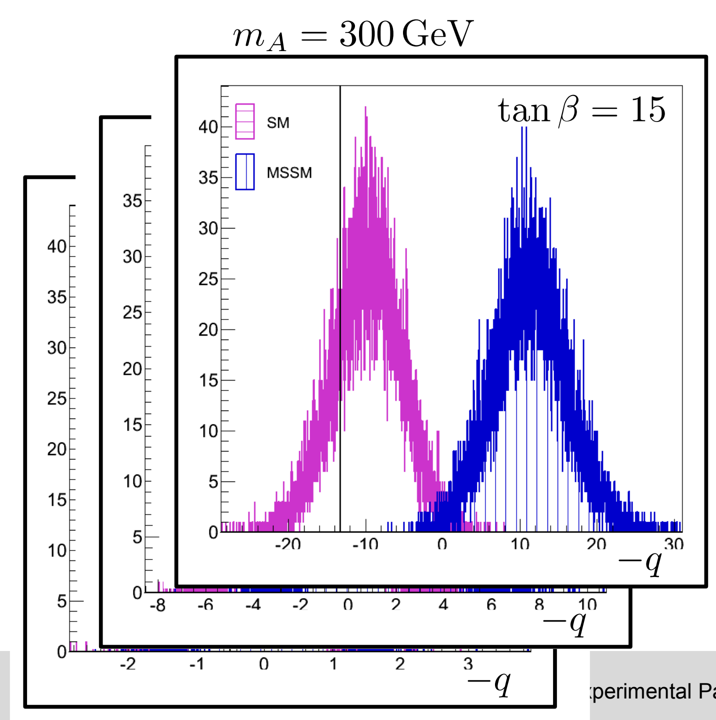

Excluding parameters

- Questioning H_1: to be conservative choose probability α that q is more BG-like than q_{obs} low.
- $\mathcal{P}(-q|_{H_1})$ usually depends on POI: $q = -2\ln\left(\frac{\mathcal{L}(\text{obs})|H_1}{\mathcal{L}(\text{obs})|H_0}\right)$

varies

fixed

fixed

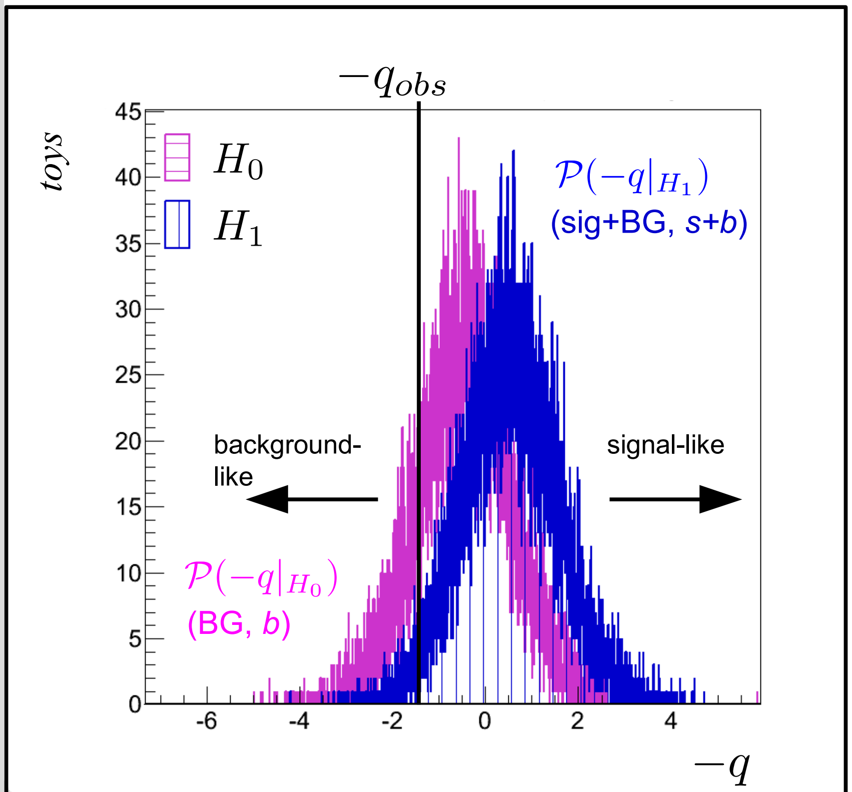

$m_A = 300\,\text{GeV}$

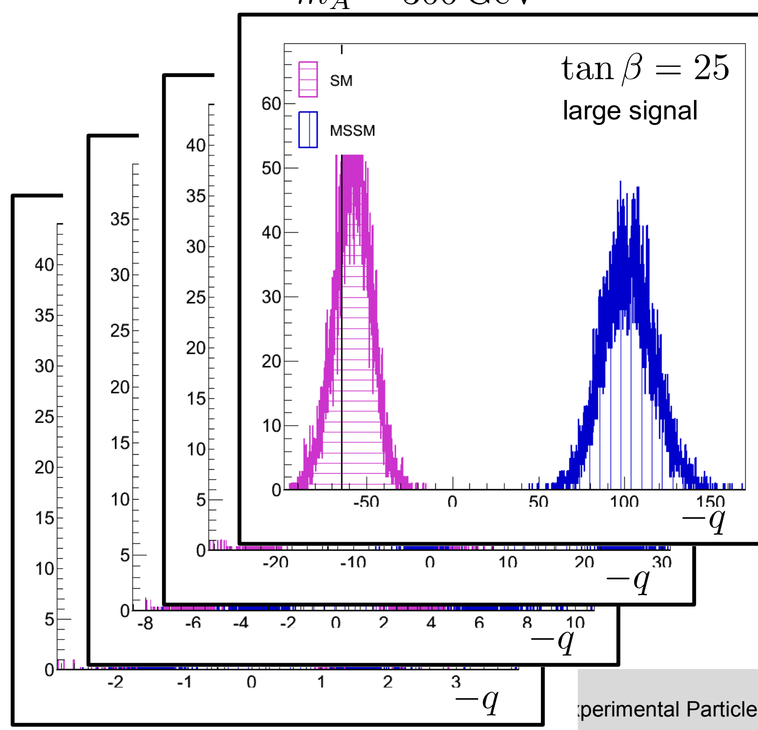

Observed exclusion

- Questioning H_1: to be conservative choose probability α that q is more BG-like than q_{obs} low.
- Traditionally we determine 95% CL exclusions on the POI ($\alpha = 0$).

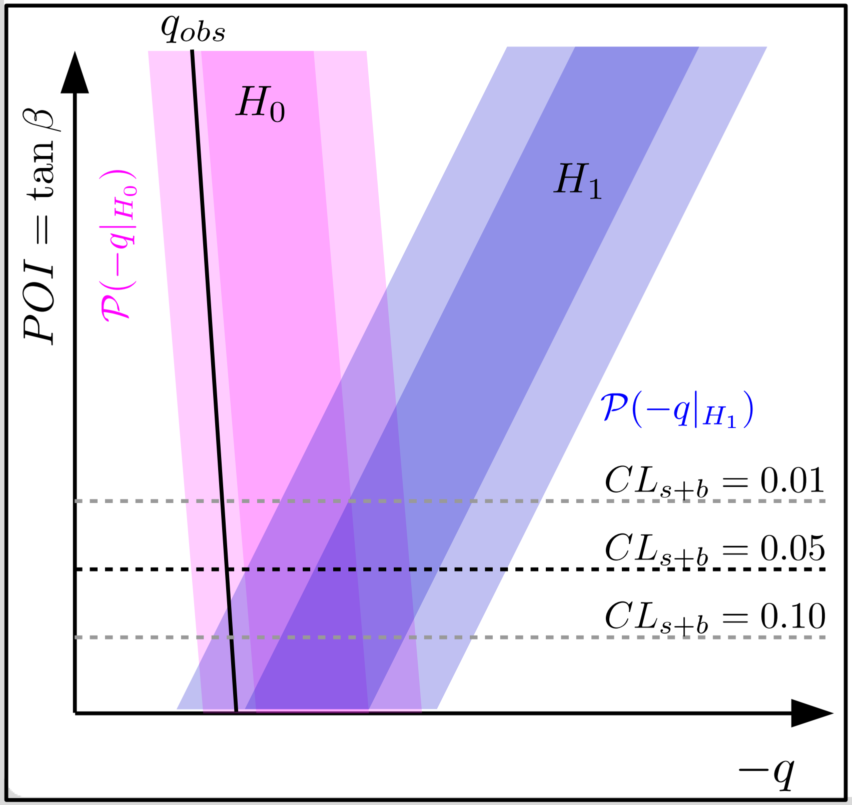

- $\mathcal{P}(-q|_{H_0})$ and $\mathcal{P}(-q|_{H_1})$ move apart from each other with increasing POI.
- The more separated $\mathcal{P}(-q|_{H_0})$ and $\mathcal{P}(-q|_{H_1})$ are the clearer H_0 and H_1 can be distinguish-ed.
- Identify value of POI for which:

  $$CL_{s+b} = \int_{-\infty}^{-q_{\mathrm{obs}}} \mathcal{P}_{s+b} = 0.05$$

  for this value q would have been more signal-like than q_{obs} with 95% probability.
- There is still a 5% chance to exclude by mistake.

Expected exclusion

- To obtain expected limit mimic calculation of observed; base it on toy datasets.
- Use fact that $\mathcal{P}(-q|_{H_0})$ and $\mathcal{P}(-q|_{H_1})$ do not depend on toys (i.e. schematic plot on the left does not change).

Throw toys under H_0 hypothesis; determine distribution of 95% CL limits on POI:

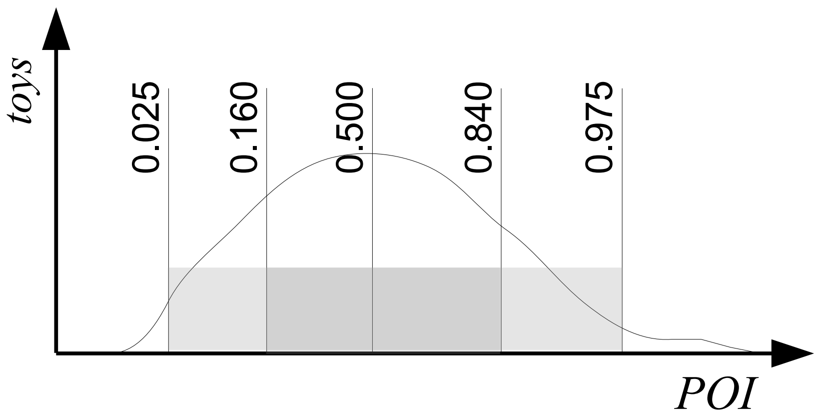

Obtain quantiles for expected exclusion from this distribution (expected limit = median).

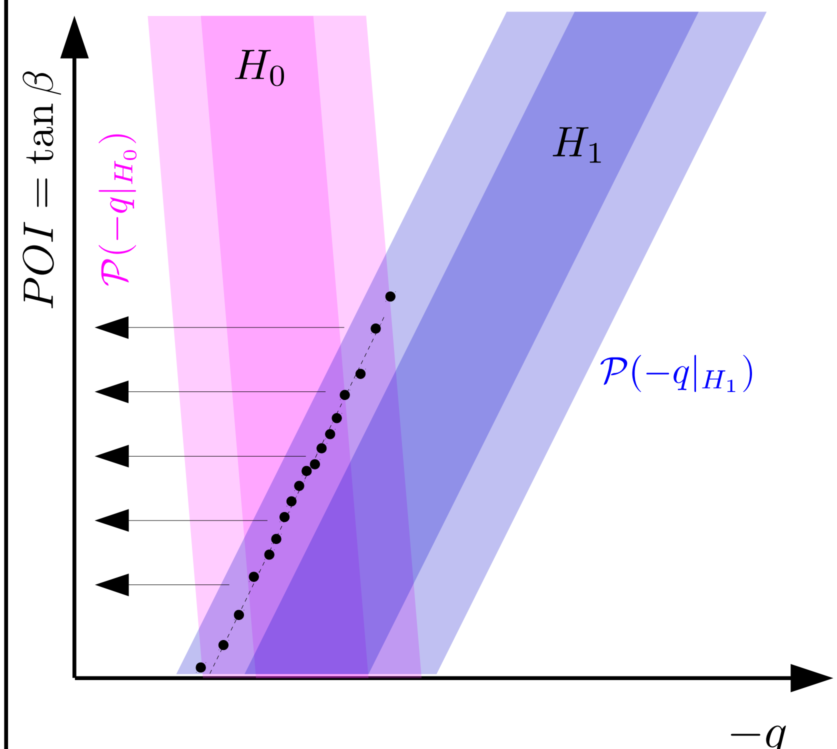

Interpretation issues (increasing pathology)

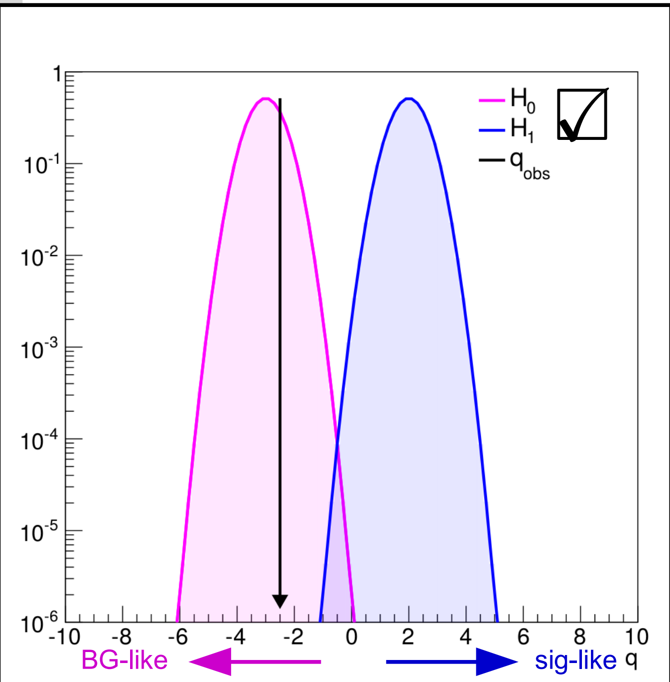

- q_{obs} compatible with BG hypothesis.
- q_{obs} incompatible with signal hypothesis.

- Signal and BG hypothesis cannot be distinguished.
- Should this outcome lead to an exclusion of the signal hypothesis?

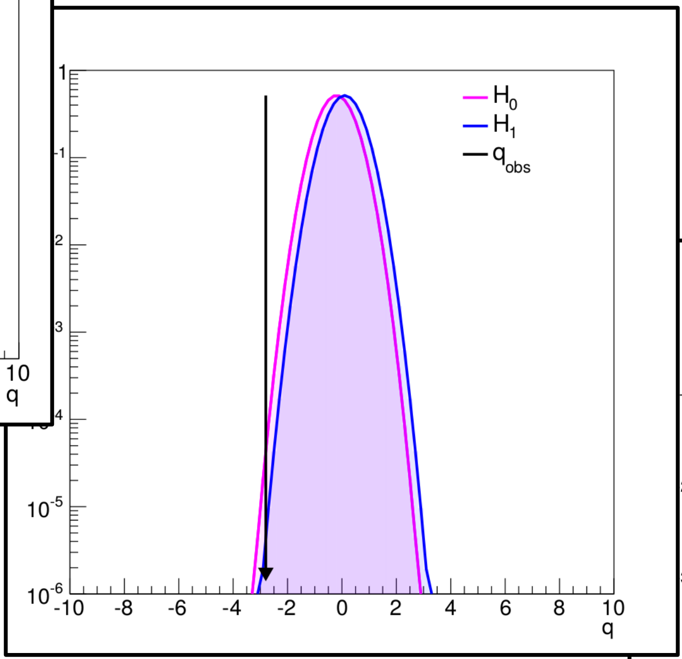

- q_{obs} incompatible both with signal and BG hypothesis.
- Should this outcome lead to an exclusion of the signal hypothesis?

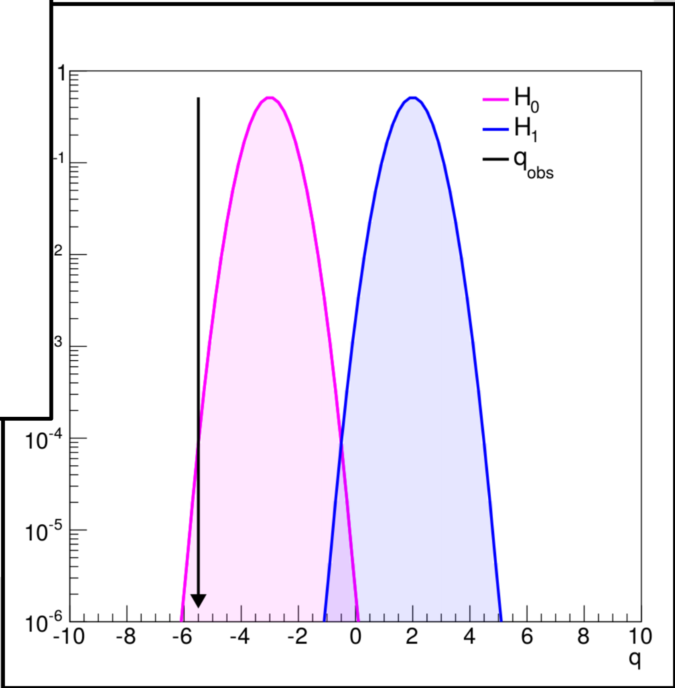

Modified frequentist exclusion method (CLs)

- In particle physics we set more conservative limits, following the CLs method:

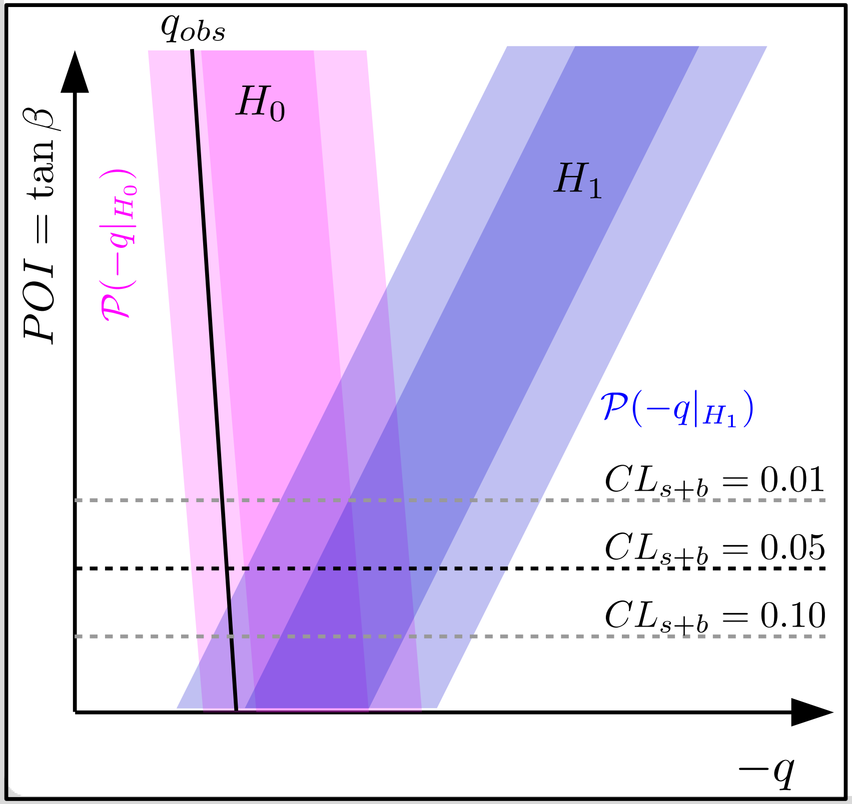

- $CL_{s+b} = \int_{-\infty}^{-q_{obs}} \mathcal{P}_{s+b}$

  $CL_b \quad = \int_{-\infty}^{-q_{obs}} \mathcal{P}_b$

- Identify value of POI for which:

  $CL_s = \frac{CL_{s+b}}{CL_b} = 0.05$

- If H_0 and H_1 are well separated (and H_0 gives a valid description of the experiment):
  $CL_s \rightarrow CL_{s+b}$

- If H_0 and H_1 become indistinguishable:
  $CL_{s+b} < CL_s \rightarrow 1$

Example: SM Higgs search (LHC)

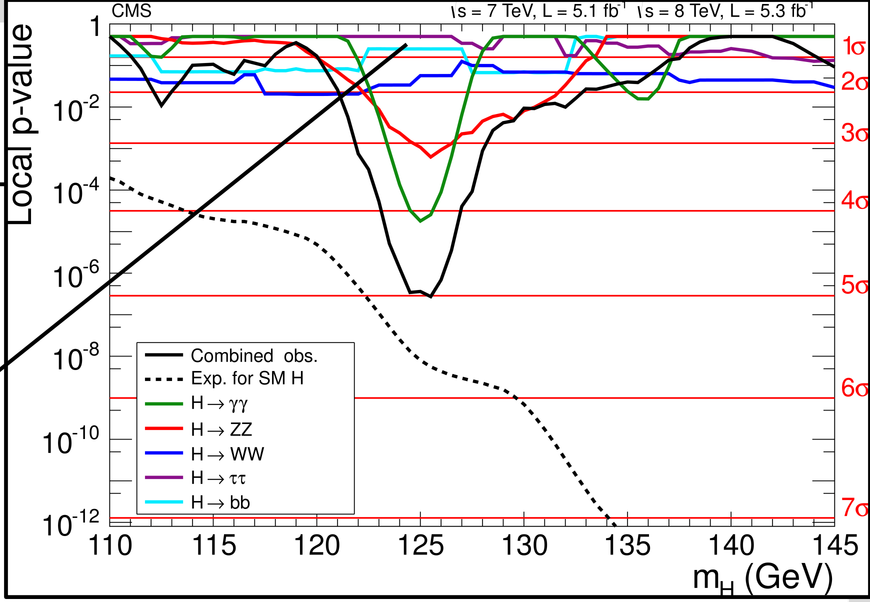

Status **July 2012**:

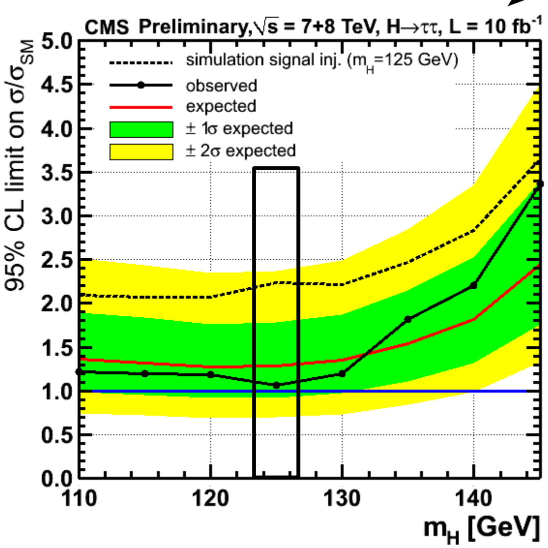

Status **March 2013**:

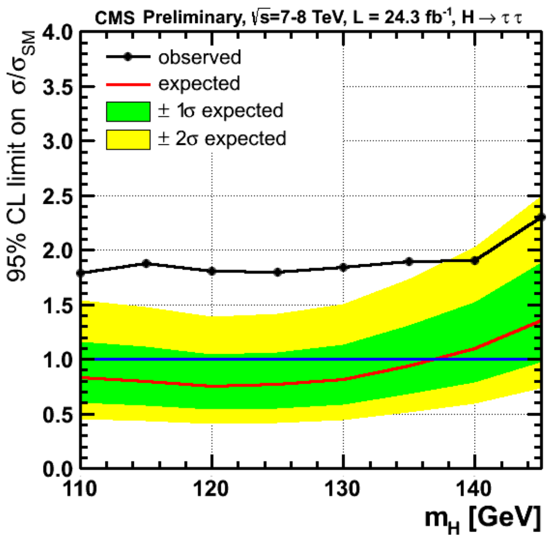

Status **Summer 2014**:

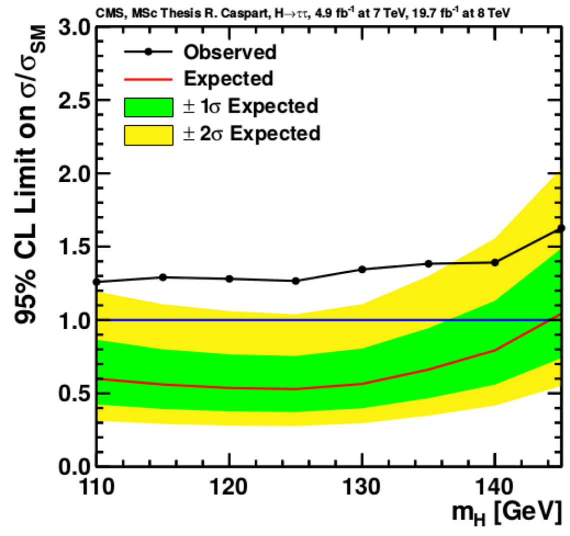

Example: MSSM Higgs search (LHC)

Typical grid in m_A-$\tan\beta$.

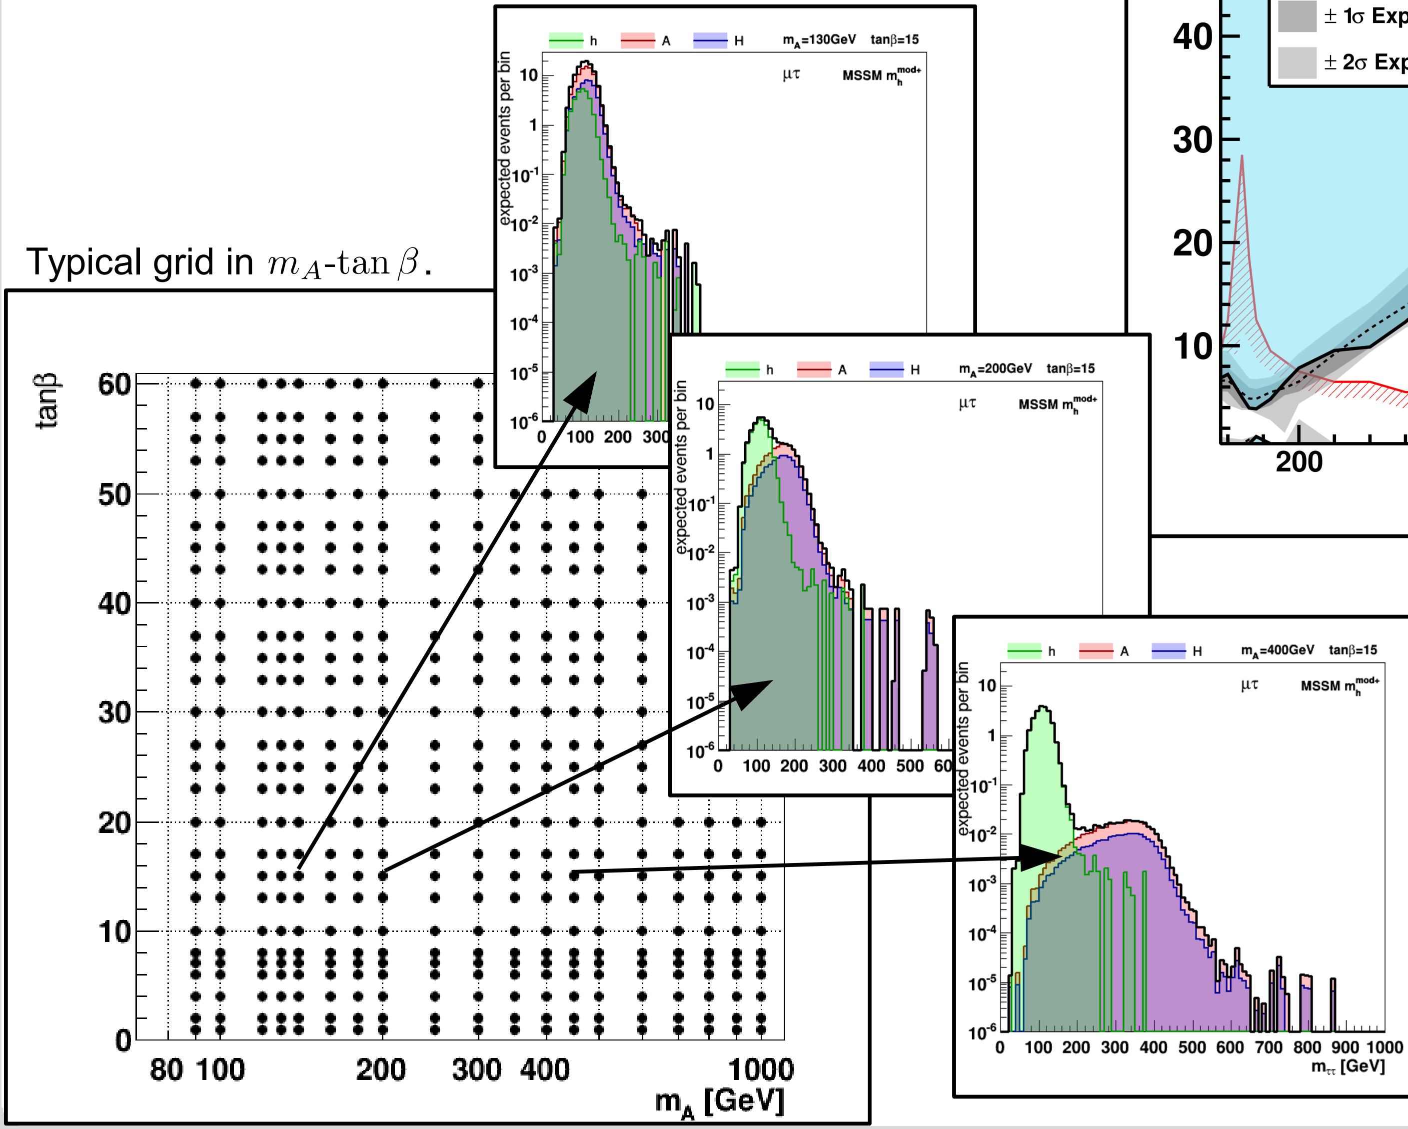

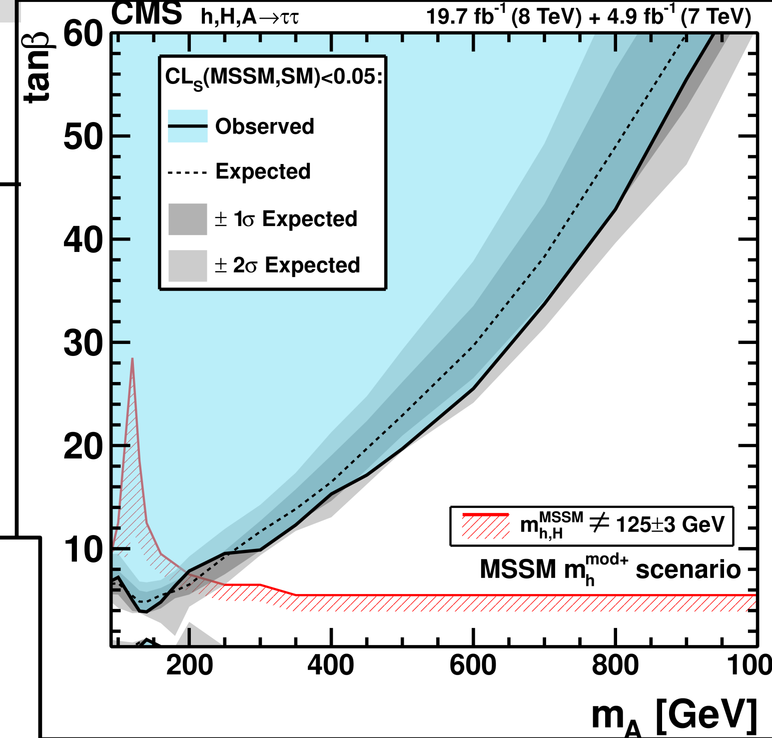

Concluding Remarks

Karlsruhe Institute of Technology

- **Statistics:** mathematical tool to distinguish true from false models.
- The most powerful tool is the likelihood ratio.
- We have discussed: likelihood models, incorporation of systematic uncertainties, maximum likelihood fits, confidence intervals, *p*-values & parameter exclusions.
- Important pillar of each modern physics branch to compare models with measurements.

Backup

Direct Higgs Boson searches @ LEP

- Main production mode in e^+e^-:

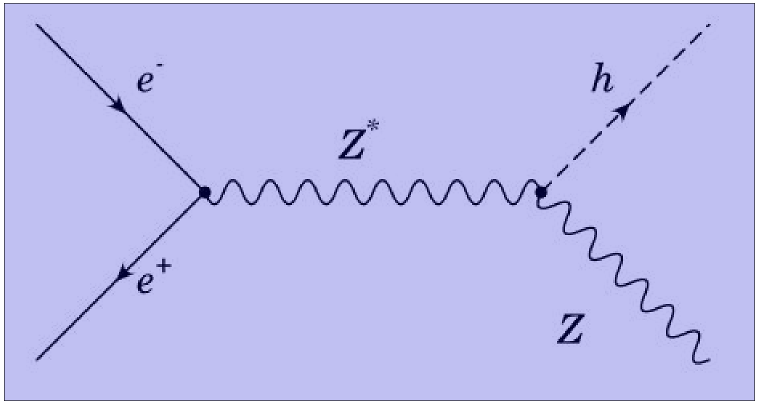

Integrated luminosities in pb^{-1}					
	ALEPH	DELPHI	L3	OPAL	LEP
$\sqrt{s} \geq 189$ GeV	629	608	627	596	2461
$\sqrt{s} \geq 206$ GeV	130	138	139	129	536

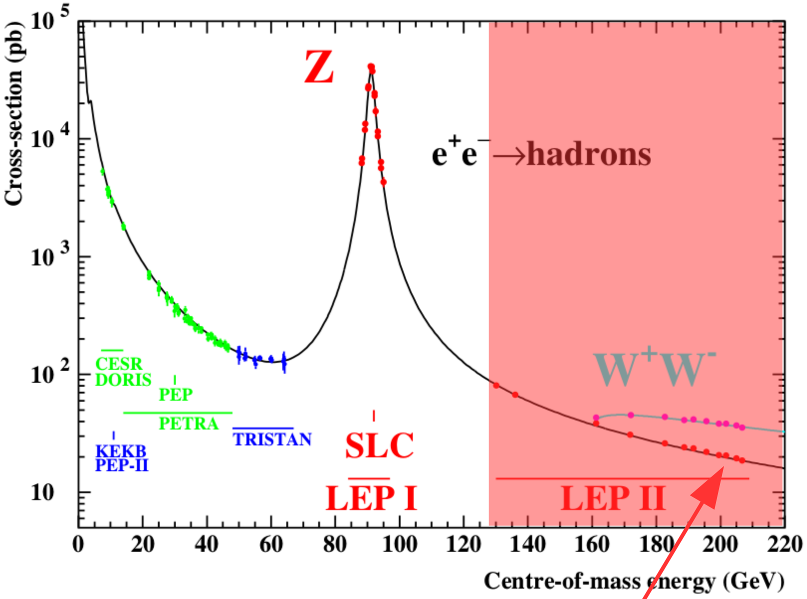

- Higgs boson couples to mass.
- Strongest coupling to heaviest objects.

Year	1996		1997	1998	1999				2000	
E_{CM} nominal [GeV]	161	172	183	189	192	196	200	202	205	207

What was the maximal reach on m_H at LEP? → $m_H \approx 117$ GeV

Result (Final Word from LEP)

p-value:

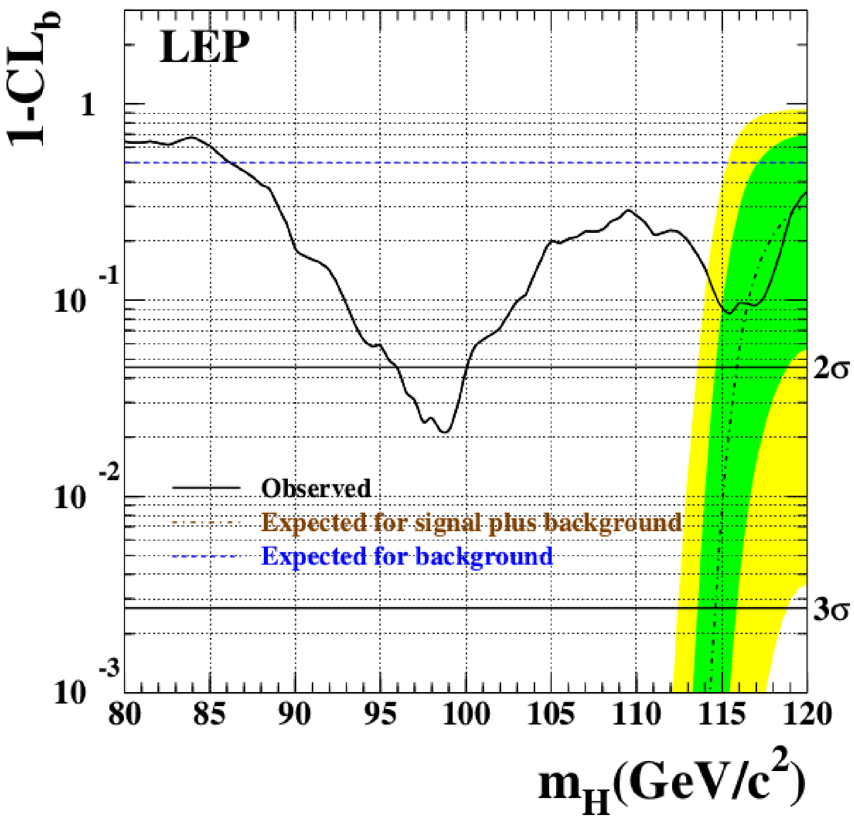

CL_s-limit ($CL_s = \frac{CL_{s+b}}{CL_b}$):

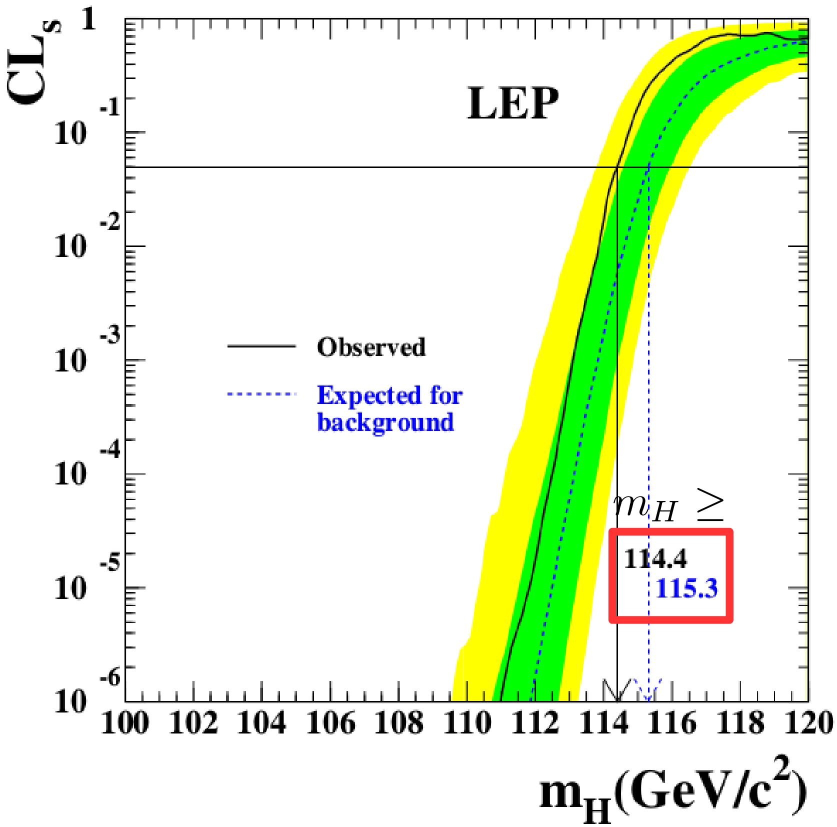

- No signal observed!

p-value vs. *Gaussian* significance

- If the measurement is normal distributed q is distributed according to a χ^2 distribution.
- The χ^2 probability can then be interpreted as a Gaussian confidence interval.

p-values:

$$\mathcal{P}(q \geq 3\sigma | H_0) = 1 \cdot 10^{-3}$$

$$\mathcal{P}(q \geq 5\sigma | H_0) = 2 \cdot 10^{-5}$$