Table 17.3: Poisson upper limits N for n_0 observed events.

n_0	$\varepsilon =$ 10%	$\varepsilon =$ 5%	n_0	$\varepsilon =$ 10%	$\varepsilon =$ 5%
0	2.30	3.00	6	10.53	11.84
1	3.89	4.74	7	11.77	13.15
2	5.32	6.30	8	13.00	14.44
3	6.68	7.75	9	14.21	15.71
4	7.99	9.15	10	15.41	16.96
5	9.27	10.51	11	16.60	18.21

17.5.3. ***Bounded physical region****:

The measurement of a physical constant α results in an estimator $\widehat{\alpha}$, together with some knowledge of experimental error and therefore knowledge of $f(\widehat{\alpha};\alpha)$, the parameterized p.d.f. that allows us to state the probability with which repeated experiments would produce results in a given range. It does *not* permit us to comment about α itself, which in this language is a constant, not a statistical variable. At the beginning of this section we introduced the confidence interval, or frequentist, approach to the problem, and were able to say that with a given probability the unknown parameter could be found between (statistical) limits c_1 and c_2. But what if a physical boundary exists? Although polarization should be less than one and mass or its square should be greater than zero, experimental results do not always fall inside such a physical boundary because of statistical fluctuations.

However one might set a limit, there is little question about how to report and combine data [14]. A given experiment finds an unbiased estimator $\widehat{\alpha} = -5 \pm 10$ for a physical constant (*e.g.* the square of the mass of a neutrino, in eV^2). This value should be reported as the primary result. In case the true value is zero, for example, this "unphysical" result would not be unlikely. It can be combined with the results of other such experiments by forming the appropriately weighted average of unbiased results, including negative ones, to find an unbiased estimator which expresses our best knowledge of the parameter.

What if we wish to extend our concept of confidence limit to such a situation? The question of how to calculate an upper limit in the vicinity of a physical boundary is one of the most divisive in high-energy physics. We present two main approaches: The confidence interval, or frequentist, method, and the Bayesian method. "Classical method" is applied to one or the other by various writers, so we avoid the term.

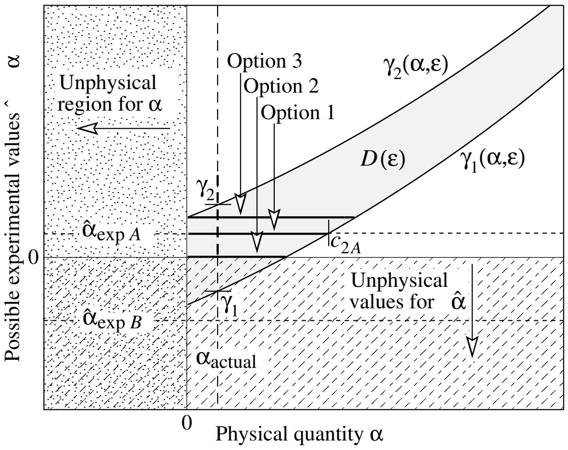

Figure 17.5: The situation near a physical boundary. In Fig. 17.1 the horizontal line for a given $\widehat{\alpha}_{\rm exp}$ crossed the domain $D(\varepsilon)$, bounded by $\gamma_1(\alpha,\epsilon)$ and $\gamma_2(\alpha,\epsilon)$ entirely in the physical region, entering at c_1 and leaving at c_2. The limits γ_1 and γ_2 cannot be defined in a region where α is not defined, so the functions cannot be continued into the unphysical region. As a result c_1 (for experiment A) or c_1 and c_2 (for experiment B) cannot be defined. Options 1, 2, and 3 label the ways one might define confidence intervals, as described in the text.

1. The method of confidence intervals [1,15]. This is the approach described in the introduction, and requires little further explanation. It is presently the method in favor [1,6,14]. For a Gaussian distribution it gives the same result as the Bayesian approach with a flat prior distribution (see below) if the region containing α with the stated probability is far from an unphysical region, as in Fig. 17.1. Two cases in which this is untrue are shown in Fig. 17.5, where as a matter of convenience we assume that α must be positive. As before, we can define limits γ_1 and γ_2 for each value of the unknown parameter α, such that we can expect that a fraction $1-\varepsilon$ of repeated experiments to produce results between these limits. Since this can be done for each value of α, the limits are described by the functions $\gamma_1(\alpha,\varepsilon)$ and $\gamma_2(\alpha,\varepsilon)$. However, these *cannot* be extended into a region in which α makes no sense. Experimental result $\widehat{\alpha}_{{\rm exp}\,A}$, indicated in Fig. 17.5, is positive, but if the true value is $\alpha_{\rm actual}$ a significant fraction of repetitions of the experiment would produce negative $\widehat{\alpha}$. In these cases there is no horizontal intercept c_1, so without further assumptions we cannot make a statement about the region which would cover $\alpha_{\rm actual}$ in a given fraction of experiments. Experimental result $\widehat{\alpha}_{{\rm exp}\,B}$ presents a more serious problem, since it is so negative that there is no physical α for which the point $(\widehat{\alpha},\alpha)$ lies in the domain $D(\varepsilon)$. The reason why the frequentist method gives no confidence interval is clear: This measured value of $\widehat{\alpha}$ would be unlikely *no matter what the true value of α was.*

There are several *ad hoc* ways to set confidence limits in such cases, although many frequentists would prefer to stop with the weighted average of unbiased results—if the outcome is exceedingly unlikely, one should look to the experiment, not to the statistics. The methods we list below all involve placing c_1 on the physical boundary, which in our example is at $\alpha = 0$.

1. If $\widehat{\alpha}_{\rm exp} > \gamma_1(0,\varepsilon)$, as in Experiment A, c_2 is defined. Use it for the upper limit, whether or not $\widehat{\alpha}_{\rm exp} > 0$.
2. If $\widehat{\alpha}_{\rm exp} < 0$, as in Experiment B, use the c_2 corresponding to $\widehat{\alpha}_{\rm exp} = 0$.
3. If c_1 is not defined, "lift up" $\widehat{\alpha}$ to $\gamma_2(0,\varepsilon)$, where $c_1 = 0$. Use the corresponding c_2 as the upper limit.

These three options are shown in in Fig. 17.5; note that there are regions where more than one of them can be used, with different results. The third option is certainly the most conservative. For Gaussian $f(\widehat{\alpha};\alpha)$ the upper limit c_2 is a one-sided Gaussian confidence limit; read the tables for a 90% two-sided limit to obtain 95% one-sided limit. Alternatively, read the intercepts of the dotted lines in Fig. 17.7. (The horizontal axis is incorrectly labeled for this application.)

2. The Bayesian approach [3]. This is the approach favored in the older literature, and has (unfortunately and incorrectly) been referred to as the "PDG method" in certain papers. To begin with, it is argued that while α is not a statistical variable, our *knowledge* of α is less than complete, and it is fair to describe our uncertainty by treating α as a statistical variable. The parameterized p.d.f. $f(\widehat{\alpha};\alpha)$ is replaced by the conditional p.d.f. $f(\widehat{\alpha}|\alpha)$. The confidence limit question can then be rephrased: Our measurements provide $f(\widehat{\alpha}|\alpha)$, that is, information about $\widehat{\alpha}$ for a fixed and unknown value of α, while we really want to know $g(\alpha|\widehat{\alpha})$, which tells us that, given our measurement $\widehat{\alpha}$, the "true answer" α lies between α and $\alpha + d\alpha$ with probability $g(\alpha|\widehat{\alpha})\,d\alpha$. The connection is provided by Bayes' theorem (Eq. (16.7):

$$g(\alpha|\widehat{\alpha}) = \frac{f(\widehat{\alpha}|\alpha)\,\pi(\alpha)}{\int f(\widehat{\alpha}|\alpha)\,\pi(\alpha)\,d\alpha}\ . \tag{17.33}$$

Here $\pi(\alpha)$ represents our "advance knowledge" of the value of α. In the usual case we claim no prior knowledge, so that before the experiment all physically reasonable values of α are equally probable: $\pi(\alpha)$ is a constant over the region of interest and zero in the unphysical region. This assumption leads to the conclusion that

$$g(\alpha|\widehat{\alpha}) = \begin{cases} f(\widehat{\alpha}|\alpha)/\int f(\widehat{\alpha}|\alpha)\,d\alpha & \text{if } \alpha \text{ is in the physical region;} \\ 0 & \text{otherwise;} \end{cases} \tag{17.34}$$

where this time the integral is over the physical region. In Fig. 17.6 we assume that an ensemble of experiments would produce values for $\widehat{\alpha}$ which distribute as shown, with a significant probability of obtaining