results with unphysical values. With our assumed step function $\pi(\alpha)$, the effect of Eq. (17.33) or (17.34) is to replace this distribution with the function shown by the shaded region, except that it is renormalized to unit area. By stating our confidence at the 90% level that α lies below the beginning of the dark shaded region, we mean that 90% of the area in the physical region is in the light shaded region.

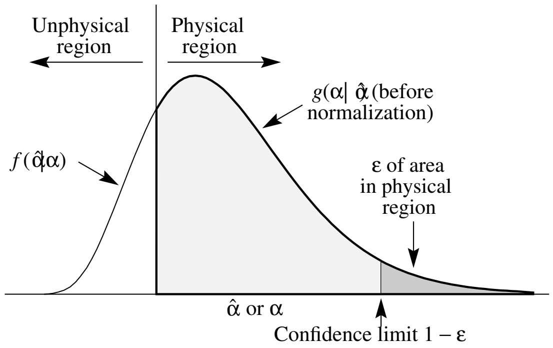

Figure 17.6: An example of a bounded physical region, in which a measurement $\widehat{\alpha}$ can fall in an unphysical region with significant probability. If we assume that α, the quantity we are trying to measure, cannot lie in the unphysical region (0 probability) but can lie anywhere in the physical region ("no prior knowledge"), then Bayes' theorem says that our new knowledge of the distribution of α, given our measurement $\widehat{\alpha}$, is given by the shaded function after appropriate renormalization.

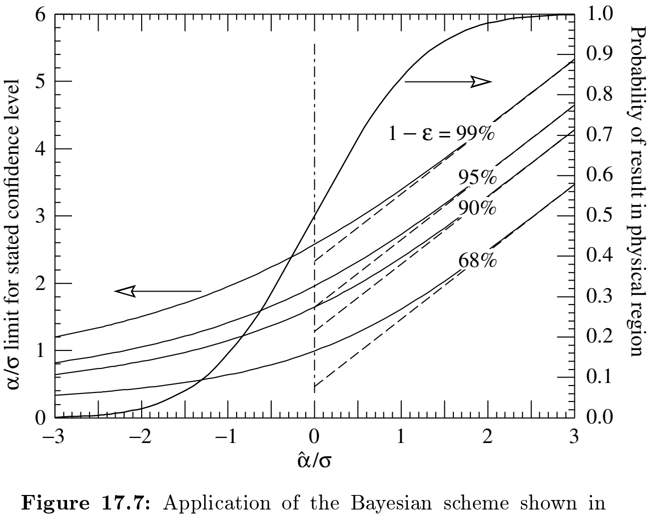

Figure 17.7: Application of the Bayesian scheme shown in Fig. 17.6 to the case of Gaussian $f(\widehat{\alpha}|\alpha)$. For example, if our measurement $\widehat{\alpha}$ is 1.0 standard deviations negative, then we conclude that $\alpha < 1.15\sigma$ with 90% probability—however, there is only a 31% probability that an experimental result as low as this would occur. Note that these are upper limits, so that the asymptote for large $\widehat{\alpha}/\sigma$ corresponds to a one-sided confidence interval, *e.g.*, the asymptote for a 95% confidence level is $\alpha < \widehat{\alpha} + 1.64\sigma$, corresponding to a 90% confidence interval for a two-sided distribution. The dashed lines show the frequentist limit; if Option 3 is used, these are extended horizontally to the right for negative $\widehat{\alpha}/\sigma$.

In most cases of interest in this Review, $\widehat{\alpha}$ is assumed to be a random value from a Gaussian distribution. Application of the procedure sketched in Fig. 17.6 then leads to the family of curves shown in Fig. 17.7. The confidence limit set by this method is always greater than the [one-sided] confidence interval set without the restriction of an unphysical region, and approaches it from above as the tail in the unphysical region becomes unimportant. It is also greater than any of the limits shown in Fig. 17.5. With a small modification (exclusion of that portion of the negative tail inside the physical region in the confidence interval definition), it smoothly approaches the usual two-sided confidence interval for Gaussian distributions.

Even so, it is not a valid confidence limit. If it were, the interval would include the true value of α with *exactly* $1-\varepsilon$ probability no matter what the true value was. If the true answer is zero, our procedure, by guaranteeing a limit greater than zero for any experiment, also guarantees that the confidence interval for any ε includes $\alpha_{\rm actual}$ with 100% probability. Only as α increases does the probability decrease toward the α-independent Gaussian result.

The error function corresponding to the right axis of Fig. 17.7 shows the probability that $\widehat{\alpha}/\sigma$ at or below the given value should occur. If the experimental value is exceeding improbable, then the formal confidence limit obtained by this or any other method means very little.

What about the arbitrariness of $\pi(\alpha)$? If the square of the neutrino mass is measured ($\alpha = m_\nu^2$), then should we not take the prior knowledge distribution as proportional to $\sqrt{\alpha}$?† There are other attractive options. Jeffreys points out that if $\pi(\alpha)\,d\alpha = d\alpha/\alpha$, then the distributions for α and α^n are proportional [16], but there are practical difficulties with this approach. Lynch has investigated prior distributions that are constant in α, α^2, and $\sqrt{\alpha}$ in the context of Gaussian $f(\widehat{\alpha}|\alpha)$, and has observed that assuming a prior distribution that is flat in α gives results that are much more satisfactory than one gets from the others [17]: All three methods have the property that the probability that the calculated limit contains the correct answer is 100% when $\alpha = 0$ and approaches the proper value when $\alpha >> \sigma$, but the approach to the proper value as α increases is much faster when the prior distribution is taken to be flat in α. In this case the approach is also monotonic, giving it the "conservative" property that for no value of α will the method produce a limit that has a probability of being correct that is less than the stated confidence limit. Although there is nothing unique about the limit calculated with a constant $\pi(\alpha)$, it has desirable features and no obvious replacement.

Summary: If there is a significant probability of obtaining an estimator corresponding to an unphysical value for a parameter, there is no universally accepted way way to make a statement of the sort "α is less that c_2 with probability $1-\varepsilon$." A variety of upper limits can be defined, but no method is entirely satisfactory. The Bayesian method with a flat prior distribution gives a reasonable upper limit which combines everything we know about the unknown quantity α into a physically reasonable value, but it does not give a complete summary of the information contained in the experiment.

17.5.4. *Poisson processes with background* [18] :

If we observe n_0 events in a Poisson process which has two components, signal and background, estimating a limit on the signal is more complicated. Let μ_S be the unknown mean (the Poisson parameter) for the signal and μ_B be the mean for the sum of all backgrounds. Assume μ_B is known with negligible error; however we don't know n_B, the actual number of events resulting from the background. We do know that $n_B \le n_0$. If $\mu_B + \mu_S$ is large, the Gaussian approximation to the Poisson distribution (see Sec. 16.3.2) is usually adequate, and one can define confidence intervals or limits as above, assuming $\widehat{n}_B \approx \mu_B$ and therefore $\widehat{\mu}_S = n_0 - \mu_B$ with variance equal to n_0 (larger than $\widehat{\mu}_S$ to allow for the error in $\widehat{n}_B$).

Otherwise an upper limit can be defined by extension of the argument of the preceding section. Let N be the desired upper limit on μ_S with confidence coefficient ε. Set N to be that value of μ_S such that any random repeat of the current experiment with $\mu_S = N$ and the same μ_B would observe *more* than n_0 events in total *and* would have $n_B \le n_0$, all with probability $1-\varepsilon$. For any assumed N and μ_B we can calculate this probability:

$$1-\varepsilon = 1 - \frac{e^{-(\mu_B+N)} \displaystyle\sum_{n=0}^{n_0} \frac{(\mu_B+N)^n}{n!}}{e^{-\mu_B} \displaystyle\sum_{n=0}^{n_0} \frac{\mu_B^n}{n!}} \,. \tag{17.35}$$