Chapter 11

The Poisson Distribution

This chapter presents a third example of a limiting distribution, the Poisson distribution, which describes the results of experiments in which we count events that occur at random but at a definite average rate. Examples of this kind of counting experiment crop up in almost every area of science; for instance, a sociologist might count the number of babies born in a hospital in a three-day period. An important example in physics is the counting of the decays of a radioactive sample; for instance, a nuclear physicist might decide to count the number of alpha particles given off by a sample of radon gas in a ten-second interval.

This kind of counting experiment was discussed in Section 3.2, where I stated but did not prove the "square-root rule": If you count the occurrences of an event of this type in a chosen time interval T and obtain ν counts, then the best estimate for the true average number in time T is, of course, ν, and the uncertainty in this estimate is $\sqrt{\nu}$.

In Sections 11.1 and 11.2, I introduce the Poisson distribution and explore some of its properties. In particular, I prove in Section 11.2 that the standard deviation of the Poisson distribution is the square root of the expected number of events. This result justifies the square-root rule of Section 3.2. Sections 11.3 and 11.4 describe some applications of the Poisson distribution.

11.1 Definition of the Poisson Distribution

As an example of the Poisson distribution, suppose we are given a sample of radioactive material and use a suitable detector to find the number ν of decay particles ejected in a two-minute interval. If the counter is reliable, our value of ν will have no uncertainty. Nevertheless, if we repeat the experiment, we will almost certainly get a different value for ν. This variation in the number ν does not reflect uncertainties in our counting; rather, it reflects the intrinsically random character of the radioactive decay process.

Each radioactive nucleus has a definite probability for decaying in any two-minute interval. If we knew this probability and the number of nuclei in our sample, we could calculate the *expected average number* of decays in two minutes. Nevertheless, each nucleus decays at a random time, and in any given two-minute interval, the number of decays may be different from the expected average number.

Obviously the question we should ask is this: If we repeat our experiment many times (replenishing our sample if it becomes significantly depleted), what distribution should we expect for the number of decays ν observed in two-minute intervals? If you have studied Chapter 10, you will recognize that the required distribution is the binomial distribution. If there are n nuclei and the probability that any one nucleus decays is p, then the probability of ν decays is just the probability of ν "successes" in n "trials," or $B_{n,p}(\nu)$. In the kind of experiment we are now discussing, however, there is an important simplification. The number of "trials" (that is, nuclei) is enormous ($n \sim 10^{20}$, perhaps), and the probability of "success" (decay) for any one nucleus is tiny (often as small as $p \sim 10^{-20}$). Under these conditions (n large and p small), the binomial distribution can be shown to be indistinguishable from a simpler function called the Poisson distribution. Specifically, it can be shown that

$$Prob(\nu \text{ counts in any definite interval}) = P_\mu(\nu), \tag{11.1}$$

where the *Poisson distribution*, $P_\mu(\nu)$, is given by

The Poisson Distribution

$$P_\mu(\nu) = e^{-\mu}\frac{\mu^\nu}{\nu!}. \tag{11.2}$$

In this definition, μ is a positive parameter ($\mu > 0$) that, as I will show directly, is just the expected mean number of counts in the time interval concerned, and $\nu!$ denotes the factorial function (with $0! = 1$).

SIGNIFICANCE OF μ AS THE EXPECTED MEAN COUNT

I will not derive the Poisson distribution (11.2) here but simply assert that it *is* the appropriate distribution for the kind of counting experiment in which we are interested.[1] To establish the significance of the parameter μ in (11.2), we have only to calculate the average number of counts, $\bar{\nu}$, expected if we repeat our counting experiment many times. This average is found by summing over all possible values of ν, each multiplied by its probability:

$$\bar{\nu} = \sum_{\nu=0}^{\infty} \nu P_\mu(\nu) = \sum_{\nu=0}^{\infty} \nu e^{-\mu}\frac{\mu^\nu}{\nu!}. \tag{11.3}$$

The first term of this sum can be dropped (because it is zero), and $\nu/\nu!$ can be replaced by $1/(\nu - 1)!$. If we remove a common factor of $\mu e^{-\mu}$, we get

$$\bar{\nu} = \mu e^{-\mu}\sum_{\nu=1}^{\infty}\frac{\mu^{\nu-1}}{(\nu-1)!}. \tag{11.4}$$

[1] For derivations, see, for example, H. D. Young, *Statistical Treatment of Experimental Data* (McGraw-Hill, 1962), Section 8, or S. L. Meyer, *Data Analysis for Scientists and Engineers* (John Wiley, 1975), p. 207.

The infinite sum that remains is

$$1 + \mu + \frac{\mu^2}{2!} + \frac{\mu^3}{3!} + \cdots = e^\mu, \tag{11.5}$$

which is just the exponential function e^μ (as indicated). Thus, the exponential $e^{-\mu}$ in (11.4) is exactly canceled by the sum, and we are left with the simple conclusion that

$$\bar{\nu} = \mu. \tag{11.6}$$

That is, the parameter μ that characterizes the Poisson distribution $P_\mu(\nu)$ is just the *average number of counts expected if we repeat the counting experiment many times.*

Sometimes, we may know in advance the average rate R at which the events we are counting should occur. In this case, the expected average number of events in a time T is just

$$\mu = \text{rate} \times \text{time} = RT.$$

Conversely, if the rate R is unknown, then by counting the number of events in a time T, we can get an estimate for μ and hence for the rate R as $R_{best} = \mu_{best}/T$.

Example: Counting Radioactive Decays

Careful measurements have established that a sample of radioactive thorium emits alpha particles at a rate of 1.5 per minute. If I count the number of alpha particles emitted in two minutes, what is the expected average result? What is the probability that I would actually get this number? What is the probability for observing ν particles for $\nu = 0, 1, 2, 3, 4$, and for $\nu \geqslant 5$?

The expected average count is just the average rate of emissions ($R = 1.5$ per minute) multiplied by the time during which I make my observations ($T = 2$ minutes):

$$(\text{expected average number}) = \mu = 3.$$

This result does not mean, of course, that I should expect to observe exactly three particles in any single trial. On the contrary, the probabilities for observing any number (ν) of particles are given by the Poisson distribution

$$Prob(\nu \text{ particles}) = P_3(\nu) = e^{-3}\frac{3^\nu}{\nu!}.$$

In particular, the probability that I would observe exactly three particles is

$$Prob(3 \text{ particles}) = P_3(3) = e^{-3}\frac{3^3}{3!} = 0.22 = 22\%.$$

Notice that although the expected average result is $\nu = 3$, we should expect to get this number only about once in every five trials.

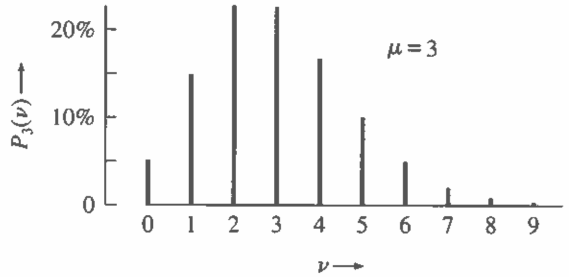

Figure 11.1. The Poisson distribution $P_3(\nu)$ gives the probabilities of observing ν events in a counting experiment for which the expected average count is 3.

The probabilities for any number ν can be calculated in the same way and are (as you might want to check):

Number ν:	0	1	2	3	4
Probability:	5%	15%	22%	22%	17%

These probabilities (up to $\nu = 9$) are plotted in Figure 11.1. The simplest way to find the probability for getting 5 or more counts is to add the probabilities for $\nu = 0, \ldots, 4$ and then subtract the sum from 100% to give[2]

$$\begin{aligned} Prob(\nu \geq 5) &= 100\% - (5 + 15 + 22 + 22 + 17)\% \\ &= 19\%. \end{aligned}$$

Quick Check 11.1. On average, each of the 18 hens in my henhouse lays 1 egg per day. If I check the hens once an hour and remove any eggs that have been laid, what is the average number, μ, of eggs that I find on my hourly visits? Use the Poisson distribution $P_\mu(\nu)$ to calculate the probabilities that I would find ν eggs for $\nu = 0$, 1, 2, 3, and $\nu = 4$ or more. What is the most probable number? What is the probability that I would find exactly μ eggs? Verify the probabilities shown in Figure 11.2.

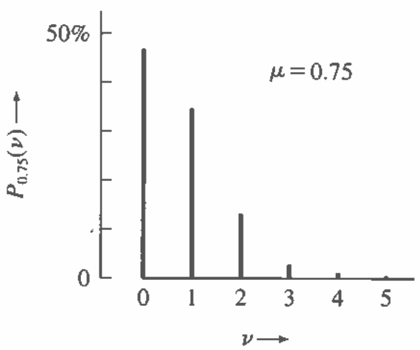

Figure 11.2. The Poisson distribution $P_{0.75}(\nu)$ gives the probabilities of observing ν events in a counting experiment for which the expected average count is 0.75.

---

[2]The correct answer is actually 18.48%, as you can check by keeping a couple of extra decimal places in all the probabilities.

11.2 Properties of the Poisson Distribution

THE STANDARD DEVIATION

The Poisson distribution $P_\mu(\nu)$ gives the probability of getting the result ν in an experiment in which we count events that occur at random but at a definite average rate. We have seen that the parameter μ is precisely the expected average count $\bar{\nu}$. The natural next question is to ask for the standard deviation of the counts ν when we repeat the experiment many times. The standard deviation of any distribution (after a large number of trials) is just the root-mean-square deviation from the mean. That is,

$$\sigma_\nu^2 = \overline{(\nu - \bar{\nu})^2},$$

or, using the result of Problems 10.15(a) or 4.5(a),

$$\sigma_\nu^2 = \overline{\nu^2} - (\bar{\nu})^2. \quad (11.7)$$

For the Poisson distribution, we have already found that $\bar{\nu} = \mu$ and a similar calculation (Problem 11.9) gives $\overline{\nu^2} = \mu^2 + \mu$. Therefore, Equation (11.7) implies that $\sigma_\nu^2 = \mu$ or

$$\boxed{\sigma_\nu = \sqrt{\mu}.} \quad (11.8)$$

That is, the Poisson distribution with mean count μ has standard deviation $\sqrt{\mu}$.

The result (11.8) justifies the square-root rule of Section 3.2. If we carry out a counting experiment once and get the answer ν, we can easily see (using the principle of maximum likelihood, as in Problem 11.11) that the best estimate for the expected mean count is $\mu_{best} = \nu$. From (11.8), it immediately follows that the best estimate for the standard deviation is just $\sqrt{\nu}$. In other words, if we make one measurement of the number of events in a time interval T and get the answer ν, our answer for the expected mean count in time T is

$$\nu \pm \sqrt{\nu}. \quad (11.9)$$

This answer is precisely the square-root rule quoted without proof in Equation (3.2).

Example: More Radioactive Decays

A student monitors the thorium sample of the previous example for 30 minutes and observes 49 alpha particles. What is her answer for the number of particles emitted in 30 minutes? What is her answer for the rate of emission, R, in particles per minute?

According to (11.9), her answer for the number of particles emitted in 30 minutes is

$$\text{(number emitted in 30 minutes)} = 49 \pm \sqrt{49} = 49 \pm 7.$$

To find the rate in particles per minute, she must divide by 30 minutes. Assuming this 30 minutes has no uncertainty, we find

$$R = \frac{49 \pm 7}{30} = 1.6 \pm 0.2 \text{ particles/min.} \tag{11.10}$$

Notice that the square-root rule gives the uncertainty in the actual counted number ($\sigma_\nu = \sqrt{\nu} = 7$, in this case). A common mistake is to calculate the rate of decay $R = \nu/T$ and then to take the uncertainty in R to be $\sqrt{R}$. A glance at (11.10) should convince you that this procedure is simply not correct. The square-root rule applies only to the actual counted number ν, and the uncertainty in $R = \nu/T$ must be found from that in ν using error propagation as in (11.10).

Quick Check 11.2. The farmer of Quick Check 11.1 observes that in a certain ten-hour period his hens lay 9 eggs. Based on this one observation, what would you quote for the number of eggs expected in ten hours? What would you give for the rate R of egg production, in eggs per hour? (Give the uncertainties in both your answers.)

GAUSSIAN APPROXIMATION TO THE POISSON DISTRIBUTION

In Chapter 10, we compared the Gauss distribution with the binomial distribution. We saw that in most ways, the two distributions are very different; nevertheless, under the right conditions, the Gauss distribution gives an excellent, extremely useful, approximation to the binomial distribution. As we will now see, almost exactly the same can be said about the Gauss and *Poisson* distributions.

The Gauss distribution $G_{X,\sigma}(x)$ gives the probabilities of the various values of a *continuous* variable x; by contrast, the Poisson distribution $P_\mu(\nu)$, like the binomial $B_{n,p}(\nu)$, gives the probabilities for a *discrete* variable $\nu = 0, 1, 2, 3, \ldots$. Another important difference is that the Gauss distribution $G_{X,\sigma}(x)$ is specified by *two* parameters, the mean X and the standard deviation σ, whereas the Poisson distribution $P_\mu(\nu)$ is specified by a single parameter, the mean μ, because the width of the Poisson distribution is automatically determined by the mean (namely, $\sigma_\nu = \sqrt{\mu}$). Finally, the Gauss distribution is always bell shaped and symmetric about its mean value, whereas the Poisson distribution has neither of these properties in general. This last point is especially clear in Figure 11.2, which shows the Poisson distribution for $\mu = 0.75$; this curve is certainly not bell shaped, nor is it even approximately symmetric about its mean, 0.75.

Figure 11.1 showed the Poisson distribution for $\mu = 3$. Although this curve is obviously not exactly bell shaped, it is undeniably more nearly so than the curve for $\mu = 0.75$ in Figure 11.2. Figure 11.3 shows the Poisson distribution for $\mu = 9$; this curve is quite nearly bell shaped and close to symmetric about its mean ($\mu = 9$). In fact, it can be proved that as $\mu \to \infty$, the Poisson distribution becomes progres-

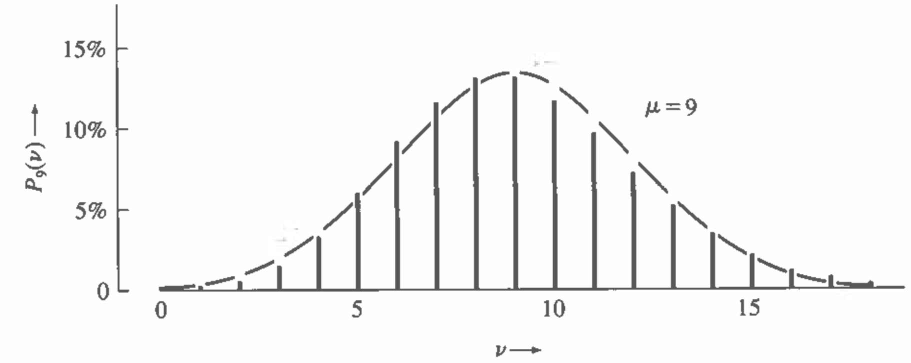

Figure 11.3. The Poisson distribution for $\mu = 9$. The dashed curve is the Gauss distribution with the same mean and standard deviation ($X = 9$ and $\sigma = 3$). As $\mu \to \infty$, the two distributions become indistinguishable; even when $\mu = 9$, they are very close.

sively more bell shaped and approaches the Gauss distribution with the same mean and standard deviation.[3] That is,

$$P_\mu(\nu) \approx G_{X,\sigma}(\nu), \qquad \text{(when } \mu \text{ is large)} \tag{11.11}$$

where

$$X = \mu \quad \text{and} \quad \sigma = \sqrt{\mu}.$$

In Figure 11.3, the dashed curve is the Gauss function with $X = 9$ and $\sigma = 3$. You can see clearly how, even when μ is only 9, the Poisson distribution is remarkably close to the appropriate Gauss function; the slight discrepancy reflects the remaining asymmetry in the Poisson function.

The approximation (11.11) is called the *Gaussian approximation to the Poisson distribution.* It is analogous to the corresponding approximation for the binomial distribution (discussed in Section 10.4) and is useful under the same conditions, namely, when the parameters involved are large.

Example: Gaussian Approximation to a Poisson Distribution

To illustrate the Gaussian approximation to the Poisson distribution, consider the Poisson distribution with $\mu = 64$. The probability of 72 counts, for example, is

$$Prob(72 \text{ counts}) = P_{64}(72) = e^{-64}\frac{(64)^{72}}{72!}, \tag{11.12}$$

which a tedious calculation gives as

$$Prob(72 \text{ counts}) = 2.9\%.$$

[3] For proof, see S. L. Meyer, *Data Analysis for Scientists and Engineers* (John Wiley, 1975), p. 227.

According to (11.11), however, the probability (11.12) is well approximated by the Gauss function

$$Prob(72 \text{ counts}) \approx G_{64,8}(72),$$

which is easily evaluated to give

$$Prob(72 \text{ counts}) \approx 3.0\%.$$

If we wanted to calculate directly the probability of 72 *or more* counts in the same experiment, an extremely tedious calculation would give

$$Prob(\nu \geq 72) = P_{64}(72) + P_{64}(73) + \ldots$$
$$= 17.3\%.$$

If we use the approximation (11.11), then we have only to calculate the Gaussian probability for getting $\nu \geq 71.5$ (because the Gauss distribution treats ν as a continuous variable). Because 71.5 is 7.5, or 0.94σ, above the mean, the required probability can be found quickly from the table in Appendix B as

$$Prob(\nu \geq 72) \approx Prob_G(\nu \geq 71.5) = Prob_G(\nu \geq X + 0.94\sigma)$$
$$= 17.4\%,$$

by almost any standard an excellent approximation.

11.3 Applications

As I have emphasized, the Poisson distribution describes the distribution of results in a counting experiment in which events are counted that occur at random but at a definite average rate. In an introductory physics laboratory, the two most common examples are counting the disintegrations of radioactive nuclei and counting the arrival of cosmic ray particles.

Another very important example is an experiment to study an expected limiting distribution, such as the Gauss or binomial distributions, or the Poisson distribution itself. A limiting distribution tells us how many events of a particular type are expected when an experiment is repeated several times. (For example, the Gaussian $G_{X,\sigma}(x)$ tells us how many measurements of x are expected to fall in any interval from $x = a$ to $x = b$.) In practice, the observed number is seldom exactly the expected number. Instead, it fluctuates in accordance with the Poisson distribution. In particular, if the expected number of events of some type is n, the observed number can be expected to differ from n by a number of order $\sqrt{n}$. We will make use of this point in Chapter 12.

In many situations, it is reasonable to expect numbers to be distributed approximately according to the Poisson distribution. The number of eggs laid in an hour on a poultry farm and the number of births in a day at a hospital would both be expected to follow the Poisson distribution at least approximately (though they would probably show some seasonal variations as well). To test this assumption, you would need to record the number concerned many times over. After plotting the resulting distribution, you could compare it with the Poisson distribution to see how close the fit is. For a more quantitative test, you would use the chi-squared test described in Chapter 12.

Example: Cosmic Ray Counting

As another example of the Poisson distribution, let us consider an experiment with cosmic rays. These "rays" originate as charged particles, such as protons and alpha particles, that enter the Earth's atmosphere from space. Many of these primary particles collide with atoms in the atmosphere and create further, secondary particles, such as mesons and positrons. Some of the particles (both primary and secondary) travel all the way to ground level and can be detected (with a Geiger counter, for example) in the laboratory. In the following problem, I exploit the fact that the number of cosmic rays hitting any given area in a given time should follow the Poisson distribution.

Student A asserts that he has measured the number of cosmic rays hitting a Geiger counter in one minute. He claims to have made the measurement repeatedly and carefully and to have found that, on average, 9 particles hit the counter per minute, with "negligible" uncertainty. To check this claim, Student B counts how many particles arrive in one minute and gets the answer 12. Does this answer cast serious doubt on A's claim that the expected rate is 9? To make a more careful check, Student C counts the number of particles that arrive in ten minutes. From A's claim, she expects to get 90 but actually gets 120. Does this value cast significant doubt on A's claim?

Let us consider B's result first. If A is right, the expected mean count is 9. Because the distribution of counts should be the Poisson distribution, the standard deviation is $\sqrt{9} = 3$. Student B's result of 12 is, therefore, only one standard deviation away from the mean of 9. This amount is certainly not far enough away to contradict A's claim. More specifically, knowing that the probability of any answer ν is supposed to be $P_9(\nu)$, we can calculate the total probability for getting an answer that differs from 9 by 3 or more. This probability turns out to be 40% (see Problem 11.7). Obviously B's result is not at all surprising, and A has no reason to worry.

Student C's result is quite a different matter. If A is right, C should expect to get 90 counts in ten minutes. Since the distribution should be Poisson, the standard deviation should be $\sqrt{90} = 9.5$. Thus, C's result of 120 is more than *three standard deviations* away from A's prediction of 90. With these large numbers, the Poisson distribution is indistinguishable from the Gauss function, and we can immediately find from the table in Appendix A that the probability of a count more than three standard deviations from the mean is 0.3%. That is, if A is right, it is extremely improbable that C would have observed 120 counts. Turning this statement around, we can say something almost certainly has gone wrong. Perhaps A was just not as careful as he claimed. Perhaps the counter was malfunctioning for A or C, introducing systematic errors into one of the results. Or perhaps A made his measurements at a time when the flux of cosmic rays was truly less than normal.

11.4 Subtracting a Background

To conclude this chapter, I discuss a problem that complicates many counting experiments. Often, the events we want to study are accompanied by other "background" events that cannot be distinguished from the events of interest. For example, in studying the disintegrations of a radioactive source, we usually cannot prevent the detector from registering particles from other radioactive materials in the vicinity or from cosmic rays. This means that the number we count includes the events of interest *plus* these background events, and we must somehow subtract out the unwanted background events. In principle at least, the remedy is straightforward: Having found the total counting rate (due to source and background), we must remove the source and find the rate of events due to the background alone; the rate of events from the source is then just the difference of these two measured rates.

In practice, it is surprisingly easy to make a mistake in this procedure, especially in the error analysis. It is usually convenient to measure the total and background counts using different time intervals. Suppose we count a total of ν_{tot} events (source *plus* background) in a time T_{tot}, and then ν_{bgd} background events in a time T_{bgd}. Obviously, we do not simply subtract ν_{bgd} from ν_{tot} because they refer to different time intervals. Instead, we must first calculate the *rates*

$$R_{tot} = \frac{\nu_{tot}}{T_{tot}} \quad \text{and} \quad R_{bgd} = \frac{\nu_{bgd}}{T_{bgd}} \tag{11.13}$$

and then calculate the rate from the source as the difference

$$R_{sce} = R_{tot} - R_{bgd}. \tag{11.14}$$

In estimating the uncertainties in the quantities involved, you must remember that the square-root rule gives the uncertainties in the *counted numbers* ν_{tot} and ν_{bgd}; the uncertainties in the corresponding rates must be found by error propagation, as in the following example.

Example: Radioactive Decays with a Background

A student decides to monitor the activity of a radioactive source by placing it in a liquid scintillation detector. In the course of 10 minutes, the detector registers 2,540 total counts. To allow for the possibility of unwanted background counts, she removes the source and notes that in 3 minutes, the detector registers a further 95 counts. To find the activity of the source, she calculates the two rates of counting, R_{tot} and R_{bgd} (in counts per minute) and their difference $R_{sce} = R_{tot} - R_{bgd}$. What are her answers with their uncertainties? (Assume the two times have negligible uncertainty.)

According to the square-root rule, the two counted numbers with their uncertainties are

$$\nu_{tot} = 2{,}540 \pm \sqrt{2{,}540} = 2{,}540 \pm 50$$

and

$$\nu_{bgd} = 95 \pm \sqrt{95} = 95 \pm 10.$$

Dividing these numbers by their corresponding times, we find the rates

$$R_{tot} = \frac{\nu_{tot}}{T_{tot}} = \frac{2{,}540 \pm 50}{10} = 254 \pm 5 \text{ counts/min}$$

and

$$R_{bgd} = \frac{\nu_{bgd}}{T_{bgd}} = \frac{95 \pm 10}{3} = 32 \pm 3 \text{ counts/min.}$$

Finally, the rate due to the source alone is

$$R_{sce} = R_{tot} - R_{bgd} = (254 \pm 5) - (32 \pm 3) = 222 \pm 6 \text{ counts/min.}$$

Notice that, in the last step, the errors are combined in quadrature, because they are certainly independent and random.

Principal Definitions and Equations of Chapter 11

THE POISSON DISTRIBUTION

The Poisson distribution describes experiments in which you count events that occur at random but at a definite average rate. If you count for a chosen time interval T, the probability of observing ν events is given by the Poisson function

$$\textit{Prob}(\nu \text{ counts in time } T) = P_\mu(\nu) = e^{-\mu}\frac{\mu^\nu}{\nu!}, \qquad \text{[See (11.2)]}$$

where the parameter μ is the expected average number of events in time T; that is,

$$\bar{\nu} = \mu \text{ (after many trials)}$$
$$= RT, \qquad \text{[See (11.6)]}$$

where R is the mean rate at which the events occur.

The standard deviation of the observed number ν is

$$\sigma_\nu = \sqrt{\mu}. \qquad \text{[See (11.8)]}$$

THE GAUSSIAN APPROXIMATION TO THE POISSON DISTRIBUTION

When μ is large, the Poisson distribution $P_\mu(\nu)$ is well approximated by the Gauss function with the same mean and standard deviation; that is,

$$P_\mu(\nu) \approx G_{X,\sigma}(\nu),$$

where $X = \mu$ and $\sigma = \sqrt{\mu}$. [See (11.11)]

SUBTRACTING A BACKGROUND

The events produced by a source subject to an unavoidable background can be counted in a three-step procedure:

(1) Count the total number ν_{tot} (source plus background) in a time T_{tot}, and calculate the total rate $R_{tot} = \nu_{tot}/T_{tot}$.
(2) Remove the source, and measure the number of background events in a time T_{bgd}; then calculate the background rate $R_{bgd} = \nu_{bgd}/T_{bgd}$.
(3) Calculate the rate of the events from the source as the difference $R_{sce} = R_{tot} - R_{bgd}$.

Finally, the uncertainties in the numbers ν_{tot} and ν_{bgd} are given by the square-root rule, and, from these values, the uncertainties in the three rates can be found using error propagation.

Problems for Chapter 11

For Section 11.1: Definition of the Poisson Distribution

11.1. ★ Compute the Poisson distribution $P_\mu(\nu)$ for $\mu = 0.5$ and $\nu = 0, 1, \ldots, 6$. Plot a bar histogram of $P_{0.5}(\nu)$ against ν.

11.2. ★ **(a)** Compute the Poisson distribution $P_\mu(\nu)$ for $\mu = 1$ and $\nu = 0, 1, \ldots, 6$, and plot your results as a bar histogram. **(b)** Repeat part (a) but for $\mu = 2$.

11.3. ★★ A radioactive sample contains 5.0×10^{19} atoms, each of which has a probability $p = 3.0 \times 10^{-20}$ of decaying in any given five-second interval. **(a)** What is the expected average number, μ, of decays from the sample in five seconds? **(b)** Compute the probability $P_\mu(\nu)$ of observing ν decays in any five-second interval, for $\nu = 0, 1, 2, 3$. **(c)** What is the probability of observing 4 or more decays in any five-second interval?

11.4. ★★ In the course of four weeks, a farmer finds that between 10:00 and 10:30 A.M., his hens lay an average of 2.5 eggs. Assuming the number of eggs laid follows a Poisson distribution with $\mu = 2.5$, on approximately how many days do you suppose he found no eggs laid between 10:00 and 10:30 A.M.? On how many days do you suppose there were 2 or less? 3 or more?

11.5. ★★ A certain radioactive sample is expected to undergo three decays per minute. A student observes the number ν of decays in 100 separate one-minute intervals, with the results shown in Table 11.1. **(a)** Make a histogram of these results, plotting f_ν (the fraction of times the result ν was found) against ν. **(b)** On the same plot, show the expected distribution $P_3(\nu)$. Do the data seem to fit the expected distribution? (For a quantitative measure of the fit, you could use the chi-squared test, discussed in Chapter 12.)

Table 11.1. Occurrences of numbers of decays in one-minute intervals; for Problem 11.5.

No. of decays ν	0	1	2	3	4	5	6	7	8	9
Times observed	5	19	23	21	14	12	3	2	1	0

11.6. ★★★ **(a)** The Poisson distribution, like all distributions, must satisfy a "normalization condition,"

$$\sum_{\nu=0}^{\infty} P_\mu(\nu) = 1. \qquad (11.15)$$

This condition asserts that the total probability of observing *all* possible values of ν must be one. Prove it. [Remember the infinite series (11.5) for e^μ.] **(b)** Differentiate (11.15) with respect to μ, and then multiply the result by μ to give an alternative proof that, after infinitely many trials, $\bar{\nu} = \mu$ as in Equation (11.6).

For Section 11.2: Properties of the Poisson Distribution

11.7. ★ **(a)** What is the standard deviation σ_ν (after a large number of trials) of the observed counts ν in a counting experiment in which the expected average count is $\mu = 9$? **(b)** Compute the probabilities $P_9(\nu)$ of obtaining ν counts for $\nu = 7, 8, \ldots, 11$. **(c)** Hence, find the probability of getting a count ν that differs from the expected mean by one or more standard deviations. **(d)** Would a count of 12 cause you to doubt that the expected mean really is 9?

11.8. ★ A nuclear physicist monitors the disintegrations of a radioactive sample with a Geiger counter. She counts the disintegrations in 15 separate five-second intervals and gets the following numbers:

7, 11, 10, 7, 5, 7, 6, 12, 12, 7, 18, 12, 13, 12, 6.

(a) Her best estimate μ_{best} for the true mean count μ is the average of her 15 counts. (For a proof of this claim, see Problem 11.15.) What is her value for μ_{best}? **(b)** The standard deviation σ_ν of her 15 counts should be close to $\sqrt{\mu}$. What is her standard deviation and how does it compare with $\sqrt{\mu_{best}}$?

11.9. ★★ **(a)** Prove that the average value of ν^2 for the Poisson distribution $P_\mu(\nu)$ is $\overline{\nu^2} = \mu^2 + \mu$. [The easiest way to do this is probably to differentiate the identity (11.15) twice with respect to μ.] **(b)** Hence, prove that the standard deviation of ν is $\sigma_\nu = \sqrt{\mu}$. [Use the identity (11.7).]

11.10. ★★ The average rate of disintegrations from a certain radioactive sample is known to be roughly 20 per minute. If you wanted to measure this rate within 4%, for approximately how long would you plan to count?

11.11. ★★ Consider a counting experiment governed by the Poisson distribution $P_\mu(\nu)$, where the mean count μ is unknown, and suppose that I make a single count and get the value ν. Write down the probability $Prob(\nu)$ for getting this value ν. According to the principle of maximum likelihood, the best estimate for the unknown μ is that value of μ for which $Prob(\nu)$ is largest. Prove that the best estimate μ_{best} is precisely the observed count ν as you would expect. (In this calculation, the unknown μ is the variable and ν is the fixed value I obtained in my one experiment.)