

Geometric and Poisson Distributions

Recall our basic probabilistic experiment of tossing a biased coin n times. This is a very simple model, yet surprisingly powerful. Many important probability distributions that are widely used to model real-world phenomena can be derived from looking at this basic coin tossing model.

The first example is the Binomial(n, p) distribution, introduced earlier. This is the distribution of the number of Heads, S_n , in n tosses of a biased coin with probability p to be Heads. Recall that the distribution of S_n is $\mathbb{P}[S_n = k] = \binom{n}{k} p^k (1-p)^{n-k}$ for $k \in \{0, 1, \dots, n\}$. The expected value is $\mathbb{E}[S_n] = np$ and the variance is $\text{Var}(S_n) = np(1-p)$. The binomial distribution frequently appears to model the number of successes in a repeated experiment.

1 Geometric Distribution

Consider repeatedly tossing a biased coin with Heads probability p . Let X denote the number of tosses until the first Head appears. Then X is a random variable that takes values in \mathbb{Z}^+ , the set of positive integers. The event that $X = i$ is equal to the event of observing Tails for the first $i - 1$ tosses and getting Heads in the i th toss, which occurs with probability $(1-p)^{i-1}p$. Such a random variable is called a geometric random variable.

The geometric distribution frequently occurs in applications because we are often interested in how long we have to wait before a certain event happens: how many runs before the system fails, how many shots before one is on target, how many poll samples before we find a Democrat, how many retransmissions of a packet before successfully reaching the destination, etc.

Definition 19.1 (Geometric Distribution). *A random variable X for which*

$$\mathbb{P}[X = i] = (1-p)^{i-1}p, \quad \text{for } i = 1, 2, 3, \dots,$$

is said to have the geometric distribution with parameter p . This is abbreviated as $X \sim \text{Geometric}(p)$.

As a quick check, we can verify that the total probability of X is equal to 1:

$$\sum_{i=1}^{\infty} \mathbb{P}[X = i] = \sum_{i=1}^{\infty} (1-p)^{i-1}p = p \sum_{i=1}^{\infty} (1-p)^{i-1} = p \times \frac{1}{1-(1-p)} = 1,$$

where in the second-to-last step we have used the formula for geometric series.

If we plot the distribution of X (i.e., the values $\mathbb{P}[X = i]$ against i) we get a curve that decreases monotonically by a factor of $1-p$ at each step, as illustrated in Figure 1.

1.1 Mean and Variance of a Geometric Random Variable

Let us now compute the expectation $\mathbb{E}[X]$. Applying the definition of expected value directly gives us:

$$\mathbb{E}[X] = \sum_{i=1}^{\infty} i \times \mathbb{P}[X = i] = p \sum_{i=1}^{\infty} i(1-p)^{i-1}.$$




Figure 1: Illustration of the Geometric(p) distribution for $p = 0.1$ and $p = 0.3$.

However, the final summation is difficult to evaluate and requires some calculus trick. Instead, we will use the following alternative formula for expectation.

Theorem 19.1 (Tail Sum Formula). *Let X be a random variable that takes values in $\{0, 1, 2, \dots\}$. Then*

$$\mathbb{E}[X] = \sum_{i=1}^{\infty} \mathbb{P}[X \geq i].$$

Proof. For notational convenience, let's write $p_i = \mathbb{P}[X = i]$, for $i = 0, 1, 2, \dots$. From the definition of expectation, we have

$$\begin{aligned} \mathbb{E}[X] &= (0 \times p_0) + (1 \times p_1) + (2 \times p_2) + (3 \times p_3) + (4 \times p_4) + \dots \\ &= p_1 + (p_2 + p_2) + (p_3 + p_3 + p_3) + (p_4 + p_4 + p_4 + p_4) + \dots \\ &= (p_1 + p_2 + p_3 + p_4 + \dots) + (p_2 + p_3 + p_4 + \dots) + (p_3 + p_4 + \dots) + (p_4 + \dots) + \dots \\ &= \mathbb{P}[X \geq 1] + \mathbb{P}[X \geq 2] + \mathbb{P}[X \geq 3] + \mathbb{P}[X \geq 4] + \dots \end{aligned}$$

In the third line, we have regrouped the terms into convenient infinite sums, and each infinite sum is exactly the probability that $X \geq i$ for each i . You should check that you understand how the fourth line follows from the third.

Let us repeat the proof more formally, this time using more compact mathematical notation:

$$\mathbb{E}[X] = \sum_{j=1}^{\infty} j \times \mathbb{P}[X = j] = \sum_{j=1}^{\infty} \sum_{i=1}^j \mathbb{P}[X = j] = \sum_{i=1}^{\infty} \sum_{j=i}^{\infty} \mathbb{P}[X = j] = \sum_{i=1}^{\infty} \mathbb{P}[X \geq i],$$

where the third equality follows from interchanging the order of summations. □

We can now use Theorem 19.1 to compute $\mathbb{E}[X]$ more easily.

Theorem 19.2. *For $X \sim \text{Geometric}(p)$, we have $\mathbb{E}[X] = \frac{1}{p}$.*

Proof. The key observation is that for a geometric random variable X ,

$$\mathbb{P}[X \geq i] = (1 - p)^{i-1} \quad \text{for } i = 1, 2, \dots \quad (1)$$

We can obtain this simply by summing $\mathbb{P}[X = j]$ for $j \geq i$. Another way of seeing this is to note that the event “ $X \geq i$ ” means at least i tosses are required. This is equivalent to saying that the first $i - 1$ tosses are all Tails, and the probability of this event is precisely $(1 - p)^{i-1}$. Now, plugging equation (1) into Theorem 19.1, we get

$$\mathbb{E}[X] = \sum_{i=1}^{\infty} \mathbb{P}[X \geq i] = \sum_{i=1}^{\infty} (1 - p)^{i-1} = \frac{1}{1 - (1 - p)} = \frac{1}{p},$$

where we have used the formula for geometric series. □

So, the expected number of tosses of a biased coin until the first Head appears is $\frac{1}{p}$. Intuitively, if in each coin toss we expect to get p Heads, then we need to toss the coin $\frac{1}{p}$ times to get 1 Head. So for a fair coin, the expected number of tosses is 2, but remember that the actual number of coin tosses that we need can be any positive integers.

Remark: Another way of deriving $\mathbb{E}[X] = \frac{1}{p}$ is to use the interpretation of a geometric random variable X as the number of coin tosses until we get a Head. Consider what happens in the first coin toss: If the first toss comes up Heads, then $X = 1$. Otherwise, we have used one toss, and we repeat the coin tossing process again; the number of coin tosses after the first toss is again a geometric random variable with parameter p . Therefore, we can calculate:

$$\mathbb{E}[X] = \underbrace{p \cdot 1}_{\text{first toss is H}} + \underbrace{(1 - p) \cdot (1 + \mathbb{E}[X])}_{\text{first toss is T, then toss again}}.$$

Solving for $\mathbb{E}[X]$ yields $\mathbb{E}[X] = \frac{1}{p}$, as claimed.

Let us now compute the variance of X .

Theorem 19.3. For $X \sim \text{Geometric}(p)$, we have $\text{Var}(X) = \frac{1-p}{p^2}$.

Proof. We will show that $\mathbb{E}[X(X - 1)] = \frac{2(1-p)}{p^2}$. Since we already know $\mathbb{E}[X] = \frac{1}{p}$, this will imply the desired result:

$$\begin{aligned} \text{Var}(X) &= \mathbb{E}[X^2] - (\mathbb{E}[X])^2 = \mathbb{E}[X(X - 1)] + \mathbb{E}[X] - (\mathbb{E}[X])^2 \\ &= \frac{2(1-p)}{p^2} + \frac{1}{p} - \frac{1}{p^2} = \frac{2(1-p) + p - 1}{p^2} = \frac{1-p}{p^2}. \end{aligned}$$

Now to show $\mathbb{E}[X(X - 1)] = \frac{2(1-p)}{p^2}$, we begin with the following identity of geometric series:

$$\sum_{i=0}^{\infty} (1 - p)^i = \frac{1}{p}.$$

Differentiating the identity above with respect to p yields (the $i = 0$ term is equal to 0 so we omit it):

$$-\sum_{i=1}^{\infty} i(1 - p)^{i-1} = -\frac{1}{p^2}.$$

Differentiating both sides with respect to p again gives us (the $i = 1$ term is equal to 0 so we omit it):

$$\sum_{i=2}^{\infty} i(i - 1)(1 - p)^{i-2} = \frac{2}{p^3}. \tag{2}$$

Now using the geometric distribution of X and identity (2), we can calculate:

$$\begin{aligned}
 \mathbb{E}[X(X-1)] &= \sum_{i=1}^{\infty} i(i-1) \times \mathbb{P}[X=i] \\
 &= \sum_{i=2}^{\infty} i(i-1)(1-p)^{i-1} p && \text{(the } i=1 \text{ term is equal to 0 so we omit it)} \\
 &= p(1-p) \sum_{i=2}^{\infty} i(i-1)(1-p)^{i-2} \\
 &= p(1-p) \times \frac{2}{p^3} && \text{(using identity (2))} \\
 &= \frac{2(1-p)}{p^2},
 \end{aligned}$$

as desired. □

1.2 Memoryless Property

Geometric distributions have a property called *memorylessness*¹. To get some intuition on this property, suppose we revisit our earlier motivational example of repeatedly tossing a biased coin, with probability of heads p . Here, let X denote the number of tosses until the first head appears, and as such $X \sim \text{Geometric}(p)$.

As we've found, the probability that we need *more than* n tosses before getting the first head is $\mathbb{P}[X > n] = (1-p)^n$, since the first n tosses must have been tails.

What if we've already tossed the coin m times, without getting a heads? What is the probability that we need more than n *additional* tosses before getting our first head? This probability turns out to be

$$\begin{aligned}
 \mathbb{P}[X > n+m \mid X > m] &= \frac{\mathbb{P}[X > n+m]}{\mathbb{P}[X > m]} \\
 &= \frac{(1-p)^{n+m}}{(1-p)^m} \\
 &= (1-p)^n \\
 &= \mathbb{P}[X > n]
 \end{aligned}$$

This gives us the formal expression of the memoryless property of geometric distributions:

$$\mathbb{P}[X > n+m \mid X > m] = \mathbb{P}[X > n].$$

Intuitively, the condition that $X > m$ implies the first m trials are tails, so for $X > n+m$, the next n trials must be tails which corresponds to $\mathbb{P}[X > n]$. In this sense, the geometric distribution is *memoryless*; previous trials do not affect the number of additional trials we need.

1.3 Application: The Coupon Collector's Problem

Suppose we are trying to collect a set of n different baseball cards. We get the cards by buying boxes of cereal: each box contains exactly one card, and it is equally likely to be any of the n cards. How many boxes do we need to buy until we have collected at least one copy of every card?

¹This property is shared by *exponential distributions*, the continuous analog of the geometric distribution, which we'll cover in a future note.

Let S_n denote the number of boxes we need to buy in order to collect all n cards. The distribution of S_n is difficult to compute directly (try it for $n = 3$). But if we are only interested in its expected value $\mathbb{E}[S_n]$, then we can evaluate it easily using linearity of expectation and what we have just learned about the geometric distribution.

As usual, we start by writing

$$S_n = X_1 + X_2 + \cdots + X_n \quad (3)$$

for suitable simple random variables X_i . What should the X_i be? Naturally, X_i is the number of boxes we buy while trying to get the i -th new card (starting immediately after we have gotten the $(i - 1)$ st new card). With this definition, make sure you believe equation (3) before proceeding.

What does the distribution of X_i look like? Well, X_1 is trivial: no matter what happens, we always get a new card in the first box (since we have none to start with). So $\mathbb{P}[X_1 = 1] = 1$, and thus $\mathbb{E}[X_1] = 1$.

How about X_2 ? Each time we buy a box, we will get the same old card with probability $\frac{1}{n}$, and a new card with probability $\frac{n-1}{n}$. So we can think of buying boxes as flipping a biased coin with Heads probability $p = \frac{n-1}{n}$; then X_2 is just the number of tosses until the first Head appears. So X_2 has the geometric distribution with parameter $p = \frac{n-1}{n}$, and

$$\mathbb{E}[X_2] = \frac{n}{n-1}.$$

How about X_3 ? This is very similar to X_2 except that now we only get a new card with probability $\frac{n-2}{n}$ (since there are now two old ones). So X_3 has the geometric distribution with parameter $p = \frac{n-2}{n}$, and

$$\mathbb{E}[X_3] = \frac{n}{n-2}.$$

Arguing in the same way, we see that, for $i = 1, 2, \dots, n$, X_i has the geometric distribution with parameter $p = \frac{n-i+1}{n}$, and hence that

$$\mathbb{E}[X_i] = \frac{n}{n-i+1}.$$

Finally, applying linearity of expectation to equation (3), we get

$$\mathbb{E}[S_n] = \sum_{i=1}^n \mathbb{E}[X_i] = \frac{n}{n} + \frac{n}{n-1} + \cdots + \frac{n}{2} + \frac{n}{1} = n \sum_{i=1}^n \frac{1}{i}. \quad (4)$$

This is an exact expression for $\mathbb{E}[S_n]$. We can obtain a tidier form by noting that the sum in (4) actually has a very good approximation,² namely:

$$\sum_{i=1}^n \frac{1}{i} \approx \ln n + \gamma_E,$$

where $\gamma_E = 0.5772\dots$ is known as *Euler's constant*.

Thus, the expected number of cereal boxes needed to collect n cards is about $n(\ln n + \gamma)$. This is an excellent approximation to the exact formula (4) even for quite small values of n . So for example, for $n = 100$, we expect to buy about 518 boxes.

2 Poisson Distribution

Consider the number of clicks of a Geiger counter, which measures radioactive emissions. The average number of such clicks per unit time, λ , is a measure of radioactivity, but the actual number of clicks fluctuates

²This is another of the little tricks you might like to carry around in your toolbox.

according to a certain distribution called the Poisson distribution. What is remarkable is that the average value, λ , completely determines the probability distribution on the number of clicks X .

Definition 19.2 (Poisson distribution). *A random variable X for which*

$$\mathbb{P}[X = i] = \frac{\lambda^i}{i!} e^{-\lambda}, \quad \text{for } i = 0, 1, 2, \dots \quad (5)$$

is said to have the Poisson distribution with parameter λ . This is abbreviated as $X \sim \text{Poisson}(\lambda)$.

To make sure this is a valid definition, let us check that (5) is in fact a distribution, i.e., that the probabilities sum to 1. We have

$$\sum_{i=0}^{\infty} \frac{\lambda^i}{i!} e^{-\lambda} = e^{-\lambda} \sum_{i=0}^{\infty} \frac{\lambda^i}{i!} = e^{-\lambda} \times e^{\lambda} = 1.$$

In the second-to-last step, we used the Taylor series expansion $e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots$.

The Poisson distribution is also a very widely accepted model for so-called “rare events,” such as disconnected phone calls, radioactive emissions, crossovers in chromosomes, the number of cases of disease, the number of births per hour, etc. This model is appropriate whenever the occurrences can be assumed to happen randomly with some constant density in a continuous region (of time or space), such that occurrences in disjoint subregions are independent. One can then show that the number of occurrences in a region of unit size should obey the Poisson distribution with parameter λ .

Example

Suppose when we write an article, we make an average of 1 typo per page. We can model this with a Poisson random variable X with $\lambda = 1$. So the probability that a page has 5 typos is

$$\mathbb{P}[X = 5] = \frac{1^5}{5!} e^{-1} = \frac{1}{120e} \approx \frac{1}{326}.$$

Now suppose the article has 200 pages. If we assume the number of typos in each page is independent, then the probability that there is at least one page with exactly 5 typos is

$$\begin{aligned} \mathbb{P}[\exists \text{ a page with exactly 5 typos}] &= 1 - \mathbb{P}[\text{every page has } \neq 5 \text{ typos}] \\ &= 1 - \prod_{k=1}^{200} \mathbb{P}[\text{page } k \text{ has } \neq 5 \text{ typos}] \\ &= 1 - \prod_{k=1}^{200} (1 - \mathbb{P}[\text{page } k \text{ has exactly 5 typos}]) \\ &= 1 - \left(1 - \frac{1}{120e}\right)^{200}, \end{aligned}$$

where in the last step we have used our earlier calculation for $\mathbb{P}[X = 5]$. □

2.1 Mean and Variance of a Poisson Random Variable

Let us now calculate the expectation and variance of a Poisson random variable.

Theorem 19.4. *For a Poisson random variable $X \sim \text{Poisson}(\lambda)$, we have $\mathbb{E}[X] = \lambda$ and $\text{Var}(X) = \lambda$.*

Proof. We can calculate $\mathbb{E}[X]$ directly from the definition of expectation:

$$\begin{aligned}
 \mathbb{E}[X] &= \sum_{i=0}^{\infty} i \times \mathbb{P}[X = i] \\
 &= \sum_{i=1}^{\infty} i \times \frac{\lambda^i}{i!} e^{-\lambda} && \text{(the } i = 0 \text{ term is equal to 0 so we omit it)} \\
 &= \lambda e^{-\lambda} \sum_{i=1}^{\infty} \frac{\lambda^{i-1}}{(i-1)!} \\
 &= \lambda e^{-\lambda} e^{\lambda} && \text{(since } e^{\lambda} = \sum_{j=0}^{\infty} \frac{\lambda^j}{j!} \text{ with } j = i - 1) \\
 &= \lambda.
 \end{aligned}$$

Similarly, we can calculate $\mathbb{E}[X(X - 1)]$ as follows:

$$\begin{aligned}
 \mathbb{E}[X(X - 1)] &= \sum_{i=0}^{\infty} i(i - 1) \times \mathbb{P}[X = i] \\
 &= \sum_{i=2}^{\infty} i(i - 1) \frac{\lambda^i}{i!} e^{-\lambda} && \text{(the } i = 0 \text{ and } i = 1 \text{ terms are equal to 0 so we omit them)} \\
 &= \lambda^2 e^{-\lambda} \sum_{i=2}^{\infty} \frac{\lambda^{i-2}}{(i-2)!} \\
 &= \lambda^2 e^{-\lambda} e^{\lambda} && \text{(since } e^{\lambda} = \sum_{j=0}^{\infty} \frac{\lambda^j}{j!} \text{ with } j = i - 2) \\
 &= \lambda^2.
 \end{aligned}$$

Therefore,

$$\text{Var}(X) = \mathbb{E}[X^2] - \mathbb{E}[X]^2 = \mathbb{E}[X(X - 1)] + \mathbb{E}[X] - \mathbb{E}[X]^2 = \lambda^2 + \lambda - \lambda^2 = \lambda,$$

as desired. □

A plot of the Poisson distribution reveals a curve that rises monotonically to a single peak and then decreases monotonically. The peak is as close as possible to the expected value, i.e., at $i = \lfloor \lambda \rfloor$. Figure 2 illustrates the $\text{Poisson}(\lambda)$ distribution for $\lambda = 1, 2, 5, 20$.

2.2 Sum of Independent Poisson Random Variables

As illustrated in Figure 2, the $\text{Poisson}(\lambda)$ distribution becomes more symmetric and resembles a “bell curve” as λ increases. As we will learn later, this phenomenon is closely related to the following useful fact regarding a sum of independent Poisson random variables.

Theorem 19.5. *Let $X \sim \text{Poisson}(\lambda)$ and $Y \sim \text{Poisson}(\mu)$ be independent Poisson random variables. Then, $X + Y \sim \text{Poisson}(\lambda + \mu)$.*

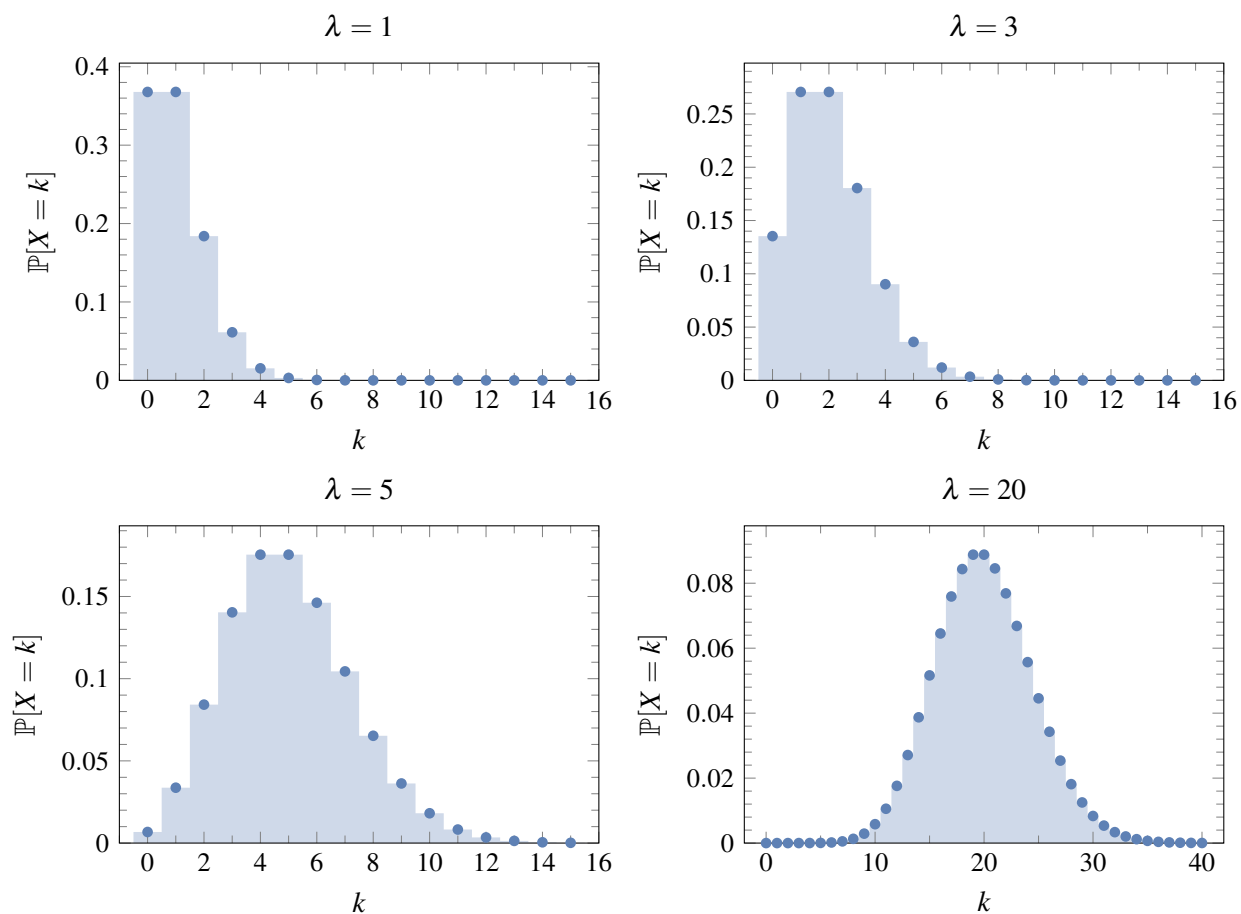


Figure 2: Illustration of the Poisson(λ) distribution for $\lambda = 1, 2, 5, 20$.

Proof. For all $k = 0, 1, 2, \dots$, we have

$$\begin{aligned}
 \mathbb{P}[X + Y = k] &= \sum_{j=0}^k \mathbb{P}[X = j, Y = k - j] \\
 &= \sum_{j=0}^k \mathbb{P}[X = j] \mathbb{P}[Y = k - j] \\
 &= \sum_{j=0}^k \frac{\lambda^j}{j!} e^{-\lambda} \frac{\mu^{k-j}}{(k-j)!} e^{-\mu} \\
 &= e^{-(\lambda+\mu)} \frac{1}{k!} \sum_{j=0}^k \frac{k!}{j!(k-j)!} \lambda^j \mu^{k-j} \\
 &= e^{-(\lambda+\mu)} \frac{(\lambda + \mu)^k}{k!},
 \end{aligned}$$

where the second equality follows from independence, and the last equality from the binomial theorem. \square

By induction, we conclude that if X_1, X_2, \dots, X_n are independent Poisson random variables with parameters $\lambda_1, \lambda_2, \dots, \lambda_n$, respectively, then

$$X_1 + X_2 + \dots + X_n \sim \text{Poisson}(\lambda_1 + \lambda_2 + \dots + \lambda_n).$$

2.3 Poisson as a limit of Binomial

To see a concrete example of how Poisson distribution arises, suppose we want to model the number of cell phone users initiating calls in a network during a time period, of duration (say) 1 minute. There are many customers in the network, and all of them can potentially make a call during this time period. However, only a very small fraction of them actually will. Under this scenario, it seems reasonable to make two assumptions:

- The probability of having more than 1 customer initiating a call in any small time interval is negligible.
- The initiation of calls in disjoint time intervals are independent events.

Then if we divide the one-minute time period into n disjoint intervals, then the number of calls X in that time period can be modeled as a binomial random variable with parameter n and probability of success p , i.e., p is the probability of having a call initiated in a time interval of length $1/n$. But what should p be in terms of the relevant parameters of the problem? If calls are initiated at an average rate of λ calls per minute, then $\mathbb{E}[X] = \lambda$ and so $np = \lambda$, i.e., $p = \lambda/n$. So $X \sim \text{Binomial}(n, \frac{\lambda}{n})$. As we shall see below, as we let n tend to infinity, this distribution tends to the Poisson distribution with parameter λ . We can also see why the Poisson distribution is a model for “rare events.” We are thinking of it as a sequence of a large number, n , of coin flips, where we expect only a finite number λ of Heads.

Now we will prove that $\text{Poisson}(\lambda)$ is the limit of $\text{Binomial}(n, \frac{\lambda}{n})$, as n tends to infinity.

Theorem 19.6. *Let $X \sim \text{Binomial}(n, \frac{\lambda}{n})$ where $\lambda > 0$ is a fixed constant. Then for every $i = 0, 1, 2, \dots$,*

$$\mathbb{P}[X = i] \longrightarrow \frac{\lambda^i}{i!} e^{-\lambda} \quad \text{as } n \rightarrow \infty.$$

That is, the probability distribution of X converges to the Poisson distribution with parameter λ .

Proof. Fix $i \in \{0, 1, 2, \dots\}$, and assume $n \geq i$ (because we will let $n \rightarrow \infty$). Then, because X has binomial distribution with parameter n and $p = \frac{\lambda}{n}$,

$$\mathbb{P}[X = i] = \binom{n}{i} p^i (1-p)^{n-i} = \frac{n!}{i!(n-i)!} \left(\frac{\lambda}{n}\right)^i \left(1 - \frac{\lambda}{n}\right)^{n-i}.$$

Let us collect the factors into

$$\mathbb{P}[X = i] = \frac{\lambda^i}{i!} \left(\frac{n!}{(n-i)!} \cdot \frac{1}{n^i}\right) \cdot \left(1 - \frac{\lambda}{n}\right)^n \cdot \left(1 - \frac{\lambda}{n}\right)^{-i}. \quad (6)$$

The first parenthesis above becomes, as $n \rightarrow \infty$,

$$\frac{n!}{(n-i)!} \cdot \frac{1}{n^i} = \frac{n \cdot (n-1) \cdots (n-i+1) \cdot (n-i)!}{(n-i)!} \cdot \frac{1}{n^i} = \frac{n}{n} \cdot \frac{(n-1)}{n} \cdots \frac{(n-i+1)}{n} \rightarrow 1.$$

From calculus, the second parenthesis in (6) becomes, as $n \rightarrow \infty$,

$$\left(1 - \frac{\lambda}{n}\right)^n \rightarrow e^{-\lambda}.$$

And since i is fixed, the third parenthesis in (6) becomes, as $n \rightarrow \infty$,

$$\left(1 - \frac{\lambda}{n}\right)^{-i} \rightarrow (1-0)^{-i} = 1.$$

Substituting these results back to (6) gives us

$$\mathbb{P}[X = i] \rightarrow \frac{\lambda^i}{i!} \cdot 1 \cdot e^{-\lambda} \cdot 1 = \frac{\lambda^i}{i!} e^{-\lambda},$$

as desired. □