

Continuous Probability Distributions

Up to now we have focused exclusively on *discrete* sample spaces Ω , where the number of sample points $\omega \in \Omega$ is either finite or countably infinite (such as the integers). As a consequence, we have only been able to talk about *discrete* random variables, which take on only a finite or countably infinite number of values.

But in real life many quantities that we wish to model probabilistically are *real-valued*; examples include the position of a particle in a box, the time at which a certain incident happens, or the direction of travel of a meteorite. In this note, we discuss how to extend the concepts we've seen in the discrete setting to this *continuous* setting. As we shall see, everything translates in a natural way once we have set up the right framework. The framework involves some elementary calculus but (at this level of discussion) nothing too daunting.

1 Continuous Uniform Probability Space

Suppose we spin a “wheel of fortune” and record the position of the pointer on the outer circumference of the wheel. Assuming that the circumference is of length ℓ and that the wheel is unbiased, the position is presumably equally likely to take on any value in the real interval $[0, \ell]$.¹ How do we model this experiment using a probability space?

Consider for a moment the analogous discrete setting, where the pointer can stop only at a finite number m of positions distributed evenly around the wheel. (If m is very large, then this is in some sense similar to the continuous setting, which we can think of as the limit $m \rightarrow \infty$.) Then we would model this situation using the discrete sample space $\Omega = \{0, \frac{\ell}{m}, \frac{2\ell}{m}, \dots, \frac{(m-1)\ell}{m}\}$, with uniform probabilities $\mathbb{P}[\omega] = \frac{1}{m}$ for each $\omega \in \Omega$.

In the continuous setting, however, we get into trouble if we try the same approach. If we let ω range over all real numbers in $\Omega = [0, \ell]$, what value should we assign to each $\mathbb{P}[\omega]$? By uniformity, this probability should be the same for all ω . But if we assign $\mathbb{P}[\omega]$ to be any positive value, then because there are infinitely many ω in Ω , the sum of all probabilities $\mathbb{P}[\omega]$ will be ∞ ! Thus, $\mathbb{P}[\omega]$ must be zero for all $\omega \in \Omega$. But if all of our sample points have probability zero, then we are unable to assign meaningful probabilities to any events!

To resolve this problem, consider instead any *interval* $[a, b] \subseteq [0, \ell]$, where $b > a$. Can we assign a non-zero probability value to this interval? Since the total probability assigned to $[0, \ell]$ must be 1, and since we want our probability to be uniform, the logical value for the probability of interval $[a, b]$ is

$$\frac{\text{length of } [a, b]}{\text{length of } [0, \ell]} = \frac{b - a}{\ell}.$$

In other words, the probability of an interval is proportional to its length.

Note that intervals are subsets of the sample space Ω and are therefore events. So in contrast to discrete probability, where we assigned probabilities to *points* in the sample space, in continuous probability we

¹As we shall see shortly, it doesn't matter whether we consider the closed interval $[0, \ell]$ or the half-open interval $[0, \ell)$.

are assigning probabilities to certain basic *events* (in this case intervals). What about probabilities of other events? By specifying the probability of intervals, we have also specified the probability of any event E which can be written as the disjoint union of (a finite or countably infinite number of) intervals, $E = \cup_i E_i$. For then we can write $\mathbb{P}[E] = \sum_i \mathbb{P}[E_i]$, in analogous fashion to the discrete case. Thus for example the probability that the pointer ends up in the first or third quadrants of the wheel is $\frac{\ell/4}{\ell} + \frac{\ell/4}{\ell} = \frac{1}{2}$. For all practical purposes, such events are all we really need.²

2 Continuous Random Variables

Recall that in the discrete setting we typically work with *random variables* and their distributions, rather than directly with probability spaces and events. The simplest example of a continuous random variable is the position X of the pointer in the wheel of fortune, as discussed above. This random variable has the *uniform* distribution on $[0, \ell]$. How, precisely, should we define the distribution of a continuous random variable? In the discrete case the distribution of a r.v. X is described by specifying, for each possible value a , the probability $\mathbb{P}[X = a]$. But for the r.v. X corresponding to the position of the pointer, we have $\mathbb{P}[X = a] = 0$ for every a , so we run into the same problem as we encountered above in defining the probability space.

The resolution is the same: instead of specifying $\mathbb{P}[X = a]$, we specify $\mathbb{P}[a \leq X \leq b]$ for all intervals $[a, b]$.³ To do this formally, we need to introduce the concept of a *probability density function* (sometimes referred to just as a “density”, or a “p.d.f.”).

Definition 20.1 (Probability Density Function). A probability density function (*p.d.f.*) for a real-valued random variable X is a function $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying:

1. f is non-negative: $f(x) \geq 0$ for all $x \in \mathbb{R}$.
2. The total integral of f is equal to 1: $\int_{-\infty}^{\infty} f(x) \, dx = 1$.

Then the distribution of X is given by:

$$\mathbb{P}[a \leq X \leq b] = \int_a^b f(x) \, dx \quad \text{for all } a < b.$$

Let us examine this definition. Note that the definite integral is just the area under the curve f between the values a and b . Thus f plays a similar role to the “histogram” we sometimes draw to picture the distribution of a discrete random variable. The first condition that f be non-negative ensures that the probability of every event is non-negative. The second condition that the total integral of f equal to 1 ensures that it defines a valid probability distribution, because the r.v. X must take on real values:

$$\mathbb{P}[X \in \mathbb{R}] = \mathbb{P}[-\infty < X < \infty] = \int_{-\infty}^{\infty} f(x) \, dx = 1. \tag{1}$$

For example, consider the wheel-of-fortune r.v. X , which has uniform distribution on the interval $[0, \ell]$. This means the density f of X vanishes outside this interval: $f(x) = 0$ for $x < 0$ and for $x > \ell$. Within the interval $[0, \ell]$ we want the distribution of X to be uniform, which means we should take f to be a constant $f(x) = c$ for $0 \leq x \leq \ell$. The value of c is determined by the requirement in (1) that the total area under f is 1:

$$1 = \int_{-\infty}^{\infty} f(x) \, dx = \int_0^{\ell} c \, dx = c\ell,$$

²A formal treatment of which events can be assigned a well-defined probability requires a discussion of *measure theory*, which is beyond the scope of this course.

³Note that it does not matter whether or not we include the endpoints a, b ; since $\mathbb{P}[X = a] = \mathbb{P}[X = b] = 0$, we have $\mathbb{P}[a < X < b] = \mathbb{P}[a \leq X \leq b]$.

which gives us $c = \frac{1}{\ell}$. Therefore, the density of the uniform distribution on $[0, \ell]$ is given by

$$f(x) = \begin{cases} 0, & \text{for } x < 0, \\ 1/\ell, & \text{for } 0 \leq x \leq \ell, \\ 0, & \text{for } x > \ell. \end{cases}$$

Remark: Following the “histogram” analogy above, it is tempting to think of $f(x)$ as a “probability.” However, $f(x)$ doesn’t itself correspond to the probability of anything! In particular, there is no requirement that $f(x)$ be bounded by 1. For example, the density of the uniform distribution on the interval $[0, \ell]$ with $\ell = \frac{1}{2}$ is equal to $f(x) = 1/(\frac{1}{2}) = 2$ for $0 \leq x \leq \frac{1}{2}$, which is greater than 1. To connect density $f(x)$ with probabilities, we need to look at a very small interval $[x, x + dx]$ close to x ; then we have

$$\mathbb{P}[x \leq X \leq x + dx] = \int_x^{x+dx} f(z) dz \approx f(x) dx. \quad (2)$$

Thus, we can interpret $f(x)$ as the “probability per unit length” in the vicinity of x .

2.1 Cumulative Distribution Function

For a continuous random variable X , one often starts the discussion with the *cumulative distribution function* (*c.d.f.*), which is the function $F(x) = \mathbb{P}[X \leq x]$. It is closely related to the probability density function for X , $f(x)$, as

$$F(x) = \mathbb{P}[X \leq x] = \int_{-\infty}^x f(z) dz. \quad (3)$$

Thus, one can describe a random variable X by its c.d.f., denoted by $F(x)$, which then gives the probability density function, $f(x)$, as

$$f(x) = \frac{dF(x)}{dx}.$$

To connect to discrete probability, one might think of approximating a continuous random variable, X , as the set of probabilities for X being in one of a countably infinite set of intervals of length dx on the real line. That is, the set of probabilities $\mathbb{P}[x_k < X \leq x_k + dx]$ where $x_k = k dx$ for $k \in \mathbb{Z}$. In this view, $\mathbb{P}[X \leq x_i] = \sum_{j \leq i} \mathbb{P}[x_j < X \leq x_j + dx]$. Connecting this to the probability density function, $f(x)$, we have

$$\mathbb{P}[x_j < X \leq x_j + dx] \approx f(x_j) dx$$

for “small” dx , and

$$F(x_i) = \mathbb{P}[X \leq x_i] \approx \sum_{j \leq i} f(x_j) dx.$$

Calculus comes in when we see the expression above is a Riemann sum. Taking the limit as dx goes to zero yields the integral we see in (3).

2.2 Expectation and Variance

As in the discrete case, we define the expectation of a continuous r.v. as follows:

Definition 20.2 (Expectation). *The expectation of a continuous r.v. X with probability density function f is*

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} xf(x) dx.$$

Note that the integral plays the role of the summation in the discrete formula $\mathbb{E}[X] = \sum_a a \mathbb{P}[X = a]$. Similarly, we can define the variance as follows:

Definition 20.3 (Variance). *The variance of a continuous r.v. X with probability density function f is*

$$\text{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2 = \int_{-\infty}^{\infty} x^2 f(x) dx - \left(\int_{-\infty}^{\infty} x f(x) dx \right)^2.$$

Example

Let X be a uniform r.v. on the interval $[0, \ell]$. Then intuitively, its expected value should be in the middle, $\frac{\ell}{2}$. Indeed, we can use our definition above to compute

$$\mathbb{E}[X] = \int_0^{\ell} x \frac{1}{\ell} dx = \frac{x^2}{2\ell} \Big|_0^{\ell} = \frac{\ell}{2},$$

as claimed. We can also calculate its variance using the above definition and plugging in the value $\mathbb{E}[X] = \frac{\ell}{2}$ to get:

$$\text{Var}(X) = \int_0^{\ell} x^2 \frac{1}{\ell} dx - \mathbb{E}[X]^2 = \frac{x^3}{3\ell} \Big|_0^{\ell} - \left(\frac{\ell}{2} \right)^2 = \frac{\ell^2}{3} - \frac{\ell^2}{4} = \frac{\ell^2}{12}.$$

The factor of $\frac{1}{12}$ here is not particularly intuitive, but the fact that the variance is proportional to ℓ^2 should come as no surprise. Like its discrete counterpart, this distribution has large variance.

2.3 Joint Density

Recall that for discrete random variables X and Y , their joint distribution is specified by the probabilities $\mathbb{P}[X = a, Y = c]$ for all possible values a, c . Similarly, if X and Y are continuous random variables, then their joint distribution is specified by the probabilities $\mathbb{P}[a \leq X \leq b, c \leq Y \leq d]$ for all $a \leq b, c \leq d$. Moreover, just as the distribution of X can be characterized by its density function, the joint distribution of X and Y can be characterized by their joint density.

Definition 20.4 (Joint Density). *A joint density function for two random variable X and Y is a function $f_{X,Y} : \mathbb{R}^2 \rightarrow \mathbb{R}$ satisfying:*

1. $f_{X,Y}$ is non-negative: $f_{X,Y}(x,y) \geq 0$ for all $x,y \in \mathbb{R}$.
2. The total integral of $f_{X,Y}$ is equal to 1: $\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{X,Y}(x,y) dx dy = 1$.

Then the joint distribution of X and Y is given by:

$$\mathbb{P}[a \leq X \leq b, c \leq Y \leq d] = \int_c^d \int_a^b f_{X,Y}(x,y) dx dy \quad \text{for all } a \leq b \text{ and } c \leq d.$$

In analogy with (2), we can connect the joint density $f_{X,Y}(x,y)$ with probabilities by looking at a very small square $[x, x+dx] \times [y, y+dy]$ close to (x,y) ; then we have

$$\mathbb{P}[x \leq X \leq x+dx, y \leq Y \leq y+dy] = \int_y^{y+dy} \int_x^{x+dx} f_{X,Y}(u,v) du dv \approx f_{X,Y}(x,y) dx dy. \quad (4)$$

Thus we can interpret $f_{X,Y}(x,y)$ as the “probability per unit area” in the vicinity of (x,y) .

As with joint distributions for discrete random variables, we can define marginal and conditional density functions for continuous random variables as follows:

Definition 20.5 (Marginal and Conditional Density). Consider a joint distribution on X and Y with joint density function $f_{X,Y}(x,y)$.

1. The marginal density function f_X of X is defined as

$$f_X(x) = \int_{-\infty}^{\infty} f_{X,Y}(x,y) dy,$$

and similarly the marginal density function f_Y of Y is defined as

$$f_Y(y) = \int_{-\infty}^{\infty} f_{X,Y}(x,y) dx.$$

2. The conditional density function $f_{Y|X=x}$ of Y given $X = x$ is defined as

$$f_{Y|X=x}(y) = \frac{f_{X,Y}(x,y)}{f_X(x)}.$$

The integration above for the marginal density function converts the joint density function on X and Y to the density function for only X (or Y).

The conditional density function $f_{Y|X=x}(y)$ is the density function for Y conditioned on the event $X = x$. Recall that the probability of event B given A is computed as $\mathbb{P}[B | A] = \frac{\mathbb{P}[A \cap B]}{\mathbb{P}[A]}$. For continuous random variables X and Y , the event B corresponds to the event $Y \in [y, y + dy]$ and the event A is $X \in [x, x + dx]$. Using (4), we can plug in $\mathbb{P}[A \cap B] \approx f_{X,Y}(x,y) dx dy$ and $\mathbb{P}[A] \approx f_X(x) dx$, which yields

$$\mathbb{P}[Y \in [y, y + dy] | X \in [x, x + dx]] \approx \frac{f_{X,Y}(x,y) dx dy}{f_X(x) dx} \approx \frac{f_{X,Y}(x,y)}{f_X(x)} dy.$$

Dividing by dy produces conditional density of Y conditioned on $X = x$. This last bit can be made formal using limits.

2.4 Independence

Recall that two discrete random variables X and Y are said to be independent if the events $X = a$ and $Y = c$ are independent for every possible values a, c . We have a similar definition for continuous random variables:

Definition 20.6 (Independence for Continuous R.V.'s). Two continuous r.v.'s X, Y are independent if the events $a \leq X \leq b$ and $c \leq Y \leq d$ are independent for all $a \leq b$ and $c \leq d$:

$$\mathbb{P}[a \leq X \leq b, c \leq Y \leq d] = \mathbb{P}[a \leq X \leq b] \cdot \mathbb{P}[c \leq Y \leq d].$$

What does this definition say about the joint density of independent r.v.'s X and Y ? Applying (4) to connect the joint density with probabilities, we get, for small dx and dy :

$$\begin{aligned} f_{X,Y}(x,y) dx dy &\approx \mathbb{P}[x \leq X \leq x + dx, y \leq Y \leq y + dy] \\ &= \mathbb{P}[x \leq X \leq x + dx] \cdot \mathbb{P}[y \leq Y \leq y + dy] && \text{(by independence)} \\ &\approx f_X(x) dx \times f_Y(y) dy \\ &= f_X(x) f_Y(y) dx dy, \end{aligned}$$

where f_X and f_Y are the (marginal) densities of X and Y respectively. So we get the following result:

Theorem 20.1. *The joint density of independent r.v.'s X and Y is the product of the marginal densities:*

$$f_{X,Y}(x,y) = f_X(x)f_Y(y) \quad \text{for all } x,y \in \mathbb{R}.$$

Exercise. Show that for independent r.v.'s X and Y that $f_{Y|X=x}(y) = f_Y(y)$.

2.5 Total Probability

Equipped with the knowledge of conditional densities, we can define the *total probability rule* similarly to discrete probability.

Recall that in discrete probability, we defined the total probability rule for an event A as

$$\mathbb{P}[A] = \sum_x \mathbb{P}[A | X = x] \mathbb{P}[X = x].$$

Working with a continuous random variable instead, the total probability rule is defined as

$$\mathbb{P}[A] = \int_{-\infty}^{\infty} \mathbb{P}[A | X = x] f_X(x) dx.$$

Suppose now that the event A is defined as $y \leq Y \leq y + dy$, where Y is another continuous random variable. Recall that the probability $\mathbb{P}[y \leq Y \leq y + dy]$ can be approximated as $f_Y(y) dy$; the conditional probability $\mathbb{P}[y \leq Y \leq y + dy | X = x]$ can similarly be approximated as $f_{Y|X=x}(y) dy$. Hence, we have

$$\begin{aligned} f_Y(y) dy &\approx \mathbb{P}[y \leq Y \leq y + dy] \\ &= \int_{-\infty}^{\infty} \mathbb{P}[y \leq Y \leq y + dy | X = x] f_X(x) dx \\ &\approx \int_{-\infty}^{\infty} (f_{Y|X=x}(y) dy) f_X(x) dx, \\ f_Y(y) dy &\approx \left(\int_{-\infty}^{\infty} f_{Y|X=x}(y) f_X(x) dx \right) dy, \\ f_Y(y) &= \int_{-\infty}^{\infty} f_{Y|X=x}(y) f_X(x) dx. \end{aligned}$$

This is the continuous analog to the total probability rule for discrete random variables. Note that this can also be viewed as computing the marginal density by utilizing the definition of a conditional density:

$$f_Y(y) = \int_{-\infty}^{\infty} f_{X,Y}(x,y) dx = \int_{-\infty}^{\infty} f_{Y|X=x}(y) f_X(x) dx.$$

3 Exponential Distribution

We have already seen one important continuous distribution, namely the uniform distribution. In this and the next sections we will see two more: the *exponential* distribution and the *normal* (or *Gaussian*) distribution. These three distributions cover the vast majority of continuous random variables arising in applications.

The exponential distribution is a continuous version of the geometric distribution, which we have already seen. Recall that the geometric distribution describes the number of tosses of a coin until the first Head appears; the distribution has a single parameter p , which is the bias (Heads probability) of the coin. Of




Figure 1: The probability density function $f(x)$ for the exponential distribution with $\lambda = 0.5, 1, 2$.

course, in real life applications we are usually not waiting for a coin to come up Heads but rather waiting for a system to fail, a clock to ring, an experiment to succeed, etc.

In such applications we are frequently not dealing with discrete events or discrete time, but rather with *continuous* time: for example, if we are waiting for an apple to fall off a tree, it can do so at any time at all, not necessarily on the tick of a discrete clock. This situation is naturally modeled by the exponential distribution, defined as follows.

Definition 20.7 (Exponential Distribution). For $\lambda > 0$, a continuous random variable X with p.d.f.

$$f(x) = \begin{cases} \lambda e^{-\lambda x}, & \text{if } x \geq 0, \\ 0, & \text{otherwise,} \end{cases}$$

is called an exponential random variable with parameter λ , and we write $X \sim \text{Exp}(\lambda)$.

Note that by definition $f(x)$ is non-negative. Moreover, we can check that it satisfies (1):

$$\int_{-\infty}^{\infty} f(x) \, dx = \int_0^{\infty} \lambda e^{-\lambda x} \, dx = -e^{-\lambda x} \Big|_0^{\infty} = 1,$$

so $f(x)$ is indeed a valid probability density function. Figure 1 shows the probability density function for the exponential distribution with a few different values of λ .

3.1 Mean and Variance of an Exponential Random Variable

Let us now compute the expectation and variance of $X \sim \text{Exp}(\lambda)$.

Theorem 20.2. Let X be an exponential random variable with parameter $\lambda > 0$. Then

$$\mathbb{E}[X] = \frac{1}{\lambda} \quad \text{and} \quad \text{Var}(X) = \frac{1}{\lambda^2}.$$

Proof. We can calculate the expected value using integration by parts:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f(x) \, dx = \int_0^{\infty} \lambda x e^{-\lambda x} \, dx = -x e^{-\lambda x} \Big|_0^{\infty} + \int_0^{\infty} e^{-\lambda x} \, dx = 0 + \left(-\frac{e^{-\lambda x}}{\lambda} \right) \Big|_0^{\infty} = \frac{1}{\lambda}.$$

To compute the variance, we first evaluate $\mathbb{E}[X^2]$, again using integration by parts:

$$\mathbb{E}[X^2] = \int_{-\infty}^{\infty} x^2 f(x) dx = \int_0^{\infty} \lambda x^2 e^{-\lambda x} dx = -x^2 e^{-\lambda x} \Big|_0^{\infty} + \int_0^{\infty} 2x e^{-\lambda x} dx = 0 + \frac{2}{\lambda} \mathbb{E}[X] = \frac{2}{\lambda^2}.$$

The variance is therefore

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

as claimed. □

3.2 As Continuous Time Analog of Geometric Distribution

Like the geometric distribution, the exponential distribution has a single parameter λ , which characterizes the *rate* at which events happen. Note that the exponential distribution satisfies, for any $t \geq 0$,

$$\mathbb{P}[X > t] = \int_t^{\infty} \lambda e^{-\lambda x} dx = -e^{-\lambda x} \Big|_t^{\infty} = e^{-\lambda t}. \quad (5)$$

In other words, the probability that we have to wait more than time t for our event to happen is $e^{-\lambda t}$, which is an exponential decay with rate λ .

Now consider a discrete-time setting in which we perform one trial every δ seconds (where δ is very small — in fact, we will take $\delta \rightarrow 0$ to make time “continuous”), and where our success probability is $p = \lambda \delta$. Making the success probability proportional to δ makes sense, as it corresponds to the natural assumption that there is a fixed *rate* of success *per unit time*, which we denote by $\lambda = p/\delta$. In this discrete setting, the number of trials until we get a success has the geometric distribution with parameter p , so if we let the r.v. Y denote the time (in seconds) until we get a success, we have

$$\mathbb{P}[Y > k\delta] = (1 - p)^k = (1 - \lambda \delta)^k, \quad \text{for any integer } k \geq 0.$$

Hence, for any $t > 0$, we have

$$\mathbb{P}[Y > t] = \mathbb{P}[Y > (\frac{t}{\delta})\delta] = (1 - \lambda \delta)^{t/\delta} \approx e^{-\lambda t}, \quad (6)$$

where this final approximation holds in the limit as $\delta \rightarrow 0$ with λ and t fixed. (We are ignoring the detail of rounding $\frac{t}{\delta}$ to an integer since we are taking an approximation anyway.)

Comparing (6) with (5), we see that this distribution has the same form as the exponential distribution with parameter λ . Thus we may view the exponential distribution as a continuous time analog of the geometric distribution, where the parameter p (success probability per trial) is replaced by the parameter λ (success probability per unit time). Note, though, that λ is not constrained to be ≤ 1 .

4 Normal Distribution

Another continuous distribution we will consider, and by far the most prevalent in applications, is called the *normal* or *Gaussian* distribution. It has two parameters, μ and σ^2 , which are the mean and variance of the distribution, respectively.

Definition 20.8 (Normal Distribution). *For any $\mu \in \mathbb{R}$ and $\sigma > 0$, a continuous random variable X with p.d.f.*

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/(2\sigma^2)}$$

is called a normal random variable with parameters μ and σ^2 , and we write $X \sim N(\mu, \sigma^2)$. In the special case $\mu = 0$ and $\sigma = 1$, X is said to have the standard normal distribution.




Figure 2: The density function for the normal distribution with several different choices for μ and σ^2 .

Let's first check that this is a valid definition of a probability density function. Clearly $f(x) \geq 0$ from the definition. For the condition in (1):

$$\int_{-\infty}^{\infty} f(x) dx = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} e^{-(x-\mu)^2/(2\sigma^2)} dx = 1. \quad (7)$$

The fact that this integral evaluates to 1 is a standard exercise in (multivariable) calculus (or feel free to look it up in any standard text on probability or on the internet).

A plot of the p.d.f. f reveals a classical "bell-shaped" curve, centered at (and symmetric around) $x = \mu$, and with "width" determined by σ . Figure 2 shows that the normal densities with different values of μ and σ are very similar to each other. Indeed, the normal distribution has the following nice property with respect to shifting and rescaling.

Lemma 20.1. *If $X \sim N(\mu, \sigma^2)$, then $Y = \frac{X-\mu}{\sigma} \sim N(0, 1)$. Equivalently, if $Y \sim N(0, 1)$, then $X = \sigma Y + \mu \sim N(\mu, \sigma^2)$.*

Proof. Given that $X \sim N(\mu, \sigma^2)$, we can calculate the distribution of $Y = \frac{X-\mu}{\sigma}$ as:

$$\mathbb{P}[a \leq Y \leq b] = \mathbb{P}[\sigma a + \mu \leq X \leq \sigma b + \mu] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{\sigma a + \mu}^{\sigma b + \mu} e^{-(x-\mu)^2/(2\sigma^2)} dx = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-y^2/2} dy,$$

by a simple change of variable $x = \sigma y + \mu$ in the integral. Hence Y is indeed standard normal. Note that Y is obtained from X just by shifting the origin to μ and scaling by σ . \square

4.1 Mean and Variance of a Normal Random Variable

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

Theorem 20.3. *For $X \sim N(\mu, \sigma^2)$,*

$$\mathbb{E}[X] = \mu \quad \text{and} \quad \text{Var}(X) = \sigma^2.$$

Proof. First consider the case when $X \sim N(0, 1)$. By definition, its expectation is

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f(x) dx = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x e^{-x^2/2} dx = \frac{1}{\sqrt{2\pi}} \left(\int_{-\infty}^0 x e^{-x^2/2} dx + \int_0^{\infty} x e^{-x^2/2} dx \right) = 0.$$

The last step follows from the fact that the function $e^{-x^2/2}$ is symmetrical about $x = 0$, so the two integrals are the same except for the sign. For the variance, we have

$$\begin{aligned}\text{Var}(X) &= \mathbb{E}[X^2] - \mathbb{E}[X]^2 = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x^2 e^{-x^2/2} dx \\ &= \frac{1}{\sqrt{2\pi}} \left(-x e^{-x^2/2} \right) \Big|_{-\infty}^{\infty} + \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-x^2/2} dx \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-x^2/2} dx = 1.\end{aligned}$$

In the first line here we used the fact that $\mathbb{E}[X] = 0$; in the second line we used integration by parts; and in the last line we used (7) in the special case $\mu = 0$, $\sigma = 1$. So the standard normal distribution has expectation $\mathbb{E}[X] = 0 = \mu$ and variance $\text{Var}(X) = 1 = \sigma^2$.

Now consider the general case when $X \sim N(\mu, \sigma^2)$. By Lemma 20.1, we know that $Y = \frac{X-\mu}{\sigma}$ is a standard normal random variable, so $\mathbb{E}[Y] = 0$ and $\text{Var}(Y) = 1$, as we have just established above. Therefore, we can read off the expectation and variance of X from those of Y . For the expectation, using linearity, we have

$$0 = \mathbb{E}[Y] = \mathbb{E}\left[\frac{X-\mu}{\sigma}\right] = \frac{\mathbb{E}[X] - \mu}{\sigma},$$

and hence $\mathbb{E}[X] = \mu$. For the variance we have

$$1 = \text{Var}(Y) = \text{Var}\left(\frac{X-\mu}{\sigma}\right) = \frac{\text{Var}(X)}{\sigma^2},$$

and hence $\text{Var}(X) = \sigma^2$. □

The bottom line, then, is that the normal distribution has expectation μ and variance σ^2 . This explains the notation for the parameters μ and σ^2 .

The fact that the variance is σ^2 (so that the standard deviation is σ) explains our earlier comment that σ determines the “width” of the normal distribution. Namely, by Chebyshev’s inequality, a constant fraction of the distribution lies within distance (say) 2σ of the expectation μ .

Note: The above analysis shows that, by means of a simple origin shift and scaling, we can relate any normal distribution to the standard normal. This means that, when doing computations with normal distributions, it’s enough to do them for the standard normal. For this reason, books and online sources of mathematical formulas usually contain tables describing the density of the standard normal. From this, one can read off the corresponding information for any normal r.v. $X \sim N(\mu, \sigma^2)$ from the formula

$$\mathbb{P}[X \leq a] = \mathbb{P}\left[Y \leq \frac{a-\mu}{\sigma}\right],$$

where Y is standard normal.

The normal distribution is ubiquitous throughout the sciences and the social sciences, because it is the standard model for aggregate data that results from averaging a large number of independent observations of the same random variable (such as the weights of mosquitos in Berkeley, or the outcome of a physical experiment). Such averaged data, as is well known, tends to cluster around its mean in a “bell-shaped” curve, with the correspondence becoming more accurate as the number of observations increases. A theoretical explanation of this phenomenon is the Central Limit Theorem, which we discuss in Section 5.

4.2 Sum of Independent Normal Random Variables

An important property of the normal distribution is that the sum of *independent* normal random variables is also normally distributed. We begin with the simple case when X and Y are independent standard normal random variables. In this case the result follows because the joint distribution of X and Y is rotationally symmetric. The general case follows from the translation and scaling property of normal distribution in Lemma 20.1.

Theorem 20.4. *Let $X \sim N(0, 1)$ and $Y \sim N(0, 1)$ be independent standard normal random variables, and suppose $a, b \in \mathbb{R}$ are constants. Then $Z = aX + bY \sim N(0, a^2 + b^2)$.*

Proof. ⁴ Since X and Y are independent, by Theorem 20.1 we know that the joint density of (X, Y) is

$$f_{X,Y}(x, y) = f_X(x) \cdot f_Y(y) = \frac{1}{2\pi} e^{-(x^2+y^2)/2}.$$

The key observation is that $f_{X,Y}(x, y)$ is rotationally symmetric around the origin, i.e., $f_{X,Y}(x, y)$ only depends on the value $x^2 + y^2$, which is the distance of the point (x, y) from the origin $(0, 0)$; see Figure 3.




Figure 3: The joint density function $f_{X,Y}(x, y) = \frac{1}{2\pi} e^{-(x^2+y^2)/2}$ is rotationally symmetric.

Thus, $f_{X,Y}(T(x, y)) = f_{X,Y}(x, y)$ where T is any rotation of the plane \mathbb{R}^2 about the origin. It follows that for any set $A \subseteq \mathbb{R}^2$,

$$\mathbb{P}[(X, Y) \in A] = \mathbb{P}[(X, Y) \in T(A)] \quad (8)$$

where T is a rotation of \mathbb{R}^2 . Now given any $t \in \mathbb{R}$, we have

$$\mathbb{P}[Z \leq t] = \mathbb{P}[aX + bY \leq t] = \mathbb{P}[(X, Y) \in A]$$

where A is the half plane $\{(x, y) \mid ax + by \leq t\}$. The boundary line $ax + by = t$ lies at a distance $d = \frac{t}{\sqrt{a^2 + b^2}}$ from the origin. Therefore, as illustrated in Figure 4, the set A can be rotated into the set

$$T(A) = \left\{ (x, y) \mid x \leq \frac{t}{\sqrt{a^2 + b^2}} \right\}.$$




Figure 4: The half plane $ax + by \leq t$ is rotated into the half plane $x \leq \frac{t}{\sqrt{a^2 + b^2}}$.

By (8), this rotation does not change the probability:

$$\mathbb{P}[Z \leq t] = \mathbb{P}[(X, Y) \in A] = \mathbb{P}[(X, Y) \in T(A)] = \mathbb{P}\left[X \leq \frac{t}{\sqrt{a^2 + b^2}}\right] = \mathbb{P}\left[\sqrt{a^2 + b^2}X \leq t\right].$$

Since the equation above holds for all $t \in \mathbb{R}$, we conclude that Z has the same distribution as $\sqrt{a^2 + b^2}X$. Since X has standard normal distribution, we know by Lemma 20.1 that $\sqrt{a^2 + b^2}X$ has normal distribution with mean 0 and variance $a^2 + b^2$. Hence we conclude that $Z = aX + bY$ also has normal distribution with mean 0 and variance $a^2 + b^2$. \square

The general case now follows easily from Lemma 20.1 and Theorem 20.4.

Corollary 20.1. *Let $X \sim N(\mu_X, \sigma_X^2)$ and $Y \sim N(\mu_Y, \sigma_Y^2)$ be independent normal random variables. Then for any constants $a, b \in \mathbb{R}$, the random variable $Z = aX + bY$ is also normally distributed with mean $\mu = a\mu_X + b\mu_Y$ and variance $\sigma^2 = a^2\sigma_X^2 + b^2\sigma_Y^2$.*

Proof. By Lemma 20.1, $Z_1 = (X - \mu_X)/\sigma_X$ and $Z_2 = (Y - \mu_Y)/\sigma_Y$ are independent standard normal random variables. We can write:

$$Z = aX + bY = a(\mu_X + \sigma_X Z_1) + b(\mu_Y + \sigma_Y Z_2) = (a\mu_X + b\mu_Y) + (a\sigma_X Z_1 + b\sigma_Y Z_2).$$

By Theorem 20.4, $Z' = a\sigma_X Z_1 + b\sigma_Y Z_2$ is normally distributed with mean 0 and variance $\sigma^2 = a^2\sigma_X^2 + b^2\sigma_Y^2$. Since $\mu = a\mu_X + b\mu_Y$ is a constant, by Lemma 20.1 we conclude that $Z = \mu + Z'$ is a normal random variable with mean μ and variance σ^2 , as desired. \square

5 The Central Limit Theorem

Recall from an earlier note the Law of Large Numbers for i.i.d. random variables $\{X_i\}$: it says that the probability of any deviation $\varepsilon > 0$, however small, of the sample average $\frac{S_n}{n}$, where $S_n = \sum_{i=1}^n X_i$, from the mean tends to zero as the number of observations n in our average tends to infinity. Thus, by taking n large enough, we can make the probability of any given deviation as small as we like.

Actually we can say something much stronger than the Law of Large Numbers: namely, the distribution of the sample average $\frac{S_n}{n}$, for large enough n , looks like a *normal distribution* with mean μ and variance $\frac{\sigma^2}{n}$. (Of course, we already know that these are the mean and variance of $\frac{S_n}{n}$; the point is that the distribution

⁴The following proof and figures are adapted from “Why Is the Sum of Independent Normal Random Variables Normal?” by B. Eisenberg and R. Sullivan, Mathematics Magazine, Vol. 81, No. 5.

becomes normal!) The fact that the standard deviation decreases with n (specifically, as $\frac{\sigma}{\sqrt{n}}$) means that the distribution approaches a sharp spike at μ .

Recall from the last section that the density of the normal distribution is a symmetrical bell-shaped curve centered around the mean μ . Its height and width are determined by the standard deviation σ as follows: the height at the mean $x = \mu$ is $\frac{1}{\sqrt{2\pi}\sigma^2} \approx \frac{0.4}{\sigma}$; 50% of the mass is contained in the interval of width 0.67σ either side of the mean, and 99.7% in the interval of width 3σ either side of the mean. (Note that, to get the correct scale, deviations are on the order of σ rather than σ^2 .)

To state the Central Limit Theorem precisely (so that the limiting distribution is a constant rather than something that depends on n), we standardize $\frac{S_n}{n}$ as

$$\frac{\frac{S_n}{n} - \mu}{\frac{\sigma}{\sqrt{n}}} = \frac{S_n - n\mu}{\sigma\sqrt{n}}.$$

The Central Limit Theorem then says that the distribution of $\frac{S_n - n\mu}{\sigma\sqrt{n}}$ converges to the *standard normal* distribution.

Theorem 20.5 (Central Limit Theorem). *Let X_1, X_2, \dots be a sequence of i.i.d. random variables with common finite expectation $\mathbb{E}[X_i] = \mu$ and finite variance $\text{Var}(X_i) = \sigma^2$. Let $S_n = \sum_{i=1}^n X_i$. Then, the distribution of $\frac{S_n - n\mu}{\sigma\sqrt{n}}$ converges to $N(0, 1)$ as $n \rightarrow \infty$. In other words, for any constant $c \in \mathbb{R}$,*

$$\mathbb{P}\left[\frac{S_n - n\mu}{\sigma\sqrt{n}} \leq c\right] \rightarrow \frac{1}{\sqrt{2\pi}} \int_{-\infty}^c e^{-x^2/2} dx \quad \text{as } n \rightarrow \infty.$$

The Central Limit Theorem is a very striking fact. What it says is the following: If we take an average of n observations of any arbitrary r.v. X , then the distribution of that average will be a bell-shaped curve centered at $\mu = \mathbb{E}[X]$. Thus all trace of the distribution of X disappears as n gets large: all distributions, no matter how complex,⁵ look like the normal distribution when they are averaged. The only effect of the original distribution is through the variance σ^2 , which determines the width of the curve for a given value of n , and hence the rate at which the curve shrinks to a spike.

The Central Limit Theorem immediately tells us that *averaged* data tends to look normal, even when the distribution it is drawn from is not itself normal. Since much of the data we work with in real life is the result of averaging (mean wealth, mean height, mean temperature etc.), it's not surprising that we deal with the normal distribution so often. But there's actually a deeper reason why even (say) the height of a population itself tends to follow a normal distribution: this is because a person's height is actually itself the result of averaging many factors (nutrition, environment, various genetic factors, etc.), so intuitively at least we would expect it to follow a normal distribution.

6 Buffon's Needle

Here is a simple yet interesting application of continuous random variables to the analysis of a classical procedure for estimating the value of π ; this is known as *Buffon's needle* problem, after its 18th century inventor Georges-Louis Leclerc, Comte de Buffon.

As illustrated in Figure 5, we are given a needle of length ℓ , and a board ruled with horizontal lines at distance ℓ apart. The experiment consists of throwing the needle randomly onto the board and observing

⁵We do need to assume that the mean and variance of X are finite.




Figure 5: Buffon's Needle.

whether or not it crosses one of the lines. We shall see below that (assuming a perfectly random throw) the probability of this event is exactly $2/\pi$. This means that, if we perform the experiment many times and record the *proportion* of throws on which the needle crosses a line, then the Law of Large Numbers tells us that we will get a good estimate of the quantity $2/\pi$, and therefore also of π ; and we can use Chebyshev's inequality as in the other estimation problems we considered in an earlier note to determine how many throws we need in order to achieve specified accuracy and confidence.

To analyze the experiment, let's consider what random variables are in play. Note that the position where the needle lands is completely specified by two random variables: the vertical distance Y between the midpoint of the needle and the closest horizontal line, and the angle Θ between the needle and the vertical. The r.v. Y ranges between 0 and $\ell/2$, while Θ ranges between $-\pi/2$ and $\pi/2$. Since we assume a perfectly random throw, we may assume that their *joint distribution* has density $f_{Y,\Theta}(y, \theta)$ that is uniform over the rectangle $[0, \ell/2] \times [-\pi/2, \pi/2]$. Since this rectangle has area $\frac{\pi\ell}{2}$, the density should be

$$f_{Y,\Theta}(y, \theta) = \begin{cases} \frac{2}{\pi\ell}, & \text{for } (y, \theta) \in [0, \ell/2] \times [-\pi/2, \pi/2], \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

Equivalently, Y and Θ are independent random variables, each uniformly distributed in their respective range. To check our answer, let's verify that the integral of this density over all possible values is indeed 1:

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{Y,\Theta}(y, \theta) \, dy \, d\theta = \int_{-\pi/2}^{\pi/2} \int_0^{\ell/2} \frac{2}{\pi\ell} \, dy \, d\theta = \int_{-\pi/2}^{\pi/2} \frac{2y}{\pi\ell} \Big|_0^{\ell/2} \, d\theta = \int_{-\pi/2}^{\pi/2} \frac{1}{\pi} \, d\theta = \frac{\theta}{\pi} \Big|_{-\pi/2}^{\pi/2} = 1.$$

This is an analog of (1) for our joint distribution; rather than the area under the curve $f(x)$, we are now computing the area under the "surface" $f_{Y,\Theta}(y, \theta)$.

Now let E denote the event that the needle crosses a line. How can we express this event in terms of the values of Y and Θ ? Well, by elementary geometry the vertical distance of the endpoint of the needle from its midpoint is $\frac{\ell}{2} \cos \Theta$, so the needle will cross the line if and only if $Y \leq \frac{\ell}{2} \cos \Theta$. Therefore we have

$$\mathbb{P}[E] = \mathbb{P}[Y \leq \frac{\ell}{2} \cos \Theta] = \int_{-\pi/2}^{\pi/2} \int_0^{(\ell/2) \cos \theta} f_{Y,\Theta}(y, \theta) \, dy \, d\theta.$$

Substituting the density $f_{Y,\Theta}(y, \theta)$ from (9) and performing the integration we get

$$\mathbb{P}[E] = \int_{-\pi/2}^{\pi/2} \int_0^{(\ell/2) \cos \theta} \frac{2}{\pi\ell} \, dy \, d\theta = \int_{-\pi/2}^{\pi/2} \frac{2y}{\pi\ell} \Big|_0^{(\ell/2) \cos \theta} \, d\theta = \frac{1}{\pi} \int_{-\pi/2}^{\pi/2} \cos \theta \, d\theta = \frac{1}{\pi} \sin \theta \Big|_{-\pi/2}^{\pi/2} = \frac{2}{\pi}.$$

This is exactly what we claimed at the beginning of the section!