

# How to Work with Normal Distributions in R Using `dnorm`, `pnorm`, `qnorm`, and `rnorm`

Authored by  
**stats writer**

December 31, 2025

## RECOMMENDED CITATION

stats writer (2025). *How to Work with Normal Distributions in R Using `dnorm`, `pnorm`, `qnorm`, and `rnorm`*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=110126>

The **R programming language** offers a robust set of functions for working with statistical distributions. Among the most fundamental are the quartet of functions designed specifically for the **normal distribution** (often called the Gaussian distribution): **`dnorm`**, **`pnorm`**, **`qnorm`**, and **`rnorm`**. These functions serve specialized roles in statistical analysis, corresponding to four distinct facets of probability modeling: calculating density, determining cumulative probabilities, finding quantiles, and generating random data, respectively.

Understanding how to effectively utilize these four functions is crucial for anyone performing statistical inference or data modeling in R. The normal distribution is arguably the most important distribution in **statistics** due to the Central Limit Theorem and its prevalence in natural phenomena. This comprehensive guide details the purpose, syntax, and practical applications of **`dnorm`**, **`pnorm`**, **`qnorm`**, and **`rnorm`**, providing clear examples for each function.

## **dnorm: The Probability Density Function**

The **`dnorm`** function is used to calculate the value of the **probability density function (PDF)** for the normal distribution at a specific point, often denoted as  $x$ . For continuous distributions like the normal distribution, the PDF does not directly give the probability of a specific value occurring (which is zero), but rather describes the relative likelihood of a **random variable** taking on that value. This resulting value represents the height of the curve at point  $x$  and is essential for visualizing the shape of the distribution.

When utilizing **`dnorm`**, you must specify the value of the random variable  $x$ , along with the two required parameters that define the specific normal distribution: the population mean ( $\mu$ ) and the population standard deviation ( $\sigma$ ). If both parameters are omitted, R defaults to the standard normal distribution, where the mean is 0 and the standard deviation is 1. The general syntax for invoking this function is structured as follows:

**`dnorm(x, mean, sd)`**

The following R code provides practical illustrations of how **`dnorm`** calculates the density. These examples show how to calculate the density for the standard normal curve and for a customized distribution:

**# Find the density function value (PDF height) for the standard normal distribution at x=0**

```
dnorm(x=0, mean=0, sd=1)
```

```
# 0.3989423
```

```
# R defaults to mean=0 and sd=1 if parameters are not explicitly provided
```

```
dnorm(x=0)
```

```
# 0.3989423
```

```
# Calculate the density for a normal distribution at x=10, given a mean of 20 and a standard
deviation of 5
dnorm(x=10, mean=20, sd=5)
# 0.01079819
```

While **`dnorm`** is rarely used for direct probability calculation--a task better suited for **`pnorm`**--its essential application is in generating the coordinates needed for graphical representation. By calculating the PDF value across a sequence of `x` values, we can accurately plot the characteristic bell curve of the **normal distribution**. The subsequent example illustrates the necessary steps to generate and plot the standard normal curve using R's built-in plotting functions:

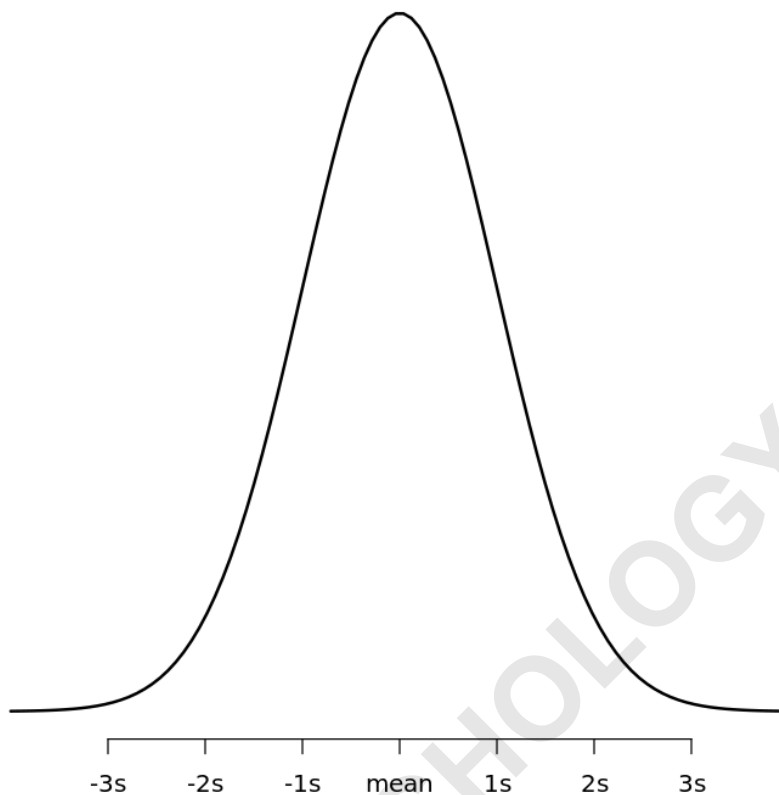
**# Create a sequence of 100 equally spaced numbers between -4 and 4 to define the x-axis range**

```
x <- seq(-4, 4, length=100)
```

```
# Calculate the corresponding height (density) of the probability distribution for each value in x
y <- dnorm(x)
```

```
# Plot x and y using a line plot (type = "l"), setting line width (lwd) and customizing the axes
plot(x,y, type = "l", lwd = 2, axes = FALSE, xlab = "", ylab = "")
axis(1, at = -3:3, labels = c("-3s", "-2s", "-1s", "mean", "1s", "2s", "3s"))
```

Executing this code produces a visualization of the standard normal curve, mapping the density values across the specified range:



### **pnorm: Calculating Cumulative Probability**

The **pnorm** function is indispensable for solving practical probability problems involving continuous distributions. It calculates the value of the **cumulative density function (CDF)**. The CDF determines the total probability that a **random variable** will assume a value less than or equal to a specific threshold  $q$ . Geometrically, this corresponds to finding the area under the **probability density function (PDF)** curve from negative infinity up to the point  $q$ .

To use **pnorm**, you supply the quantile  $q$  (the value used as the upper limit for accumulation), the mean ( $\mu$ ), and the standard deviation ( $\sigma$ ) of the population. The basic syntax mirrors that of **dnorm**:

**pnorm(q, mean, sd)**

A highly useful feature of **pnorm** is the optional argument **lower.tail**. By default, **pnorm** calculates the area to the left of  $q$  ( $P(X \leq q)$ ), equivalent to **lower.tail = TRUE**. If the desired probability involves the area to the right of  $q$  ( $P(X > q)$ ), we can specify **lower.tail = FALSE**. This ability to

directly calculate the upper-tail probability simplifies the workflow, preventing the need for the manual calculation of 1 minus the left-tail probability.

```
pnorm(q, mean, sd, lower.tail = FALSE)
```

## **pnorm Applications: Solving Probability Problems**

The versatility of **pnorm** allows us to address various probability scenarios commonly encountered in statistical modeling, ranging from simple tail probabilities to calculating probabilities within a defined interval.

### **Example 1: Upper-Tail Probability (Taller than 74 inches)**

Imagine that the height of males in a school setting is **normally distributed** with a mean of 70 inches and a standard deviation of 2 inches. To find the percentage of males taller than 74 inches, we must calculate the area in the upper tail, which we achieve by setting **lower.tail = FALSE**:

```
# Find the percentage of males taller than 74 inches (upper tail) in a population with mean = 70 and sd = 2
```

```
pnorm(74, mean=70, sd=2, lower.tail=FALSE)
```

```
# 0.02275013
```

The result indicates that approximately 0.02275, or **2.275%**, of males at this school are taller than 74 inches.

### **Example 2: Lower-Tail Probability (Weight less than 22 lbs)**

Suppose the weight of a certain species of otters is normally distributed with a mean of 30 lbs and a standard deviation of 5 lbs. If we seek the percentage of otters weighing less than 22 lbs, we are calculating the area in the lower tail. We can use the default settings of **pnorm** for this calculation:

```
# Find the percentage of otters weighing less than 22 lbs (lower tail) in a population with mean = 30 and sd = 5
```

```
pnorm(22, mean=30, sd=5)
```

```
# 0.05479929
```

The output shows that 0.054799, or roughly **5.48%**, of this species of otters weigh less than 22 lbs.

### **Example 3: Probability Within a Range (Between 10 and 14 inches)**

If the height of plants in a region is normally distributed with a mean of 13 inches and a standard deviation of 2 inches, finding the proportion of plants between 10 and 14 inches requires an area subtraction approach. We calculate the cumulative probability up to the upper bound (14 inches) and subtract the cumulative probability up to the lower bound (10 inches):  $P(10 \leq X \leq 14) = P(X \leq 14) - P(X \leq 10)$ .

```
# Calculate P(X <= 14) and subtract P(X <= 10) to find the probability within the specified range
```

```
# Based on a population with mean = 13 and sd = 2
```

```
pnorm(14, mean=13, sd=2) - pnorm(10, mean=13, sd=2)
```

```
# 0.6246553
```

This calculation yields 0.624655, meaning approximately **62.47%** of plants in this region are expected to have a height between 10 and 14 inches.

## qnorm: The Inverse Cumulative Function (Quantile Function)

The `qnorm` function performs the inverse operation of `pnorm`. Instead of taking a value and returning a probability, `qnorm` takes a probability ( $p$ , representing the area under the curve) and returns the corresponding value (the quantile) on the x-axis. This function is officially known as the **Quantile Function** or the inverse **Cumulative Distribution Function (CDF)**.

The primary use of `qnorm` is to determine the critical boundary value associated with a specific cumulative probability, which is particularly vital for constructing confidence intervals or setting thresholds for statistical testing. When applied to the standard **normal distribution** (mean=0, sd=1), the result returned by `qnorm` is the **Z-score** corresponding to the  $p$ -th percentile.

The syntax for `qnorm` requires specifying the target probability  $p$  (which must be between 0 and 1), along with the mean and standard deviation of the population distribution:

```
qnorm(p, mean, sd)
```

The following examples illustrate how `qnorm` is used to calculate Z-scores associated with common percentile thresholds. Since these examples rely on the standard normal distribution, the mean and standard deviation parameters are often omitted, relying on R's defaults:

```
# Find the Z-score (value) corresponding to the 99th percentile (p=0.99) of the standard normal distribution
```

```
qnorm(.99, mean=0, sd=1)
```

```
# 2.326348
```

```
# R uses mean=0 and sd=1 by default when arguments are omitted
```

```
qnorm(.99)
```

```
# 2.326348
```

```
# Determine the Z-score for the 95th percentile
```

```
qnorm(.95)
```

```
# 1.644854
```

```
# Find the Z-score associated with the 10th percentile (p=0.10)
```

```
qnorm(.10)
```

```
# -1.281552
```

## **rnorm: Random Generation of Normal Data**

The **rnorm** function is a powerful tool in computational statistics, designed for stochastic simulation by generating a vector of random numbers that adhere to a specified normal distribution. This capability is fundamental in Monte Carlo methods, simulation studies, bootstrapping, and for testing statistical models with synthetic data.

To employ **rnorm**, the user must specify three essential arguments: the desired number of observations (the vector length  $n$ ), the population mean ( $\mu$ ), and the population standard deviation ( $\sigma$ ). The resulting output is a collection of **random variables** whose empirical distribution approximates the specified normal distribution parameters. The function's syntax is:

```
rnorm(n, mean, sd)
```

The following R code demonstrates generating both small and large samples. The final example uses two large samples with differing standard deviations to visualize how the variance parameter impacts the spread of the generated data when plotted using histograms:

```
# Generate a vector of 5 normally distributed random variables with mean=10 and sd=2
```

```
five <- rnorm(5, mean = 10, sd = 2)
```

```
five
```

```
# 10.658117 8.613495 10.561760 11.123492 10.802768
```

```
# Generate a vector of 1000 observations centered at 50 with a narrow standard deviation (sd=15)
```

```
narrowDistribution <- rnorm(1000, mean = 50, sd = 15)
```

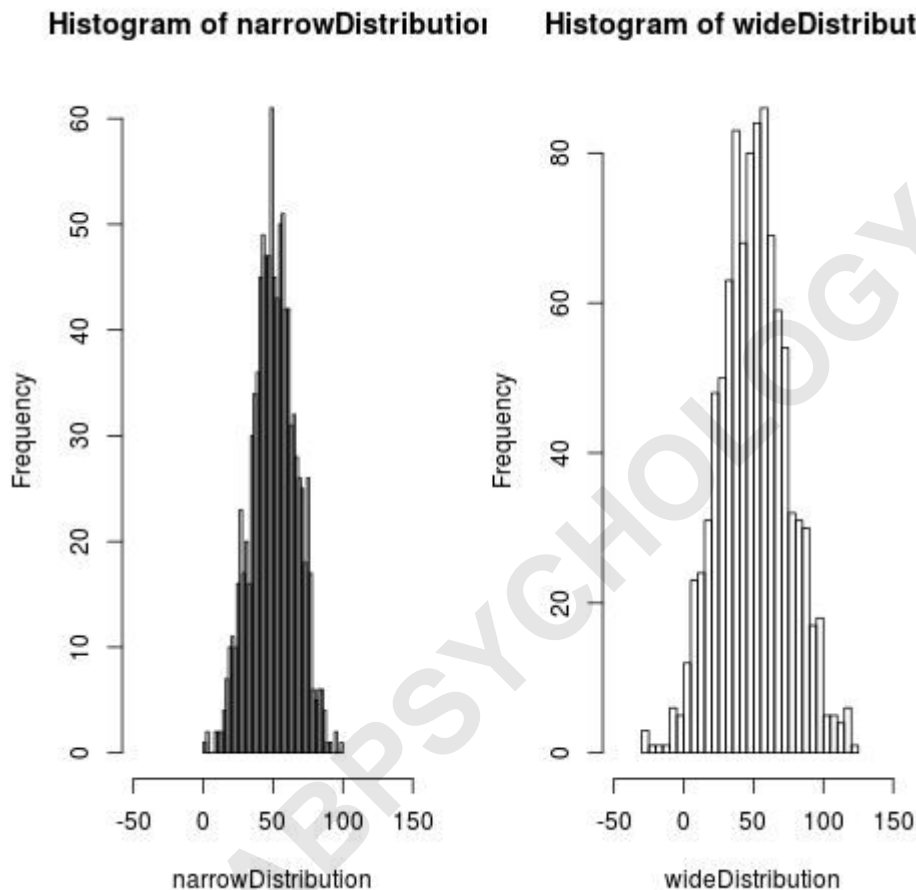
```
# Generate a vector of 1000 observations centered at 50 with a wide standard deviation (sd=25)
```

```
wideDistribution <- rnorm(1000, mean = 50, sd = 25)
```

```
# Visualize both distributions side-by-side using histograms for comparison.
```

```
par(mfrow=c(1, 2)) # Set graphical parameters to display plots in one row and two columns
hist(narrowDistribution, breaks=50, xlim=c(-50, 150), main="Narrow Distribution (SD=15)")
hist(wideDistribution, breaks=50, xlim=c(-50, 150), main="Wide Distribution (SD=25)")
```

This code generates the following histograms, vividly demonstrating the impact of the standard deviation:



As observed in the visualization, the "Wide Distribution" (SD=25) exhibits significantly greater variability and spread across the x-axis compared to the "Narrow Distribution" (SD=15). Despite this difference in dispersion, both histograms remain correctly centered around the specified mean of 50. This confirms that **rnorm** accurately generates synthetic data sets according to the input parameters, making it an essential tool for simulation and parameter sensitivity analysis.