# How to Easily Generate Random Numbers with rnorm() and runif() in R

Authored by

**stats writer**

November 20, 2025

The distinction between the R functions `rnorm()` and `runif()` is fundamental to understanding probability distributions in computational statistics. Both functions are essential tools for generating random numbers, yet they draw their samples from distinctly different theoretical models. Specifically, `rnorm()` sources values from a Normal distribution (also known as the Gaussian distribution), characterized by its ubiquitous bell-shaped curve. Conversely, `runif()` generates values based on the Uniform distribution, where every potential outcome within a defined interval is equally likely, resulting in a flat, rectangular probability density function when visualized.

The choice between these two functions hinges entirely on the underlying probabilistic structure required for a simulation or statistical analysis. The Normal distribution allows for the precise control of central tendency and dispersion through parameters like the mean and standard deviation, which are crucial when modeling naturally occurring phenomena or errors. In contrast, the Uniform distribution is employed when modeling processes where there is no inherent preference for one value over another within a specified range. Understanding these mathematical differences is critical for ensuring the validity and accuracy of any simulation performed in R.

## Introduction to Random Number Generation in R

Generating random variates is a cornerstone of modern statistics, enabling everything from Monte Carlo simulations to bootstrapping techniques and hypothesis testing. Within the R environment, a family of functions exists to handle various distributions, collectively known as the "r-functions" (e.g., `rnorm`, `runif`, `rpois`, `rexp`). These functions utilize complex algorithms, typically pseudo-random number generators, to produce sequences of numbers that exhibit statistical randomness based on the chosen probability density function (PDF). While these numbers are not truly random--they are generated by deterministic algorithms initialized by a seed--they serve as excellent approximations for stochastic processes.

The utility of these tools lies in their ability to model real-world uncertainty and variability computationally. By generating synthetic data that adheres to known distributional properties, analysts can test statistical models, evaluate the performance of estimators, and explore complex scenarios without relying solely on limited observational data. Both the **rnorm()** and **runif()** functions are indispensable members of this family, providing the capability to generate random values for fundamentally different statistical purposes. When initiating any simulation, the first crucial step is determining which theoretical distribution best describes the phenomenon being modeled, as this decision dictates whether we use **rnorm()** or **runif()**.

## The Core Distinction: Distribution Type

The principal difference between **rnorm()** and **runif()** lies in the underlying mathematical distribution from which they sample random numbers. The `rnorm()` function, designed for the

Normal distribution, assumes that values cluster symmetrically around a central point (the mean), with probabilities rapidly decreasing as values move further into the tails. This distribution is characterized by two essential parameters: the mean (controlling the location of the center) and the standard deviation (controlling the spread or scale). The majority of generated values will fall close to the mean, aligning with the concept of common outcomes being more likely than extreme ones.

In stark contrast, `runif()` samples from the Uniform distribution, meaning the probability density function is constant across the entire defined interval. If you specify a range (from minimum to maximum), any number within that interval has the exact same likelihood of being selected. There is no clustering around a mean or median; the distribution is perfectly flat. This makes `runif()` ideal for scenarios where true randomness within a bounded range is needed, such as selecting random coordinates or modeling physical processes where outcomes are equally probable across a fixed scale. The visual representation of the Normal distribution is the classic bell curve, while the Uniform distribution is a perfect rectangle.

## Deep Dive into rnorm(): The Normal Distribution

The Normal distribution is arguably the most important distribution in statistics due to the Central Limit Theorem, which dictates that the distribution of sample means tends towards normality, irrespective of the shape of the population distribution, provided the sample size is sufficiently large. Consequently, **rnorm()** is indispensable for modeling data that naturally follows this bell-shaped pattern, such as human height, measurement errors, IQ scores, and many financial market fluctuations. The function's parameters allow for detailed control over the generated data set, ensuring it accurately reflects the specific population characteristics being studied.

When using **rnorm(n, mean, sd)**, the output vector of length `n` will exhibit statistical properties closely mirroring the specified mean and standard deviation. A smaller standard deviation produces a narrower, taller bell curve, indicating less variability, while a larger standard deviation results in a wider, flatter curve, signifying greater spread. Furthermore, because the Normal distribution is defined across the entire real number line (from negative infinity to positive infinity), **rnorm()** can generate values far outside the specified mean, albeit with extremely low probability, reflecting the theoretical tails of the distribution. This flexibility makes it the go-to function for hypothesis testing and creating synthetic data sets for regression analysis and other advanced statistical analysis techniques.

## Parameters and Characteristics of rnorm()

The standard syntax for the **rnorm()** function is: `rnorm(n, mean = 0, sd = 1)`. The first argument, `n`, specifies the number of observations to generate. The `mean` parameter sets the expected value, or central location, of the distribution. By default, if the `mean` is not specified, it

defaults to 0, creating a standard normal distribution. The `sd` parameter specifies the standard deviation, controlling the dispersion of the data around the mean. If `sd` is not specified, it defaults to 1, again resulting in the standard normal distribution, which is central to z-score calculations and standardized testing.

It is paramount to understand that these parameters define the target population's characteristics. If an analyst knows that a certain population of scores has an average of 100 and a standard deviation of 15, using `rnorm(1000, mean=100, sd=15)` will generate 1,000 random numbers that collectively possess the same expected mean and standard deviation. If no parameters are provided beyond `n`, the function generates values from the standard Normal distribution N(0, 1). Mastery of these parameters is essential for accurately simulating processes where the variance and central location are known or hypothesized.

## Deep Dive into runif(): The Uniform Distribution

The Uniform distribution is a simple yet powerful model, often used as the default expectation of randomness when no prior information about skew or central tendency is available. Unlike the Normal distribution, the Uniform distribution has no preference for any value within its boundaries; the probability density function is constant. This means that if you are sampling from 5 to 10, the probability of selecting 5.1 is identical to the probability of selecting 8.9. This concept of equal probability across the interval makes **runif()** ideal for applications such as basic simulation initialization, selecting random samples without bias, or generating initial seeds for complex algorithms.

The generated values from **runif()** are constrained by the specified minimum and maximum values. This strict bounding is a key feature differentiating it from **rnorm()**, which theoretically spans infinity. Because the distribution is flat, the mean and median will always be located exactly halfway between the minimum and maximum values: (min + max) / 2. The simplicity of its parameterization, relying only on the limits of the interval, makes it highly predictable, though less flexible than **rnorm()** for modeling complex real-world variables that usually exhibit some degree of central tendency.

## Parameters and Characteristics of runif()

The syntax for the **runif()** function is: `runif(n, min = 0, max = 1)`. Here, `n` again denotes the number of observations. The `min` parameter defines the lower boundary of the interval, and `max` defines the upper boundary. If these boundaries are not explicitly specified, **runif()** defaults to generating random numbers between 0 and 1. This default setting, `runif(n)`, is frequently used to generate probability values or to scale other variables, as it directly produces values from the standard Uniform distribution U(0, 1).

It is crucial to note that the parameters `min` and `max` completely define the shape and range of the output distribution. For instance, if you require a random value between 100 and 200, you use `runif(n, min=100, max=200)`. Unlike **rnorm()**, you cannot directly input a mean or standard deviation; these statistical properties are derived solely from the `min` and `max` values. This clear distinction in parametrization reinforces the difference in their theoretical foundations: one is centered on location and scale (Normal), the other on boundaries (Uniform).

## Practical Implementation: Using rnorm()

To solidify the understanding of **rnorm()**, let us examine a practical example in R. We will generate a sequence of random values from a Normal distribution where the population mean is set to 10 and the standard deviation is 2. This simulates a scenario, perhaps measuring the output of a machine, where outputs center around 10 but vary slightly. The use of `set.seed()` ensures that the exact sequence of "random" numbers can be reproduced by anyone running the code, which is an indispensable practice for reproducible research and statistical analysis.

The following code shows how to use the **rnorm()** function to generate 100 random values from a Normal distribution with a mean of 10 and a standard deviation of 2:

**#make this example reproducible**
**set.seed(0)**

#create vector of 100 random values from normal distribution
random_values <- rnorm(n=100, mean=10, sd=2)

#view first six values
head(random_values)

12.525909 9.347533 12.659599 12.544859 10.829283 6.920100

After generating the data, the most effective way to confirm that the distribution is indeed Normal and centered correctly is through visualization. We can use the **hist()** function, R's built-in tool for creating histograms, to display the frequency distribution of the generated random values. This step provides immediate visual evidence supporting the theoretical claims of the Normal distribution, showing how values concentrate around the mean parameter we supplied.

We can also use the **hist()** function to create a histogram to visualize the distribution of random values we just generated:

**#create histogram to visualize distribution of values**
**hist(random_values)**

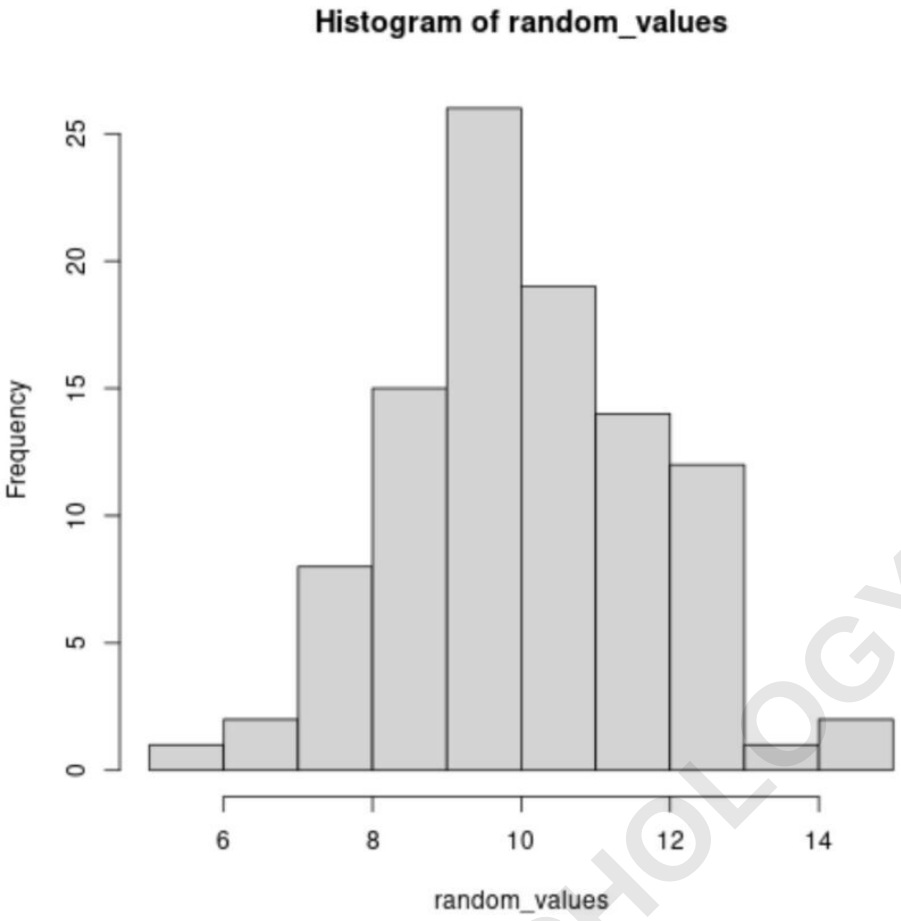




The resulting histogram clearly displays the frequency distribution of the 100 values drawn from the Normal distribution. Notice how the shape strongly approximates the characteristic bell curve, with the highest bars clustered near the center. This visual confirmation is vital in simulation work. We observe that the distribution's peak, representing the highest frequency of occurrences, is located around 10, which precisely matches the mean value we specified (`mean=10`) in the **rnorm()** function call. The distribution smoothly tapers off symmetrically on both sides, illustrating the effect of the standard deviation of 2.

## Practical Implementation: Using runif()

Next, we explore **runif()**, simulating a scenario where every outcome within a 20-unit range is equally likely, such as generating points in a fixed two-dimensional space. We will define the interval from 5 (minimum) to 25 (maximum). In this simulation, there should be no discernible peak; the histogram should show approximately equal frequency counts across all bins within the defined interval, reflecting the flatness of the Uniform distribution.

The following code shows how to use the **runif()** function to generate 100 random values from a Uniform distribution with a minimum value of 5 and a maximum value of 25:

**#make this example reproducible**
**set.seed(0)**

#create vector of 100 random values from uniform distribution
random_values <- runif(n=100, min=5, max=25)

#view first six values
head(random_values)

22.933944 10.310173 12.442478 16.457067 23.164156 9.033639

Again, we visualize the results using the **hist()** function. The expected output is a histogram where the height of the bars is roughly consistent across the entire width of the plot, illustrating the fundamental property of the Uniform distribution: constant probability density. This provides a stark contrast to the bell-shaped curve observed when using **rnorm()**, emphasizing the visual difference between location-based and boundary-based distributions.

We can also use the **hist()** function to create a histogram to visualize the distribution of random values we just generated:

**#create histogram to visualize distribution of values**
**hist(random_values)**

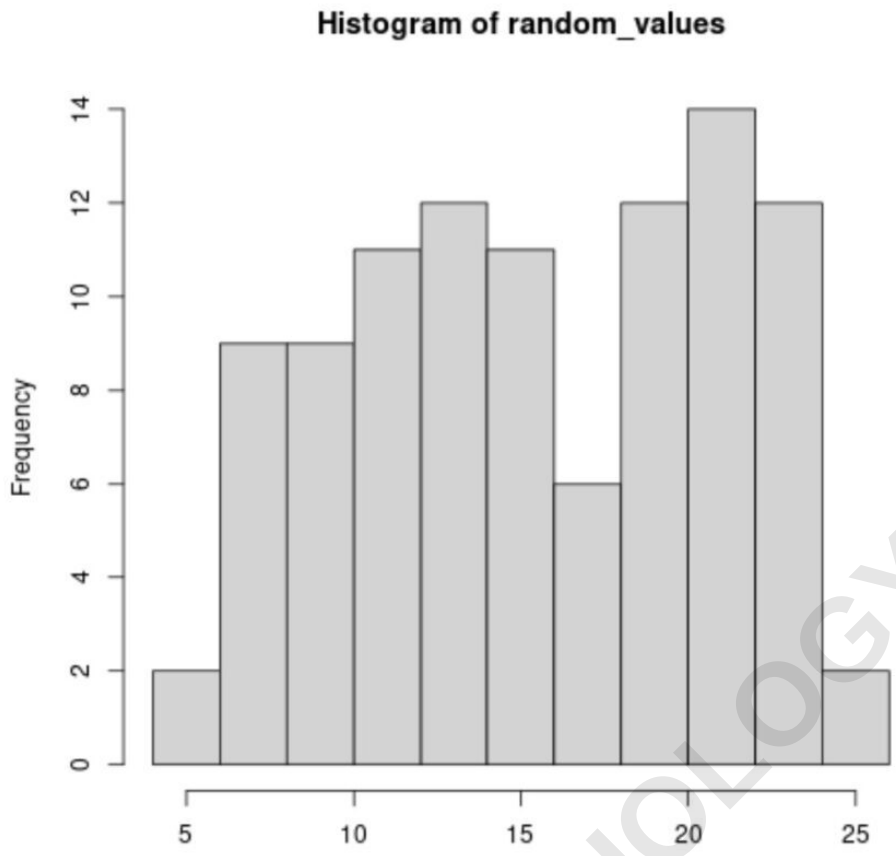




The resulting histogram confirms the theoretical shape of the Uniform distribution. Unlike the bell shape seen previously, the bars here are relatively flat and consistent in height, indicating that the 100 generated values were spread out evenly across the designated range. Critically, the distribution starts precisely at 5 and ends near 25, which perfectly corresponds to the minimum and maximum values specified in the **runif()** function call. The values are strictly bounded by these parameters, a key difference from **rnorm()** where extreme values are theoretically possible.

## Choosing the Right Function for Your Statistical Analysis

The decision to use **rnorm()** or **runif()** is a fundamental step in designing valid computational experiments in R. If the data being modeled represents a continuous measurement subject to natural variation, errors, or averages--such as test scores, physical properties, or residuals in a model--the Normal distribution and **rnorm()** are the appropriate choices, allowing the simulation to incorporate a known mean and standard deviation. This aligns with most parametric statistical analysis requirements, particularly those relying on t-tests or ANOVA, which assume normally distributed data.

Conversely, if the simulation requires generating a random starting point, selecting a parameter from an unknown range without bias, or modeling a process where all outcomes are equally

probable within a boundary, the Uniform distribution and **runif()** are necessary. A typical application is in numerical integration methods or cryptographic seeding, where unpredictability across the entire range is crucial. Furthermore, **runif()** serves as the backbone for many complex sampling techniques, such as the inverse transform method, where uniform random variates are transformed into random variates of other distributions. Ultimately, both functions are essential, but their application is dictated by whether the data's probability density is centralized (Normal) or constant (Uniform).