

Learning the Student's t-Distribution in R: A Practical Guide to `dt()`, `qt()`, `pt()`, and `rt()` Functions

Authored by
Mohammed loot

November 9, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning the Student's t-Distribution in R: A Practical Guide to `dt()`, `qt()`, `pt()`, and `rt()` Functions*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=14217>

The [Student t distribution](#) is foundational in statistical inference, particularly when sample sizes are small or population standard deviations are unknown. Mastering its associated functions in [R](#) is essential for any data analyst or statistician. This comprehensive guide details the practical application of the four core functions--**dt()**, **qt()**, **pt()**, and **rt()**--which allow users to work seamlessly with the **Student t distribution** within the R environment.

These four functions follow a standard naming convention in [R](#) for statistical distributions: 'd' for density, 'p' for probability (cumulative), 'q' for quantile, and 'r' for random generation. Understanding this convention simplifies the navigation of R's vast statistical toolkit, enabling efficient calculation of probabilities, critical values, and the simulation of statistical data required for hypothesis testing and confidence interval construction. Each function requires the specification of **degrees of freedom** (*df*), the parameter that defines the exact shape of the distribution.

Understanding dt(): The Probability Density Function

The function **dt()** is designed to return the value of the [probability density function](#) (PDF) for the **Student t distribution** at a specified point *x*. The PDF value itself does not represent a probability but rather the relative likelihood of a continuous random variable taking on that specific value. This function is critical for visualizing the shape of the distribution and understanding how the spread changes based on the **degrees of freedom**.

The fundamental syntax for utilizing **dt()** requires two primary arguments: the random variable *x* and the **degrees of freedom** (*df*). The **degrees of freedom**, a parameter unique to the t-distribution, dictates its shape, controlling how "heavy" the tails are compared to a standard normal distribution. As *df* increases, the [Student t distribution](#) gradually approaches the standard normal distribution. The basic structure is straightforward:

dt(x, df)

Let's explore practical examples demonstrating how **dt()** operates, showing both explicit and positional argument usage. Calculating the height of the curve at the mean (*x*=0) provides a quick way to see how the distribution peaks based on the degrees of freedom. Note how R automatically assumes the order of arguments if names are omitted:

Find the value of the Student t distribution pdf at x = 0 with 20 degrees of freedom

dt(x = 0, df = 20)

```
# 0.3939886
```

```
# By default, R assumes the first argument is x and the second argument is df
```

```
dt(0, 20)
```

```
# 0.3939886
# Find the value of the Student t distribution pdf at x = 1 with 30 degrees of freedom
dt(1, 30)

# 0.2379933
```


While **dt()** provides the height of the curve, for solving traditional probability questions (e.g., finding the area under the curve), the **pt()** function is generally preferred. However, **dt()** is indispensable for graphical representation. It generates the necessary density points for plotting the bell-shaped curve of the [Student t distribution](#), enabling visual comparison across different degrees of freedom. The following code snippet demonstrates how to generate a smooth plot for the distribution using 20 degrees of freedom:

```
# Create a sequence of 100 equally spaced numbers between -4 and 4
x <- seq(-4, 4, length=100)

# Create a vector of values that shows the height of the probability distribution
# for each value in x, using 20 degrees of freedom
y <- dt(x = x, df = 20)

# Plot x and y as a scatterplot with connected lines (type = "l") and add
# an x-axis with custom labels
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 generates a clear visual representation of the t-distribution curve, centered around zero, illustrating the probability density across the range of values:



Using `pt()`: The Cumulative Distribution Function

The function `pt()` calculates the [cumulative density function](#) (CDF) of the Student t distribution. The CDF provides a direct measure of probability: specifically, the probability that a random observation drawn from the distribution will be less than or equal to a given value x . This is equivalent to finding the area under the PDF curve from negative infinity up to x , making `pt()` the primary tool for calculating P-values and determining confidence intervals.

The standard syntax for `pt()` requires the value x and the **degrees of freedom** df . By default, `pt(x, df)` returns the area to the left of x ($P(T \leq x)$). This left-tail calculation is essential for one-tailed hypothesis tests where the alternative hypothesis specifies a 'less than' relationship. The inherent structure is:

`pt(x, df)`

If the objective is to find the area to the right of x ($P(T \geq x)$), R offers a convenient argument: **`lower.tail = FALSE`**. This simplifies the calculation, avoiding the need to subtract the left-tail probability from 1. Utilizing this argument is crucial when performing right-tailed tests or calculating upper confidence limits, enhancing code clarity and computational efficiency. The syntax for calculating the upper tail is as follows:

`pt(x, df, lower.tail = FALSE)`

Consider a practical scenario involving probability calculation: finding the area to the left of a specific t-statistic. Suppose we have a t-statistic with a value of -0.785 and 14 **degrees of freedom**. Using **pt()** directly provides the desired probability for a left-tailed test:

Example 1: Find the area to the **left** of a t-statistic with value of -0.785 and 14 degrees of freedom.

```
pt(-0.785, 14)
```

```
# 0.2227675
```

If we needed the area to the **right** of that t-statistic, we can use the **lower.tail = FALSE** argument, which produces an equivalent result to subtracting the left-tail probability from one. This flexibility ensures that users can easily tackle different forms of probability questions, which is vital for performing right-tailed hypothesis testing:

The following approaches produce equivalent results for the area to the right:

```
# 1 - Area to the left (traditional method)
```

```
1 - pt(-0.785, 14)
```

```
# 0.7772325
```

```
# Area to the right (using the specific argument in R)
```

```
pt(-0.785, 14, lower.tail = FALSE)
```

```
# 0.7772325
```

Example 3: Finding two-tailed probability is crucial for standard two-sided hypothesis testing. We often need to find the total probability that lies in the extreme tails (both left and right) beyond a certain positive and negative t-score. For a symmetrical distribution like the Student t distribution, this is calculated by summing the area to the left of the negative score and the area to the right of the positive score. This cumulative probability represents the P-value in a two-sided test.

Example 3: Find the total area in a **Student t distribution** with 14 degrees of freedom that lies to the left of -0.785 or to the right of 0.785.

```
pt(-0.785, 14) + pt(0.785, 14, lower.tail = FALSE)
```

```
# 0.4455351
```

Exploring qt(): The Inverse Cumulative Function (Quantiles)

The function **qt()** performs the inverse operation of **pt()**. It is known as the inverse [cumulative density function](#) (CDF) or the [quantile](#) function. Instead of taking a t-score and returning a probability, **qt()** takes a probability (or quantile) p and returns the corresponding t-score (the **critical value**) associated with that cumulative area. This is the essential tool for establishing **critical values** needed to construct confidence intervals and define rejection regions in hypothesis testing.

The primary role of **qt()** is to determine the boundary score--the point on the x-axis--below which a specified percentage of the distribution area lies. For example, if we seek the 95th percentile, **qt()** will tell us the exact t-score that separates the bottom 95% of the data from the top 5%. The syntax requires the probability p (often denoted as x in R's documentation) and the **degrees of freedom** df :

qt(p, df)

Understanding **qt()** is synonymous with understanding **critical values**. When calculating a 95% confidence interval, we need the t-scores that cut off the top 2.5% and the bottom 2.5% of the distribution. For a distribution with 20 **degrees of freedom**, we can find the 99th, 95th, and 90th percentile t-scores, which are critical for various levels of significance (α). These values are foundational for making statistical decisions.

Find the t-score of the 99th quantile of the Student t distribution with df = 20

```
qt(.99, df = 20)
```

```
# 2.527977
```

Find the t-score of the 95th quantile of the Student t distribution with df = 20

```
qt(.95, df = 20)
```

```
# 1.724718
```

Find the t-score of the 90th quantile of the Student t distribution with df = 20

```
qt(.9, df = 20)
```

```
# 1.325341
```

It is important to recognize that the **critical values** determined by **qt()** directly correspond to the values found in traditional printed t-tables used in statistics textbooks. Furthermore, these values can be compared to the critical values derived from the equivalent functions for the standard normal distribution (e.g., **qnorm()**), demonstrating the convergence principle of the t-distribution

when the **degrees of freedom** are sufficiently high.

Generating Random Data with rt()

The final function, **rt()**, is used for generating pseudo-**random variables** that follow the specified [Student t distribution](#). This capability is vital for simulation studies, Monte Carlo analysis, bootstrapping, and verifying statistical concepts empirically. By repeatedly sampling from the t-distribution, researchers can model real-world phenomena where data may exhibit heavier tails than the normal distribution.

The **rt()** function requires two arguments: *n*, which specifies the number of **random variables** (the vector length) to be generated, and *df*, the **degrees of freedom** that define the shape and spread of the distribution. The output is a vector containing *n* independently and identically distributed (i.i.d.) random numbers conforming to the specified t-distribution parameters. The syntax is simply:

rt(n, df)

A few immediate examples demonstrate the output of **rt()**, showing how a small sample of t-distributed values is generated. Due to the nature of **random variables**, each execution of this code will yield a different set of numbers, though they will collectively adhere to the shape dictated by the **degrees of freedom**.

Generate a vector of 5 random variables that follow a Student t distribution

with df = 20

rt(n = 5, df = 20)

```
# -1.7422445 0.9560782 0.6635823 1.2122289 -0.7052825
```

A more powerful application involves generating large samples to visually compare how the **degrees of freedom** influence the distribution's shape. Generating 1000 observations for a 'narrow' distribution (high *df*=40) and 1000 observations for a 'wide' distribution (low *df*=5) allows us to plot histograms side-by-side, clearly illustrating the impact of *df* on the variance and kurtosis of the sampled data.

Generate a vector of 1000 random variables that follow a Student t distribution

with df = 40 (Narrower distribution)

narrowDistribution <- rt(1000, 40)


Generate a vector of 1000 random variables that follow a Student t distribution

with df = 5 (Wider distribution, heavier tails)

wideDistribution <- rt(1000, 5)

```
# Generate two histograms to view these two distributions side by side, and specify
# 50 bars in histogram for granularity
par(mfrow=c(1, 2)) # One row, two columns layout
hist(narrowDistribution, breaks=50, xlim = c(-6, 4))
hist(wideDistribution, breaks=50, xlim = c(-6, 4))
```

The resulting histograms vividly demonstrate the theoretical relationship between the degrees of freedom and the distributional spread:



As illustrated, the distribution generated with fewer **degrees of freedom** ($df=5$, the 'wide' distribution) is noticeably more spread out and exhibits thicker, heavier tails compared to the distribution generated with higher degrees of freedom ($df=40$). This visual confirmation reinforces the principle that the t-distribution becomes flatter and more prone to extreme values when based on smaller samples.

Conclusion and Further Resources

The four functions--**dt()**, **pt()**, **qt()**, and **rt()**--form the complete toolkit for working with the **Student t distribution** in the R statistical environment. Whether you are plotting the density curve,

calculating a P-value for a hypothesis test, finding a **critical value** for a confidence interval, or simulating random data, R provides intuitive and powerful capabilities.

We have covered how to use the density function (d), the cumulative probability function (p), the [quantile](#) function (q), and the random generation function (r). This framework is consistently applied across almost all probability distributions implemented in R, including the Normal, Binomial, and Chi-Squared distributions. By applying the principles learned here, you are well-equipped to explore other fundamental distributions in statistical modeling.

Further Reading:

[A Guide to dnorm, pnorm, qnorm, and rnorm in R](#)

[A Guide to dbinom, pbinom, qbinom, and rbinom in R](#)