Continue

T distribution degrees of freedom

A student’s t-distribution is a bell-shaped probability distribution symmetrical about its mean. It is regarded as the most suitable distribution to use in the construction of confidence intervals in the following instances: When dealing with small samples of less than 30 elements. When the population variance is unknown. When the distribution involved is either normal or approximately normal. Apart from being used in the construction of confidence intervals, a t-distribution is used to test the following: Single population mean. The differences between two population means. The mean difference between paired (dependent) populations. The population correlation coefficient. In the absence of explicit normality of a given distribution, a t-distribution may still be appropriate for use if the sample size is large enough for the central limit theorem to be applied. In such a case, the distribution is considered approximately normal. A t-statistic, also called the t-score, is given by: $$ t = \cfrac {(x – \mu)}{\left(\cfrac {S}{\sqrt n} \right)} $$ Where: \(x\) = Sample mean. \(μ\) = Population mean. \(S\) = Sample standard deviation. \(n\) = Sample size. Relationship between the t-distribution and the Normal Distribution A t-distribution allows us to analyze distributions that are not perfectly normal. A t-distribution has thicker tails relative to a normal distribution. The shape of a t-distribution is dependent on the number of degrees of freedom. It, therefore, follows that as the number of d.f. increases, the distribution becomes more ‘spiked,’ and its tails become thinner, closer to those of the normal distribution. The t-distribution has the following properties: A t-distribution is symmetrical. It is a bell-shaped distribution that assumes the shape of a normal distribution and has a mean of zero. A t-distribution is defined by one parameter, that is, degrees of freedom (df) \(v= n – 1\), where \(n\) is the sample size. Its \(\text {variance} = \frac {v}{ \left(\frac {v}{2} \right) }\), where \(v\) represents the number of degrees of freedom and \(v ≥ 2\). The variance is greater than 1 at all times. Note, however, that it gets very close to one when there are many degrees of freedom. With a large number of degrees of freedom, a t-distribution resembles a normal distribution. The tails of a t-distribution are fatter and less peaked than those of a normal distribution, indicating more probability in the tails. The shape of a t-distribution changes with the change in the degrees of freedom. The more the d.f. increase, the more the shape of a t-distribution looks like a standard normal distribution. The Degrees of Freedom A t-distribution, just like several other distributions, has only one parameter: The degrees of freedom (d.f.). The number of degrees of freedom refers to the number of independent observations (total number of observations less 1): $$ v = n-1 $$ Hence, a sample of 10 observations or elements would be analyzed using a t-distribution with 9 degrees of freedom. Similarly, a 6 d.f. distribution would be used for a sample size of 7 observations. Notations It is the standard practice for statisticians to use \(t_α\) to represent the t-score with a cumulative probability of \((1 – α)\). Therefore, if we were interested in a t-score with a 0.9 cumulative probability, α would be equal to 1 – 0.9 = 0.1. We would denote the statistic as \(t_0.1\). However, the value of \(t_α\) depends on the number of degrees of freedom and is often written as \(t_{α,n-1}\). For example, we could write \(t_{0.05,2}= 2.92\) where the second subscript (2) represents the number of d.f. Important Relationships Just like the normal distribution, the t-distribution is symmetrical about the mean. As such, $$ t_{\alpha}= -t_{1 – \alpha} \text{ and } t_{1 – \alpha} = -t_{\alpha} $$ The table below represents one-tailed confidence intervals and various probabilities for a range of degrees of freedom. Which of the following statements regarding a t-distribution is most likely correct? A t distribution is less spread out than a standard normal distribution. A t-distribution is symmetric about zero. As the degrees-of-freedom increase, a t-distribution becomes wider and flat. Solution The correct answer is B. A t-distribution is a symmetrical, bell-shaped distribution that looks like a normal distribution and has a mean of zero. A is incorrect. A t-distribution is more spread out than a standard normal distribution. C is incorrect. As the degrees-of-freedom increase, a t-distribution becomes narrower, taller, and approaches a standard normal distribution. Probability distribution This article is about the mathematics of Student's t-distribution. For its uses in statistics, see Student's t-test. Student's t Probability density function Cumulative distribution functionParameters ν > 0 {\displaystyle u >0} degrees of freedom (real, almost always a positive integer)Support x ∈ (− ∞ , ∞) {\displaystyle x\in (-\infty ,\infty)} PDF Γ (ν + 1 2) π ν Γ (ν 2) (1 + x 2 ν) − ν + 1 2 {\displaystyle {\frac {\Gamma \left({\frac {u +1}{2}}\right)}{{\sqrt {\pi u }}\,\Gamma \left({\frac {u }{2}}\right)}}\left(1+{\frac {x^{2}}{u }}\right)^{-{\frac {u +1}{2}}}} CDF 1 2 + x Γ (ν + 1 2) × 2 F 1 (1 2 , ν + 1 2 ; 3 2 ; − x 2 ν) π ν Γ (ν 2) , {\displaystyle {\begin{aligned}& {\frac {1}{2}}+x\Gamma \left({\frac {u +1}{2}}\right)\times \\&\quad {\frac {{}_{2}F_{1}\!\left({\frac {1}{2}},{\frac {u +1}{2}};{\frac {3}{2}};-{\frac {x^{2}}{u }}\right)}{{\sqrt {\pi u }}\,\Gamma \left({\frac {u }{2}}\right)}},\end{aligned}}} where 2 F 1 {\displaystyle {}_{2}F_{1}} is the hypergeometric functionMean 0 {\displaystyle 0} for ν > 1 , {\displaystyle u >1,} otherwise undefinedMedian 0 {\displaystyle 0} Mode 0 {\displaystyle 0} Variance ν ν − 2 {\displaystyle {\frac {u }{u -2}}} for ν > 2 , {\displaystyle u >2,} ∞ for 1 < ν ≤ 2 , {\displaystyle 1 3 , {\displaystyle \ u >3\ ,} otherwise undefinedExcess kurtosis 6 ν − 4 {\displaystyle {\frac {6}{u -4}}} for ν > 4 , {\displaystyle u >4,} ∞ for 2 < ν ≤ 4 , {\displaystyle 2 0 {\displaystyle u >0} , where K ν {\displaystyle K_{u }} is the modified Bessel function of the second kind[1]Expected shortfall μ + s ((ν + [T − 1 (1 − p)] 2) × τ (T − 1 (1 − p)) (ν − 1) (1 − p)) , {\displaystyle \mu +s\left({\frac {{\big (}u +[T^{-1}(1-p)]^{2}{\big)}\times \tau {\big (}T^{-1}(1-p){\big)}}{(u -1)(1-p)}}\right),} where T − 1 {\displaystyle T^{-1}} is the inverse standardized Student t CDF, and τ {\displaystyle \tau } is the standardized Student t PDF.[2] In probability theory and statistics, Student's t distribution (or simply the t distribution) t ν {\displaystyle t_{u }} is a continuous probability distribution that generalizes the standard normal distribution. Like the latter, it is symmetric around zero and bell-shaped. However, t ν {\displaystyle t_{u }} has heavier tails, and the amount of probability mass in the tails is controlled by the parameter ν {\displaystyle u }. For ν = 1 {\displaystyle u =1} the Student's t distribution t ν {\displaystyle t_{u }} becomes the standard Cauchy distribution, which has very "fat" tails; whereas for ν → ∞ {\displaystyle u \to \infty } it becomes the standard normal distribution N (0 , 1) , {\displaystyle {\mathcal {N}}(0,1),} which has very "thin" tails. The name "Student" is a pseudonym used by William Sealy Gosset in his scientific paper publications during his work at the Guinness Brewery in Dublin, Ireland. The Student's t distribution plays a role in a number of widely used statistical analyses, including Student's t-test for assessing the statistical significance of the difference between two sample means, the construction of confidence intervals for the difference between two population means, and in linear regression analysis. In the form of the location-scale t distribution ℓ s t (μ , τ 2 , ν) {\displaystyle \operatorname {ell st} (\mu ,\tau ^{2},u)} it generalizes the normal distribution and also arises in the Bayesian analysis of data from a normal family as a compound distribution when marginalizing over the variance parameter. Student's t distribution has the probability density function (PDF) given by f (t) = Γ (ν + 1 2) π ν Γ (ν 2) (1 + t 2 ν) − (ν + 1) / 2 , {\displaystyle f(t)={\frac {\Gamma \left({\frac {u +1}{2}}\right)}{{\sqrt {\pi u }}\Gamma \left({\frac {u }{2}}\right)}}\left(1+{\frac {t^{2}}{u }}\right)^{-(u +1)/2},} where ν {\displaystyle u } is the number of degrees of freedom, and Γ {\displaystyle \Gamma } is the gamma function. This may also be written as f (t) = 1 ν B (1 2 , ν 2) (1 + t 2 ν) − (ν + 1) / 2 , {\displaystyle f(t)={\frac {1}{{\sqrt {u }}\,\mathrm {B} \left({\frac {1}{2}},{\frac {u }{2}}\right)}}\left(1+{\frac {t^{2}}{u }}\right)^{-(u +1)/2},} where B {\displaystyle \mathrm {B} } is the beta function. In particular for integer valued degrees of freedom ν {\displaystyle u } we have: For ν > 1 {\displaystyle u >1} and even, Γ (ν + 1 2) π ν Γ (ν 2) = 1 2 ν · (ν − 1) · (ν − 3) ⋯ 5 · 3 (ν − 2) · (ν − 4) ⋯ 4 · 2 . {\displaystyle {\frac {\Gamma \left({\frac {u +1}{2}}\right)}{{\sqrt {\pi u }}\,\Gamma \left({\frac {u }{2}}\right)}}={\frac {1}{2{\sqrt {u }}}}\cdot {\frac {(u -1)\cdot (u -3)\cdots 5\cdot 3}{(u -2)\cdot (u -4)\cdots 4\cdot 2}}.} For ν > 1 {\displaystyle u >1} and odd, Γ (ν + 1 2) π ν Γ (ν 2) = 1 π ν · (ν − 1) · (ν − 3) ⋯ 4 · 2 (ν − 2) · (ν − 4) ⋯ 5 · 3 . {\displaystyle {\frac {\Gamma \left({\frac {u +1}{2}}\right)}{{\sqrt {\pi u }}\,\Gamma \left({\frac {u }{2}}\right)}}={\frac {1}{\pi {\sqrt {u }}}}\cdot {\frac {(u -1)\cdot (u -3)\cdots 4\cdot 2}{(u -2)\cdot (u -4)\cdots 5\cdot 3}}.} The probability density function is symmetric, and its overall shape resembles the bell shape of a normally distributed variable with mean 0 and variance 1, except that it is a bit lower and wider. As the number of degrees of freedom grows, the t distribution approaches the normal distribution with mean 0 and variance 1. For this reason ν {\displaystyle u } is also known as the normality parameter.[3] The following images show the density of the t distribution for increasing values of ν . {\displaystyle u .} The normal distribution is shown as a blue line for comparison. Note that the t distribution (red line) becomes closer to the normal distribution as ν {\displaystyle u } increases. Density of the t distribution (red) for 1, 2, 3, 5, 10, and 30 degrees of freedom compared to the standard normal distribution (blue).Previous plots shown in green.1 degree of freedom2 degrees of freedom3 degrees of freedom5 degrees of freedom10 degrees of freedom30 degrees of freedom The cumulative distribution function (CDF) can be written in terms of I, the regularized incomplete beta function. For t > 0 , F (t) = ∫ − ∞ t f (u) d u = 1 − 1 2 I x (t) (ν 2 , 1 2) , {\displaystyle F(t)=\int _{-\infty }^{t}\ f(u)\ \operatorname {d} u~=~1-{\frac {1}{2}}I_{x(t)}\!\left({\frac {\ u \ }{2}},\ {\frac {\ 1\ }{2}}\right)\ ,} where x (t) = ν t 2 + ν . {\displaystyle x(t)={\frac {u }{\ t^{2}+u \ }}~.} Other values would be obtained by symmetry. An alternative formula, valid for t 2 < ν , {\displaystyle \ t^{2} 1 , {\displaystyle u >1 ,} the raw moments of the t distribution are E { T k } = { 0 k odd , 0 < k < ν , 1 π Γ (ν 2) [Γ (k + 1 2) Γ (ν − k 2) ν k 2] k even , 0 < k < ν . {\displaystyle \operatorname {mathbb {E} } \left\{\ T^{k}\ \right\}={\begin{cases}\quad 0&k{\text{ odd }},\quad 04~.} Student's t-distribution with ν {\displaystyle u } degrees of freedom can be defined as the distribution of the random variable T with[5][6] T = Z V / ν = Z ν V , {\displaystyle T={\frac {Z}{\sqrt {V/u }}}=Z{\sqrt {\frac {u }{V}}},} where Z is a standard normal with expected value 0 and variance 1; V has a chi-squared distribution (χ2-distribution) with ν {\displaystyle u } degrees of freedom; Z and V are independent; A different distribution is defined as that of the random variable defined, for a given constant μ, by (Z + μ) ν V . {\displaystyle (Z+\mu){\sqrt {\frac {u }{V}}}.} This random variable has a noncentral t-distribution with noncentrality parameter μ. This distribution is important in studies of the power of Student's t-test. Suppose X1, ..., Xn are independent realizations of the normally-distributed, random variable X, which has an expected value μ and variance σ2. Let X ¯ n = 1 n (X 1 + ⋯ + X n) {\displaystyle {\overline {X}}_{n}={\frac {1}{n}}(X_{1}+\cdots +X_{n})} be the sample mean, and s 2 = 1 n − 1 ∑ i = 1 n (X i − X ¯ n) 2 {\displaystyle s^{2}={\frac {1}{n-1}}\sum _{i=1}^{n}\left(X_{i}-{\overline {X}}_{n}\right)^{2}} be an unbiased estimate of the variance from the sample. It can be shown that the random variable V = (n − 1) s 2 σ 2 {\displaystyle V=(n-1){\frac {s^{2}}{\sigma ^{2}}}} has a chi-squared distribution with ν = n − 1 {\displaystyle u =n-1} degrees of freedom (by Cochran's theorem).[7] It is readily shown that the quantity Z = (X ¯ n − μ) n σ {\displaystyle Z=\left({\overline {X}}_{n}-\mu \right){\frac {\sqrt {n}}{\sigma }}} is normally distributed with mean 0 and variance 1, since the sample mean X ¯ n {\displaystyle {\overline {X}}_{n}} is normally distributed with mean μ and variance σ2/n. Moreover, it is possible to show that these two random variables (the normally distributed one Z and the chi-squared-distributed one V) are independent. Consequently[clarification needed] the pivotal quantity T ≡ Z V / ν = (X ¯ n − μ) n s , {\textstyle T≡{\frac {Z}{\sqrt {V/u }}}=\left({\overline {X}}_{n}-\mu \right){\frac {\sqrt {n}}{s}},} which differs from Z in that the exact standard deviation σ is replaced by the sample standard error s, has a Student's t-distribution as defined above. Notice that the unknown population variance σ2 does not appear in T, since it was in both the numerator and the denominator, so it canceled. Gosset intuitively obtained the probability density function stated above, with ν {\displaystyle u } equal to n − 1, and Fisher proved it in 1925.[8] The distribution of the test statistic T depends on ν {\displaystyle u } , but not μ or σ; the lack of dependence on μ and σ is what makes the t-distribution important in both theory and practice. The t distribution arises as the sampling distribution of the t statistic. Below the one-sample t statistic is discussed, for the corresponding two-sample t statistic see Student's t-test. Let x 1 , ... , x n ~ N (μ , σ 2) {\displaystyle \ x_{1},\ldots ,x_{n}\sim {\mathcal {N}}(\mu ,\sigma ^{2})\ } be independent and identically distributed samples from a normal distribution with mean μ {\displaystyle \mu } and variance σ 2 . {\displaystyle \ \sigma ^{2}~.} The sample mean and unbiased sample variance are given by: x ¯ = x 1 + ⋯ + x n n , s 2 = 1 n − 1 ∑ i = 1 n (x i − x ¯) 2 . {\displaystyle {\begin{aligned}{\bar {x}}&={\frac {\ x_{1}+\cdots +x_{n}\ }{n}}\ ,\\[5pt]s^{2}&={\frac {1}{\ n-1\ }}\ \sum _{i=1}^{n}(x_{i}-{\bar {x}})^{2}~.\end{aligned}}} The resulting (one sample) t statistic is given by t = x ¯ − μ s / n ~ t n − 1 . {\displaystyle t={\frac {{\bar {x}}-\mu }{\ s/{\sqrt {n\ }}\ }}\ ~\mathrm {t} _{n-1}~.} and is distributed according to a Student's t distribution with n − 1 {\displaystyle \ n-1\ } degrees of freedom. Thus for inference purposes the t statistic is a useful "pivotal quantity" in the case when the mean and variance (μ , σ 2) {\displaystyle (\mu ,\sigma ^{2})} are unknown population parameters, in the sense that the t statistic has then a probability distribution that depends on neither μ {\displaystyle \mu } nor σ 2 . {\displaystyle \ \sigma ^{2}~.} Instead of the unbiased estimate s 2 {\displaystyle \ s^{2}\ } we may also use the maximum likelihood estimate s M L 2 = 1 n ∑ i = 1 n (x i − x ¯) 2 {\displaystyle \ s_{\mathrm {ML}}^{2}={\frac {1 }{n}}\ \sum _{i=1}^{n}(x_{i}-{\bar {x}})^{2}\ } yielding the statistic t M L = x ¯ − μ s M L 2 / n = n n − 1 t . {\displaystyle \ t_{\mathrm {ML}}={\frac {{\bar {x}}-\mu }{\sqrt {s_{\mathrm {ML}}^{2}/n\ }}}={\sqrt {{\frac {n}{n-1}}\ }}\ t~.} This is distributed according to the location-scale t distribution: t M L ~ ℓ s t (0 , τ 2 = n / (n − 1) , n − 1) . {\displaystyle t_{\mathrm {ML}}\sim \operatorname {ell st} (0,\ \tau ^{2}=n/(n-1),\ n-1)~.} The location-scale t distribution results from compounding a Gaussian distribution (normal distribution) with mean μ {\displaystyle \ \mu \ } and unknown variance, with an inverse gamma distribution placed over the variance with parameters a = ν 2 {\displaystyle \ a={\frac {\ u \ }{2}}\ } and b = ν τ 2 2 . {\displaystyle b={\frac {\ u \ \tau ^{2}\ }{2}}~.} In other words, the random variable X is assumed to have a Gaussian distribution with an unknown variance distributed as inverse gamma, and then the variance is marginalized out (integrated out). Equivalently, this distribution results from compounding a Gaussian distribution with a scaled-inverse-chi-squared distribution with parameters ν {\displaystyle u } and τ 2 . {\displaystyle \ \tau ^{2}~.} The scaled-inverse-chi-squared distribution is exactly the same distribution as the inverse gamma distribution, but with a different parameterization, i.e. ν = 2 a , τ 2 = b a . {\displaystyle \ u =2\ a,\;{\tau }^{2}={\frac {\ b\ }{a}}~.} The reason for the usefulness of this characterization is that in Bayesian statistics the inverse gamma distribution is the conjugate prior distribution of the variance of a Gaussian distribution. As a result, the location-scale t distribution arises naturally in many Bayesian inference problems.[9] Student's t distribution is the maximum entropy probability distribution for a random variate X having a certain value of E { ln (ν + X 2) } {\displaystyle \ \operatorname {mathbb {E} } \left\{\ ln(u +X^{2})\ \right\}\ } .[10][clarification needed][better source needed] This follows immediately from the observation that the pdf can be written in exponential family form with ν + X 2 {\displaystyle u +X^{2}} as sufficient statistic. The function A(t | ν) is the integral of Student's probability density function, f(t) between -t and t, for t ≥ 0 . It thus gives the probability that a value of t less than that calculated from observed data would occur by chance. Therefore, the function A(t | ν) can be used when testing whether the difference between the means of two sets of data is statistically significant, by calculating the corresponding value of t and the probability of its occurrence if the two sets of data were drawn from the same population. This is used in a variety of situations, particularly in t tests. For the statistic t, with ν degrees of freedom, A(t | ν) is the probability that t would be less than the observed value if the two means were the same (provided that the smaller mean is subtracted from the larger, so that t ≥ 0). It can be easily calculated from the cumulative distribution function Fν(t) of the t distribution: A (t | ν) = F ν (t) − F ν (− t) = 1 − I ν ν + t 2 (ν 2 , 1 2) , {\displaystyle A(t\mid u)=F_{u }(t)-F_{u }(-t)=1-I_{\frac {u }{u +t^{2}}}\!\left({\frac {u }{2}},{\frac {1}{2}}\right),} where Ix(a, b) is the regularized incomplete beta function. For statistical hypothesis testing this function is used to construct the p-value. The noncentral t distribution generalizes the t distribution to include a noncentrality parameter. Unlike the nonstandardized t distributions, the noncentral distributions are not symmetric (the median is not the same as the mode). The discrete Student's t distribution is defined by its probability mass function at r being proportional to:[11] ∏ j = 1 k 1 (r + j + a) 2 + b 2 r = ... , − 1 , 0 , 1 , ... {\displaystyle \prod _{j=1}^{k}{\frac {1}{(r+j+a)^{2}+b^{2}}}\quad \quad r=\ldots ,-1,0,1,\ldots ~.} Here a, b, and k are parameters. This distribution arises from the construction of a system of discrete distributions similar to that of the Pearson distributions for continuous distributions.[12] One can generate Student A(t | ν) samples by taking the ratio of variables from the normal distribution and the square-root of the χ² distribution. If we use instead of the normal distribution, e.g., the Irwin–Hall distribution, we obtain over-all a symmetric 4 parameter distribution, which includes the normal, the uniform, the triangular, the Student t and the Cauchy distribution. This is also more flexible than some other symmetric generalizations of the normal distribution. t distribution is an instance of ratio distributions. The square of a random variable distributed tn is distributed as Snedecor's F distribution F1,n. Student's t distribution generalizes to the three parameter location-scale t distribution ℓ s t (μ , τ 2 , ν) {\displaystyle \operatorname {ell st} (\mu ,\ \tau ^{2},\ u)\ } by introducing a location parameter μ {\displaystyle \ \mu \ } and a scale parameter τ . {\displaystyle \ \tau ~.} With T ~ t ν {\displaystyle \ T\sim t_{u }\ } and location-scale family transformation X = μ + τ T {\displaystyle \ X=\mu +\tau \ T\ } we get X ~ ℓ s t (μ , τ 2 , ν) . {\displaystyle \ X\sim \operatorname {ell st} (\mu ,\ \tau ^{2},\ u)~.} The resulting distribution is also called the non-standardized Student's t distribution. The location-scale t distribution has a density defined by:[13] p (x | ν , μ , τ) = Γ (ν + 1 2) Γ (ν 2) π ν τ (1 + 1 ν (x − μ τ) 2) − (ν + 1) / 2 {\displaystyle p(x\mid u ,\mu ,\tau)={\frac {\ \Gamma \left({\frac {\ u +1\ }{2}}\right)\ }{\ \Gamma \left({\frac {\ u \ }{2}}\right)\ {\sqrt {\pi \ u \ }}\ \tau \ }}\ \left(1+{\frac {\ 1\ }{u }}\ {\frac {\ (x-\mu)^{2}\ }{\ \tau ^{2}\ }}\ \right)^{-(u +1)/2}\ } Equivalently, the density can be written in terms of τ 2 {\displaystyle {\tau ^{2}} : p (x | ν , μ , τ 2) = Γ (ν + 1 2) Γ (ν 2) π ν τ 2 (1 + 1 ν (x − μ) 2 τ 2) − (ν + 1) / 2 {\displaystyle \ p(x\ \mid \ u ,\ \mu ,\ \tau ^{2})={\frac {\ \Gamma \left({\frac {u +1}{2}}\right)\ }{\ \Gamma \left({\frac {\ u \ }{2}}\right)\ {\sqrt {\pi \ u \ \tau ^{2}}}\ }}\ \left(\ 1+{\frac {\ 1\ }{\ u \ }}\ {\frac {\ (x-\mu)^{2}\ }{\ \tau ^{2}\ }}\ \right)^{-(u +1)/2}\ } Other properties of this version of the distribution are:[13] E { X } = μ for ν > 1 , var { X } = τ 2 ν ν − 2 for ν > 2 , mode { X } = μ . {\displaystyle \begin{aligned}\operatorname {mathbb {E} } \{\ X\ \}&=\mu &\{\text{ for }}u >1\ ,\\\operatorname {var} \{\ X\ \}&=\tau ^{2}{\frac {u }{u -2}}&\{\text{ for }}u >2\ ,\\\operatorname {mode} \{\ X\ \}&=\mu ~.\end{aligned}}} If X {\displaystyle \ X\ } follows a location-scale t distribution X ~ ℓ s t (μ , τ 2 , ν) {\displaystyle \ X\sim \operatorname {ell st} \left(\mu ,\ \tau ^{2},\ u \right)\ } then for ν → ∞ {\displaystyle \ u \rightarrow \infty \ } X {\displaystyle \ X\ } is normally distributed X ~ N (μ , τ 2) {\displaystyle X\sim \mathrm {N} \left(\mu ,\tau ^{2}\right)} with mean μ {\displaystyle \mu } and variance τ 2 . {\displaystyle \ \tau ^{2}~.} The location-scale t distribution ℓ s t (μ , τ 2 , ν = 1) {\displaystyle \ \operatorname {ell st} \left(\mu ,\ \tau ^{2},\ u =1\right)\ } with degree of freedom ν = 1 {\displaystyle \ u =1\ } is equivalent to the Cauchy distribution C a u (μ , τ) . {\displaystyle \ \mathrm {Cau} \left(\mu ,\tau \right)~.} The location-scale t distribution ℓ s t (μ = 0 , τ 2 = 1 , ν) {\displaystyle \ \operatorname {ell st} \left(\mu =0,\ \tau ^{2}=1,\ u \right)\ } with μ = 0 {\displaystyle \ \mu =0\ } and τ 2 = 1 {\displaystyle \ \tau ^{2}=1\ } reduces to the Student's t distribution t ν . {\displaystyle \ t_{u }~.} Student's t distribution arises in a variety of statistical estimation problems where the goal is to estimate an unknown parameter, such as a mean value, in a setting where the data are observed with additive errors. If (as in nearly all practical statistical work) the population standard deviation of these errors is unknown and has to be estimated from the data, the t distribution is often used to account for the extra uncertainty that results from this estimation. In most such problems, if the standard deviation of the errors were known, a normal distribution would be used instead of the t distribution. Confidence intervals and hypothesis tests are two statistical procedures in which the quantiles of the sampling distribution of a particular statistic (e.g. the standard score) are required. In any situation where this statistic is a linear function of the data, divided by the usual estimate of the standard deviation, the resulting quantity can be rescaled and centered to follow Student's t distribution. Statistical analyses involving means, weighted means, and regression coefficients all lead to statistics having this form. Quite often, textbook problems will treat the population standard deviation as if it were known and thereby avoid the need to use the Student's t distribution. These problems are generally of two kinds: (1) those in which the sample size is so large that one may treat a data-based estimate of the variance as if it were certain, and (2) those that illustrate mathematical reasoning, in which the problem of estimating the standard deviation is temporarily ignored because that is not the point that the author or instructor is then explaining. A number of statistics can be shown to have t distributions for samples of moderate size under null hypotheses that are of interest, so that the t distribution forms the basis for significance tests. For example, the distribution of Spearman's rank correlation coefficient ρ, in the null case (zero correlation) is well approximated by the t distribution for sample sizes above about 20.[citation needed] Suppose the number A is so chosen that P { − A < T < A } = 0.9 , {\displaystyle \ \operatorname {\mathbb {P} } \left\{\ -A