Published in final edited form as:

Commun Stat Theory Methods. 2017 ; 46(19): 9834–9841. doi:10.1080/03610926.2016.1222433.

A general approximation to quantiles

Chang Yu* and

Department of Biostatistics, Vanderbilt University School of Medicine, Nashville, TN 37232, U.S.A

Daniel Zelterman

Department of Biostatistics, Yale University School of Public Health New Haven, CT 06520, U.S.A

Abstract

For many continuous distributions, a closed-form expression for their quantiles does not exist. Numerical approximations for their quantiles are developed on a distribution-by-distribution basis. This work develops a general approximation for quantiles using the Taylor expansion. Our method only requires that the distribution has a continuous probability density function and its derivatives can be derived to a certain order (usually 3 or 4). We demonstrate our unified approach by approximating the quantiles of the normal, exponential, and chi-square distributions. The approximation works well for these distributions.

Keywords

Inverse distribution function; Percentage points; Taylor expansion

1 Introduction

For many continuous distributions, a closed-form expression for their quantiles does not exist. Examples include the normal, Student t, and chi-square distributions, among many others. Numerical approximations for their quantiles are developed on a distribution-by-distribution basis. Abramowitz and Stegun (1970, page 933) provide approximations for the quantiles of the normal distribution. For other distributions, e.g. the Student t distribution and the chi-square distribution, they express their quantile as functions of the normal quantile.

This work develops a general approximation to quantiles based on the Taylor expansion. Our method only requires that the distribution has a continuous probability density function (PDF) and its derivatives can be derived up to a certain order. This is the case for many of the common distributions. Usually order 3 or 4 is sufficient for the approximation to be accurate as we shall demonstrate in a few examples. Kafadar and Tukey (1988) suggest the quantiles for the t distribution can be approximated as a linear function of $\log_{\mathcal{A}}(p)$ with base $G = 10^{-0.1}$, where p is the cumulative probability of the distribution at the quantile. We show that their result is a special case of our general approximation. In this work, the log function is the natural logarithm unless the base is explicitly specified.

*Contact for corresponding author: chang.yu@vanderbilt.edu, Tel: (615)-322-8422, Fax: (615)-343-4924.

We demonstrate our method using the normal, exponential, and chi-square distributions. For the normal distribution, there have been approximations to its quantiles (Abramowitz and Stegun, 1970, page 933; and Wichura, 1988). For the exponential distribution, there is a closed-form expression. It is difficult to obtain the quantiles for the noncentral chi-square distribution, although there are approximations in the literature (Abramowitz and Stegun, 1970, page 941–2). These approximations are rarely used, and in the popular statistical software R, an unpublished algorithm is used to compute the chi-square quantiles (R Core Team, 2014). Recently, Ferreira, Oliveira and Toledo (2015) developed a R package “ncg” to compute the noncentral gamma distribution-related quantities. The demonstration here is not to show which approximation is more accurate, rather, to demonstrate our general approximation can be flexibly implemented for all these distributions and desired accuracy of the approximation can be achieved. Through these examples we demonstrate that our general approximation can serve as a unified approach for quantile approximation for a broad range of distributions. Furthermore, in the event we have a new continuous distribution, for which a closed-form expression for its quantile does not exist, our method can be readily implemented to obtain its quantiles.

The rest of the manuscript is organized as follows. In Section 2, we derive our approximation and we construct an algorithm to implement the approximation. We leave the error evaluation of our approximation to the Appendix. In Section 3 we demonstrate the method using three common distributions. The R programs that implement the approximation can be obtained from the authors. We conclude the manuscript with a discussion.

2 A general approximation for the quantile function

Suppose X is a random variable with probability density function $f(\cdot)$ and cumulative distribution function $F(\cdot)$. Assume that $f(x)$ is continuous for all x for which it is defined and let $x_p = F^{-1}(p)$ denote the quantile function for $0 < p < 1$, but x_p does not have a closed form expression. Let $\beta = -\log(p)$.

The first five terms of the Taylor series expansion of x_p in $\beta = -\log(p)$ about $\beta_0 = -\log(p_0)$ are

$$\begin{aligned} x_p = & x_{p_0} - (\partial/\partial\beta) x_p \log(p/p_0) + (\partial/\partial\beta)^2 x_p (\log(p/p_0))^2/2! - \\ & (\partial/\partial\beta)^3 x_p (\log(p/p_0))^3/3! + (\partial/\partial\beta)^4 x_p (\log(p/p_0))^4/4! + \dots, \end{aligned} \quad (1)$$

where the derivatives are evaluated at x_{p_0} . Successive derivatives of x_p with respect to β are

$$\begin{aligned}
(\partial/\partial\beta)x_p &= -e^{-\beta}/f(x_p) \\
(\partial/\partial\beta)^2x_p &= \left[e^{-\beta} - f'(x_p) \left(\frac{\partial x_p}{\partial\beta} \right)^2 \right] / f(x_p) \\
(\partial/\partial\beta)^3x_p &= \left[-e^{-\beta} - 3f'(x_p) \frac{\partial x_p}{\partial\beta} \frac{\partial^2 x_p}{\partial\beta^2} - f''(x_p) \left(\frac{\partial x_p}{\partial\beta} \right)^3 \right] / f(x_p) \\
(\partial/\partial\beta)^4x_p &= \left[e^{-\beta} - 4f'(x_p) \frac{\partial x_p}{\partial\beta} \frac{\partial^3 x_p}{\partial\beta^3} - 3f''(x_p) \left(\frac{\partial^2 x_p}{\partial\beta^2} \right)^2 - 6f'''(x_p) \left(\frac{\partial x_p}{\partial\beta} \right)^2 \frac{\partial^2 x_p}{\partial\beta^2} - f''''(x_p) \left(\frac{\partial x_p}{\partial\beta} \right)^4 \right] / f(x_p),
\end{aligned}$$

where $f(\cdot)$ and its derivatives are evaluated at x_p .

The Taylor series expansion of x_p about $\beta_0 = -\log(p_0)$ can then be obtained by plugging these derivatives of x_p into (1). This approximation is a polynomial of $\log(p)$. The coefficients in the polynomial are functions of the probability density function and its derivatives evaluated at x_{p_0} . Kafadar and Tukey (1988) propose to approximate the quantiles

of the Student t distribution using a linear function of $\log_G(p)$ with base $G = 10^{-0.1}$. Their approximation is a special case of (1) since $\log_G(p) = (-10/\log(10)) \log(p)$.

When we calculate x_p using (1), we start with a p_0 at which the quantile x_{p_0} is known, e.g.

$p_0 = 0, 0.5$, or 1 . We also want to choose p_0 as close to p as possible. Suppose we start at $p_0 = 0.5$, we first calculate x_{p_1} for $p_1 = p_0 + h$, where h is a very small quantity. Then we can

repeatedly apply (1) to obtain x_{p_k} from $x_{p_{k-1}}$, where $p_k = p_{k-1} + h$, and $x_{p_{k-1}}$ is calculated in the previous step, for $k = 2, 3, \dots, K$. The step length h is chosen so that there are $K = (p - p_0)/h$ steps in the algorithm and $p_K = p$ is the target at which we want to calculate the quantile. This algorithm is an analogy to a walk in which we move K steps from p_0 to p_K with each step having length of h .

The accumulated error in approximation x_{p_k} depends on the error in each step, i.e. the error in Taylor expansion (1), and the number of steps k that is needed to reach p_k from p_0 . The error in the Taylor expansion depends on the step length h and the highest order term in expansion (1). In the Appendix, we show that the error in x_{p_k} is in the order of

$$\epsilon = O([h/\min(p_{k-1}, p_0)]^I), \quad (2)$$

where I is the highest order in the Taylor expansion. Specifically, I is 4 in our approximation (1), thus the error is bounded by $O([h/\min(p_{k-1}, p_0)]^4)$. To minimize the accumulated error, we would naturally choose p_0 close to p_k and h that is sufficiently small. The examples in Section 3 demonstrate the impact of h on the accumulated error.

3 Numerical examples

In this section, we implement our approximation for the normal, exponential, and chi-square distributions. In R, there are functions to approximate the normal and chi-square quantiles. For the exponential distribution, there is a closed-form expression for its quantiles. These provide a benchmark to evaluate our general approximation for these three distributions.

EXAMPLE 1: For a standard normal distribution, its PDF is

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

with derivatives of $f'(x) = -x f(x)$, $f''(x) = (x^2 - 1)f(x)$, $f'''(x) = (3x - x^3)f(x)$, and so on.

There are a range of approximations for the normal quantiles, including those described in Abramowitz and Stegun (1970, page 933), and Wichura (1988), among others. Abramowitz and Stegun (1970) describe approximations that are rational functions of $\{-2 \log(p)\}^{1/2}$. We compare our approximation (1) with the normal quantile function *qnorm* in R. This R function approximates the normal quantile by piece-wise minimax rational functions. The minimax algorithm minimizes the maximum deviation between the rational function and the function to be approximated by iterating through a set of reference points in the range of the function being defined. Curtis and Osborne (1966) provides a survey on constructing minimax rational approximations to functions. The R approximations are ratio of two polynomials of $(p - 0.5)$. The order of the polynomials depends on the range of p (Wichura, 1988).

Table 1 summarizes our approximation in comparison with R function *qnorm*. Using the 5-term expansion (1), our calculated quantiles agree with the R function to 5 or more decimal places using a step length $h = 0.0001$. The approximation is much more accurate if we use $h = 0.00001$. In the calculations, we start at $p_0 = 0.5$ and move toward the target p by steps of length h .

EXAMPLE 2: In this example, we calculate the quantile for the exponential distribution. The purpose is to examine how accurate approximation (1) is for a distribution whose quantile has a closed-form expression. For an exponential distribution with rate λ , its PDF is $f(x; \lambda) = \lambda e^{-\lambda x}$ and its distribution function is $F(x; \lambda) = 1 - e^{-\lambda x}$ for $x \geq 0$. Its p -th quantile is $x_p = -\log(1 - p)/\lambda$. The difference between our approximation and the exact quantile is less than 10^{-5} for $0.0005 \leq p \leq 0.9995$ using a step length $h = 0.0001$. When we use $h = 0.00001$, the difference is less than 10^{-9} . In these calculations, we start at $p_0 = 0.5$ and the quantile is $\log(2)/\lambda$.

EXAMPLE 3: In this example, we calculate the quantile for the chi-square distribution. The (central) chi-square distribution with ν degrees of freedom (DF) has a PDF

$$f(x; \nu) = \frac{1}{2^{\nu/2} \Gamma(\nu/2)} x^{\nu/2 - 1} e^{-x/2} \quad (3)$$

for $x = 0$. The PDF of a noncentral chi-square distribution with ν DF and noncentrality parameter $\lambda = 0$ can be expressed as a weighted sum of central chi-square densities with the probabilities of a Poisson distribution with mean $\lambda/2$ as the weights

$$g(x; \nu, \lambda) = \sum_{j=0}^{\infty} \left\{ \frac{(\lambda/2)^j}{j!} e^{-\lambda/2} \right\} f(x; \nu + 2j) \quad (4)$$

(*c.f.* Johnson, Kotz, and Balakrishnan, 1995, page 436; and Fisher, 1928). This distribution is more difficult to work with than the two previous examples. Its only known finite quantile is the 0th quantile $x_0 = 0$. For these chi-square PDF's, their K^{th} or higher order derivatives do not exist at $x = 0$, where K is the smallest integer greater than $(\nu/2 - 1)$. In order to apply approximation (1), we have to start at a tiny quantity $\epsilon > 0$. In the following calculations, we started at $\epsilon = h/2$. There are various ways to obtain the quantile at ϵ for the initial step in the algorithm. For example, there are approximations in Abramowitz and Stegun (1970, page 941) that are summarized in the Discussion section. Or more accurately, we can implement the series expansion of the chi-square distribution function (Abramowitz and Stegun, 1970, results 26.4.4 and 26.4.5, page 941) and conduct a root-solving. Here for the purpose of illustrating our algorithm, we use the R function *qchisq* to obtain this initial quantile at ϵ .

The central chi-square density (3) has derivatives,

$$\begin{aligned} f'(x; \nu) &= (\nu/2 - 1)/x - 1/2) f(x; \nu), \\ f''(x; \nu) &= -(\nu/2 - 1)/x^2 f(x; \nu) + (\nu/2 - 1)/x - 1/2) f'(x; \nu), \\ f'''(x; \nu) &= 2(\nu/2 - 1)/x^3 f(x; \nu) - 2(\nu/2 - 1)/x^2 f'(x; \nu) + (\nu/2 - 1)/x - 1/2) f''(x; \nu), \end{aligned}$$

and so on. Thus, the derivatives of the non-central chi-square density $g(x; \nu, \lambda)$ can also be expressed as weighted sum of these central chi-square derivatives. These are sums of infinite terms like in (4). However, the summands diminish quickly since the Poisson probability weight becomes small as j increases beyond two or three time of the Poisson mean $\lambda/2$. In general, the limit of j needs to be sufficiently large relative to $\lambda/2$ so that the omitted terms are negligible. Our numerical evaluation indicate that a limit of 30 for j is sufficiently large for our examples. We have included this limit as a variable in our R function so that users can choose their own limit based on the non-centrality parameter λ .

The absolute difference between our approximation and R function *qchisq* is less than 10^{-5} except for $p = 0.9990$ and $p = 0.9995$ when we used $h = 0.0001$ as shown in Table 2 for the noncentral chi-square distribution with $\nu = 10$ DF and noncentrality parameter $\lambda = 2$. The difference is 10^{-2} to 10^{-1} times smaller when $h = 0.00001$. As p increases, the quantile also increases. Therefore we also calculated the relative difference. The relative difference was in the range of 10^{-7} to 10^{-4} percent and 10^{-8} to 10^{-6} percent for $h = 0.0001$ and $h = 0.00001$, respectively. The difference between our approximation and R function *qchisq* for the (central) chi-square quantiles was similarly small as for the non-central chi-square quantiles. A referee brought the R package “ncg” (Oliveira and Ferreira, 2013; Ferreira, Oliveira and Toledo, 2015) for computing the noncentral gamma distribution-related quantities to our

attention. Non-central chi-square variable (X) is a scaled non-central gamma variable (Z) through a transformation $Z = X/2$. Based on this relationship, we used R function *qgammanc* in the “ncg” package to obtain these quantiles. The absolute difference between our approximation using $h = 0.00001$ and R function *qgammanc* is less than 10^{-5} for the noncentral chi-square distribution with $\nu = 10$ DF and noncentrality parameter $\lambda = 2$, indicating an excellent agreement between the two methods.

We conducted all the computing on a Hewlett Packard EliteBook laptop with a 2.70GHz Intel i7-3820QM CPU running Microsoft Windows 7. We implemented our method in R (version 3.1.0). We used R function *system.time* to obtain the user, system, and elapsed time for each method. The computing time of obtaining the 14 quantiles is reported in Table 3. All the methods were very fast, however, our method with step $h = 0.00001$ was slower than the other methods. We would like to note that there is a slight in-comparability between our method and the R functions for the comparisons reported in Table 3. Our method actually calculates all the quantiles at increase of h from a known starting quantile, e.g. the median 0 ($p = 0.5$) for normal distribution, to the target p . To put the comparison in a more comparable scale, the computing times (user, system, elapsed) to obtain the quantiles for the same series at increase of $h = 0.0001$ are (0.16, 0.00, 0.15) second and (2.12, 0.00, 2.12) second for the chi-square and the non-central chi-square distributions, respectively, using R function *qchisq*. These are more comparable with our method.

4 Discussion

In R, the noncentral chi-square quantile is computed through inversion of its distribution function. The noncentral chi-square distribution function is computed as a Poisson mixture of central chi-square distribution functions for noncentrality parameter less than 80, and for noncentrality parameter larger than 80, it is computed using an algorithm described in Ding (1992). The central chisquare distribution functions are computed through the gamma distribution since a simple transformation of the central chi-square variate would turn (3) into a gamma density. However, the gamma distribution function was computed using “an unpublished algorithm ‘mainly by Morten Welinder’ (R help file, R Core Team, 2014)”. This is somewhat surprising. Recently, Ferreira, Oliveira and Toledo (2015) developed a R package “ncg” to compute the noncentral gamma distribution-related quantities. The current work provides a clear mathematical basis to approximate quantiles for the chi-square distribution while using it as one example to demonstrate our method.

In Abramowitz and Stegun (1970), the central chi-square quantile with $\nu > 30$ DF is approximated by

$$\chi_p^2 \approx \nu \left[1 - 2/(9\nu) + x_p \sqrt{2/(9\nu)} \right]^3, \quad (5)$$

where χ_p^2 and x_p are the upper p^{th} quantile of the chi-square and the normal distributions, respectively (c.f. Abramowitz and Stegun, 1970, result 26.4.17, p 941). For $\nu = 30$, no

approximations are provided in Abramowitz and Stegun (1970). The noncentral chi-square quantile with ν DF and noncentrality parameter λ is approximated by

$$\chi_p'^2 \approx a \left[1 - \frac{2(1+b)}{9a} + x_p \left(\frac{2(1+b)}{9a} \right)^{1/2} \right]^3, \quad (6)$$

where $a = \nu + \lambda$ and $b = \lambda(\nu + \lambda)$ (c.f. Abramowitz and Stegun, 1970, result 26.4.32, p 942). These results essentially approximate the chi-square quantiles using functions of the normal quantile. We implemented approximations (5) and (6) and they appear to be less accurate than R function *qchisq*. Thus, we decide to compare our approximation (1) with R function *qchisq*.

In conclusion, our general approximation for quantiles works well for the three examples, each of which has its own specific method to calculate or approximate its quantiles. Even though our method does not appear to offer an advantage in computing time when compared with existing R functions for the three examples, with today's desktop computers, the computing time does not appear to be an issue. Our method provides a unified approach to approximating quantiles for various distributions with a continuous PDF. In the event of a new continuous distribution for which no closed-form expression for the quantiles exists, our approximation may be implemented to accurately obtain its quantiles if the derivatives of its PDF can be derived up to a certain order (usually 3 or 4).

Acknowledgment

This work was supported in part by Vanderbilt CTSA grant UL1 TR000445 from NIH/NCATS, R01 CA149633 from NIH/NCI, R21 HL129020, PPG HL108800, R01HL111259, R21HL123829 from NIH/NHLBI, R01HS022093 from NIH/AHRQ (CY) and by grants R01CA131301, R01CA157749, R01CA148996, R01CA168733, and PC50CA196530 from NIH/NCI (DZ).

Appendix: Error in the approximation of x_{pK}

From (1), the error in one step of the approximation from $x_{p_{k-1}}$ to x_{p_k} , where $p_k = p_{k-1} + h$, is

$$\epsilon_k = (\partial/\partial\beta)^5 x_{p_{k-1}} (\log(p_k/p_{k-1}))^{5/5!} + \dots$$

plus higher order terms of $\log(p_k/p_{k-1})$. Using Taylor expansion, we have

$$\log(p_k/p_{k-1}) = h/p_{k-1} - (h/p_{k-1})^2/2 + (h/p_{k-1})^3/3 + \dots$$

Then

$$\epsilon_k \approx (\partial/\partial\beta)^5 x_{p_{k-1}} (h/p_{k-1})^{5/5!}.$$

The accumulated error in the K steps to calculate x_{p_K} from x_{p_0} is

$$\epsilon = \sum_{k=1}^K \epsilon_k \leq K O\left(\left[h/\min(p_{K-1}, p_0)\right]^5\right) \approx O\left(\left[h/\min(p_{K-1}, p_0)\right]^4\right)$$

where K and h are chosen so that $K = (p_K - p_0)/h$ and h is sufficiently small.

In general, if the highest order in the Taylor expansion (1) is I , the error in x_{p_K} is in the order of $O\left[h/\min(p_{K-1}, p_0)\right]^I$.

References

- Abramowitz M and Stegun IA (1970). Handbook of Mathematical Functions : with Formulas, Graphs, and Mathematical Tables. (Applied Mathematics Series, No. 55), Washington, DC: National Bureau of Standards.
- Curtis A and Osborne MR (1966). The construction of minimax rational approximations to functions. The Computer Journal 9(3): 286–293.
- Ding CG (1992). Algorithm AS 275: Computing the non-central chi-squared distribution function. Applied Statistics 41: 478–482.
- Ferreira DF, Oliveira IRC and Toledo FH (2015). Package ‘ncg’. R help file. <https://cran.r-project.org/web/packages/ncg/ncg.pdf>.
- Fisher RA (1928). The general sampling distribution of the multiple correlation coefficient. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character, Vol. 121, No. 788 (12 1, 1928), pp. 654–673.
- Johnson NL, Kotz S, and Balakrishnan N (1995). Continuous Univariate Distributions. Volume 2, Second edition. New York: John Wiley & Sons.
- Kafadar K and Tukey JW (1988). A bidec t table. Journal of the American Statistical Association 83: 532–539.
- Oliveira IRC and Ferreira DF (2013). Computing the noncentral gamma distribution, its inverse and the noncentrality parameter. Computational Statistics 28: 1663–80.
- R Core Team (2014). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria URL <http://www.R-project.org/>.
- Wichura MJ (1988). Algorithm AS 241: The percentage points of the normal distribution. Applied Statistics 37: 477–484.

Table 1:

Quantiles calculated using approximation (1) for the standard normal distribution in comparison with R function *qnorm*.

p	R function	Approximation (1)	Difference *	Approximation (1)	Difference *
	<i>qnorm</i>	(h=0.0001)	($\times 10^5$)	(h=0.00001)	($\times 10^{10}$)
0.0005	-3.2905	-3.2905	0.0004	-3.2905	0.0126
0.0010	-3.0902	-3.0902	0.0000	-3.0902	0.0045
0.0050	-2.5758	-2.5758	0.0000	-2.5758	0.0009
0.0100	-2.3263	-2.3263	0.0000	-2.3263	0.0005
0.0500	-1.6449	-1.6449	0.0000	-1.6449	0.0006
0.1000	-1.2816	-1.2816	0.0000	-1.2816	0.0004
0.3000	-0.5244	-0.5244	0.0000	-0.5244	0.0001
0.7000	0.5244	0.5244	0.0000	0.5244	0.0001
0.9000	1.2816	1.2816	0.0000	1.2816	-0.0004
0.9500	1.6449	1.6449	0.0000	1.6449	-0.0023
0.9900	2.3263	2.3263	0.0000	2.3263	-0.0077
0.9950	2.5758	2.5758	-0.0003	2.5758	-0.0185
0.9990	3.0902	3.0902	-0.1436	3.0902	-1.7221
0.9995	3.2905	3.2905	-1.8787	3.2905	-25.1879

* Difference from R function *qnorm*

Table 2:

Quantiles calculated using approximation (1) for the noncentral chi-square distribution with $\nu = 10$ DF and noncentrality parameter $\lambda = 2$ in comparison with R function *qchisq*.

p	R function	Approximation (1)	Difference *	Approximation (1)	Difference *
	<i>qchisq</i>	(h=0.0001)	($\times 10^5$)	(h=0.00001)	($\times 10^5$)
0.0005	1.5414	1.5414	0.2167	1.5414	0.0127
0.0010	1.8012	1.8012	0.1303	1.8012	0.0076
0.0050	2.6229	2.6229	0.0404	2.6229	0.0023
0.0100	3.1103	3.1103	0.0249	3.1103	0.0014
0.0500	4.7796	4.7796	0.0088	4.7796	0.0005
0.1000	5.8929	5.8929	0.0060	5.8929	0.0003
0.3000	8.7711	8.7711	0.0040	8.7711	0.0002
0.7000	14.1344	14.1344	0.0054	14.1344	0.0003
0.9000	19.0887	19.0887	0.0130	19.0887	0.0008
0.9500	21.8055	21.8055	0.0241	21.8055	0.0014
0.9900	27.5151	27.5151	0.1077	27.5151	0.0062
0.9950	29.8092	29.8092	0.2111	29.8092	0.0120
0.9990	34.8901	34.8900	2.6279	34.8901	0.0570
0.9995	36.9975	36.9973	23.9978	36.9975	0.1144

* Difference from R function *qchisq*

Table 3:

Computing time to obtain the 14 quantiles shown in Table 1 for the normal distribution and for the exponential and the chi-square distributions.

Distribution	Method	Computing time (second)		
		User	System	Elapsed
Normal (0,1)	Approximation (1), $h=0.0001$	0.25	0	0.25
	Approximation (1), $h=0.00001$	2.47	0	2.47
	R function <i>qnorm</i>	0.00	0.00	0.29
Exponential (1)	Approximation (1), $h=0.0001$	0.25	0	0.25
	Approximation (1), $h=0.00001$	2.42	0	2.41
	R function <i>qexp</i>	0	0	0
Central Chi-square 10 DF	Approximation (1), $h=0.0001$	0.75	0.00	0.75
	Approximation (1), $h=0.00001$	7.38	0.00	7.38
	R function <i>qchisq</i>	0	0	0
Non-central Chi-square 10 DF $\lambda = 2$	Approximation (1), $h=0.0001$	0.78	0.00	0.78
	Approximation (1), $h=0.00001$	7.70	0.02	7.71
	R function <i>qchisq</i>	0.01	0.00	0.02