ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

The Mellin Transform to Manage Quadratic Forms in Normal Random Variables

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:
Gardini, A., Greco, F., Trivisano, C. (2022). The Mellin Transform to Manage Quadratic Forms in Normal Random Variables. JOURNAL OF COMPUTATIONAL AND GRAPHICAL STATISTICS, 31(4), 1416-1425 [10.1080/10618600.2022.2034639].

Availability:
This version is available at: https://hdl.handle.net/11585/880552 since: 2022-03-31

Published:
DOI: http://doi.org/10.1080/10618600.2022.2034639

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/).
When citing, please refer to the published version.

(Article begins on next page)

31 July 2024

This is the final peer-reviewed accepted manuscript of:

Aldo Gardini, Fedele Greco, Carlo Trivisano. (2022). "The Mellin Transform to Manage Quadratic Forms in Normal Random Variables". ***Journal of Computational and Graphical Statistics*****, Vol. 31, Issue 4, pp. 1416-1425.**

The final published version is available online at:
https://doi.org/10.1080/10618600.2022.2034639

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/)

When citing, please refer to the published version.

The Mellin transform to manage quadratic forms in normal random variables

Aldo Gardini, Fedele Greco, Carlo Trivisano
Department of Statistical Sciences, University of Bologna, Italy

January 21, 2022

Abstract

The problem of computing the distribution of quadratic forms in normal variables has a long tradition in the statistical literature. Well-established numerical algorithms that deal with this task rely on the inversion of Fourier transforms or series representations. In this paper, the Mellin transform is proposed as a tool to compute both the density and the cumulative distribution functions of a positive definite quadratic form: an outline of the numerical algorithm is presented, providing details on the error analysis. The algorithm's characteristics allow us to propose an efficient way to compute the random variables' quantiles. From the theoretical point of view, the analytic properties of the Mellin transform are exploited to provide a novel representation of the distribution of the ratio of independent quadratic forms as a mixture of beta random variables of the second kind. Moreover, algorithms are proposed for computations related to ratios of both independent and dependent quadratic forms. The methods are tested and compared to popular numerical algorithms in terms of computational times and accuracy. The `R` package `QF` implementing all the proposed algorithms is also made available. Supplementary information is available online.

Keywords: Chi-Square Linear Combination, Integral Transform, Numerical Inversion, Computational Probability

1 Introduction

The study of the distribution of quadratic forms (q.f.s) in Gaussian random variables (r.v.s) is of particular interest since several test statistics and estimators are q.f.s or a ratio of q.f.s. The book by Provost and Mathai (1992) provides a comprehensive overview of the topic, discussing also several applications in Chapter 7. The pioneering paper by Robbins and Pitman (1949) deals with the distributions of positive definite q.f.s, expressing them as mixtures. The fundamental contribution by Ruben (1962) provides the basis for our paper: he expressed the cumulative distribution function (c.d.f.) of a linear combination of central and non-central chi-square r.v.s with positive weights as an infinite linear combination of chi-square c.d.f.s.

A relevant research question concerns the computational aspects related to q.f. distributions: Imhof (1961) and Davies (1980) provided algorithms for computing the c.d.f. of indefinite q.f.s in Gaussian variables based on the numerical inversion of the Fourier transform. On the other hand, Farebrother (1984) proposed an algorithm for computing both the probability density function (p.d.f.) and the c.d.f. of a positive q.f. exploiting the infinite series representation introduced by Ruben (1962). Many of these algorithms are available in the `R` (R Core Team, 2020) package `CompQuadForm` (De Micheaux, 2017). Other works focused on approximating the c.d.f. of q.f.s: for instance, Box (1954) proposed a solution based on moments matching, whereas Kuonen (1999) derived a saddle point approximation.

The ratio of q.f.s also received much attention in the literature. Kim et al. (2006) proposed new computational methods and analytical approximations for the ratio between a chi-squared r.v. and a positive q.f., independently distributed. Usually, computing the c.d.f. of a ratio of dependent q.f.s is led back to the problem of computing the c.d.f. of an indefinite q.f. (Imhof, 1961). Other remarkable works in this framework are Lieberman (1994), which proposed a saddle-point approximation, and Broda and Paolella (2009), which numerically evaluate the p.d.f..

Ratios of dependent q.f.s are widely used to test for the presence of correlation structures in the residuals of a linear regression: Anderson (1948) studied a general form of test-statistic that has such distribution under the null hypothesis. The weights of the q.f. at the numerator depends on the kind of structure assumed for the error: some examples of such tests are provided in Table S1 of the Supplementary material. On the other hand, several examples of tests, characterized by a test statistic distributed as a q.f. in Gaussian r.v.s under the null hypothesis, are reported in Table S2. Moreover, some statistics are q.f.s under the Gaussian assumption, such as the intralclass correlation coefficient, Cronbach's alpha, and the sample variance.

The target in this paper is the exact computation of the p.d.f., c.d.f., and quantile function of positive definite q.f.s and the ratio of positive definite q.f.s in normal r.v.s. Our findings rely on the Mellin transform of the p.d.f.. This mathematical tool enjoys several properties useful in statistics and probability (Epstein, 1948), especially to derive the distribution of product or ratio of random variables with positive support. As an example, Provost and Rudiuk (1994) used the Mellin transform to derive the exact p.d.f. of the ratio of dependent q.f.s, but the obtained expressions are not convenient from a computational perspective. Provost (1989) exploited the Mellin transform to derive the distribution of the sum of independent Gamma r.v.s and their ratio. Our proposal is to retrieve the target quantities by numerically inverting the Mellin transform, controlling for the numerical error of the algorithm.

The proposed algorithms are built starting from the evaluation of the Mellin transform, which is the most computationally intensive step. However, once this task has been accomplished, the information contained in the transform can be exploited for computing the p.d.f. and the c.d.f. with a computational effort that can be sensibly lower than the one required by other existing algorithms: the convenience of our proposal depends on the structure of the q.f. weights, as we will discuss in Section 5. Moreover, the algorithm allows building an efficient routine to evaluate the quantile function by means of a simple Newton-Raphson algorithm. To facilitate the usage of the algorithms, the `R` package `QF` (Gardini et al., 2021) is available, where the routines are implemented in the `C++` language through the `Rcpp` package (Eddelbuettel and François, 2011). In addition, expressing the distribution as an inverse Mellin transform led us to a novel representation of the distribution of the ratio of independent q.f.s as a mixture of beta r.v.s of the second kind.

The rest of the paper is organized as follows. In Section 2, some key concepts about the Mellin transform and the q.f. distribution are briefly outlined. These notions are crucial to develop the numerical algorithm aimed at computing the p.d.f. and the c.d.f. of a positive definite q.f.s (Section 3). Section 4 deals with ratios of q.f.s. In Section 5, the proposed methods are compared to other popular numerical algorithms in terms of computational times and accuracy, and the results of some applications are presented. Finally, concluding remarks are reported in Section 6.

2 Preliminaries

2.1 Basics of the Mellin Transform

The Mellin transform is a mathematical tool largely used throughout the paper. It is a transformation strictly related to the more famous Laplace and Fourier transforms, and its properties have been already exploited in statistics (Epstein, 1948), mainly to deal with the distribution of the product (or ratio) of r.v.s and for its connection to moments.

In what follows, given a random variable X, the p.d.f. and c.d.f. are denoted as $f_X(\cdot)$ and $F_X(\cdot)$, respectively. Moreover, given a function $g(\cdot)$, its Mellin transform (Paris and Kaminski, 2001; Poularikas, 2018) is denoted as $\widehat{g}(\cdot)$. For a r.v. having the positive real axis as support, the Mellin transform of its p.d.f. is defined as:

$$\widehat{f}_X(z) = \int_0^{+\infty} x^{z-1} f_X(x) \mathrm{d}x,$$

where $z = h + iy \in \mathbb{C}$. The density function $f_X(\cdot)$ can be recovered from $\widehat{f}_X(\cdot)$ by the inverse Mellin transform:

$$f_X(x) = \frac{1}{2\pi i} \int_{h-i\infty}^{h+i\infty} x^{-z} \widehat{f}_X(z) \mathrm{d}z,$$

where $h = \Re(z)$ individuates a Bromwich path of integration included in the strip of analyticity of $\widehat{f}_X(\cdot)$, that is denoted as $\mathcal{S}\big(\widehat{f}_X\big)$. Once the strip of analyticity is identified, absolute convergence of both the previous integrals is guaranteed if $h \in \mathcal{S}\big(\widehat{f}_X\big)$.

Some properties of the Mellin transform that turn out to be useful for managing computations related to q.f.s distributions are summarized in the following.

- The Mellin transform is a linear operator, i.e. if $g(x) = af_1(x) + bf_2(x)$, then:

$$\widehat{g}(z) = a\widehat{f}_1(z) + b\widehat{f}_2(z) \text{ with strip of analyticity } \mathcal{S}\big(\widehat{g}\big) = \mathcal{S}\big(\widehat{f}_1\big) \cap \mathcal{S}\big(\widehat{f}_2\big). \tag{1}$$

- Let V_1 and V_2 be two independent r.v.s and let $W = V_1/V_2$. The Mellin transform of $f_W(\cdot)$ can be obtained from $\widehat{f}_{V_1}(\cdot)$ and $\widehat{f}_{V_2}(\cdot)$ as:

$$\widehat{f}_W(z) = \widehat{f}_{V_1}(z)\widehat{f}_{V_2}(2-z) \text{ with strip of analyticity } \mathcal{S}\big(\widehat{f}_W\big) = \mathcal{S}\big(\widehat{f}_{V_1}\big) \cap \mathcal{S}\big(\widehat{f}_{V_2}\big). \tag{2}$$

- In the strip of analyticity, it holds that $\widehat{f}(z) \to 0$ as $\Im(z) \to \infty$, and for a suitable constant $M < \infty$, the following inequality holds:

$$\left|\widehat{f}(z)\right| < M|z|^{-2}. \tag{3}$$

- The general property concerning the Mellin transform of a function's derivative allows formalizing a relationship between the Mellin transform of the c.d.f. $F_X(x)$ and $\widehat{f}_X(z)$:

$$\widehat{f}_X(z) = -(z-1)\widehat{F}_X(z-1),$$

from which it follows that the Mellin transform of the c.d.f. can be expressed in terms of $\widehat{f}_X(z)$:

$$\widehat{F}_X(z-1) = -\frac{\widehat{f}_X(z)}{z-1} \text{ with strip of analyticity } \mathcal{S}\big(\widehat{F}_X\big) = \mathcal{S}\big(\widehat{f}_X\big) \cap \{z : \Re(z) < 1\}, \tag{4}$$

since the strip of analyticity of $-(z-1)^{-1}$ is $\Re(z-1) < 0$.

2.2 Quadratic forms in Gaussian variables

To define a q.f. in Gaussian r.v.s, consider a p-dimensional random vector distributed as a multivariate Gaussian: $\mathbf{X} \sim \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The canonical expression of a q.f. in $\mathbf{X}$ is:

$$Q(\mathbf{X}) = \mathbf{X}^T \mathbf{A} \mathbf{X},$$

where $\mathbf{A} \in \mathbb{R}^{p \times p}$ is a symmetric matrix.

Defining the matrix decomposition $\mathbf{\Sigma} = \mathbf{L}^T\mathbf{L}$, $Q(\mathbf{X})$ can be expressed as a function of the standardized Gaussian random vector $\mathbf{Z} \in \mathbb{R}^p$:

$$Q(\mathbf{X}) = (\mathbf{Z} + \mathbf{L}^{-1}\boldsymbol{\mu})^T\mathbf{L}^T\mathbf{A}\mathbf{L}(\mathbf{Z} + \mathbf{L}^{-1}\boldsymbol{\mu}). \tag{5}$$

A fundamental tool for studying the distribution of $Q(\mathbf{X})$ is the spectral decomposition of the matrix $\mathbf{L}^T\mathbf{A}\mathbf{L}$, i.e. $\mathbf{P}\mathbf{\Lambda}\mathbf{P}^T$, where $\mathbf{P}$ is the orthogonal matrix whose columns are the eigenvectors of $\mathbf{L}^T\mathbf{A}\mathbf{L}$, and $\mathbf{\Lambda} = \text{diag}(\boldsymbol{\lambda})$ is a diagonal matrix. The vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_r, \lambda_{r+1}, \ldots, \lambda_p)$ contains the eigenvalues of the matrix in descending order, where r is the rank of $\mathbf{L}^T\mathbf{A}\mathbf{L}$. If $r < p$, then $\lambda_i = 0$, $i = r+1, \ldots, p$. Defining $\tilde{\boldsymbol{\eta}} = \mathbf{P}^T\mathbf{L}^{-1}\boldsymbol{\mu}$ and the transformed random vector $\mathbf{U} = \mathbf{P}^T\mathbf{Z}$, such that $\mathbb{E}[\mathbf{U}] = \mathbf{0}$ and $\text{Cov}[\mathbf{U}] = \mathbf{I}$, the q.f. can be written as:

$$Q(\mathbf{X}) = (\mathbf{U} + \tilde{\boldsymbol{\eta}})^T\mathbf{\Lambda}(\mathbf{U} + \tilde{\boldsymbol{\eta}}) = \sum_{i=1}^{r} \lambda_i (U_i + \tilde{\eta}_i)^2,$$

where $(U_i + \tilde{\eta}_i)^2 \sim \chi^2_{1,\eta_i}$, i.e. a chi-square r.v. with 1 degree of freedom and non-centrality parameter $\eta_i = \tilde{\eta}_i^2$. If the original random vector $\mathbf{X}$ has zero mean ($\boldsymbol{\mu} = \mathbf{0}$), then $\eta_i = 0$, $\forall i$.

A q.f. can be labeled as *positive definite* if the eigenvalues $\boldsymbol{\lambda}$ are non-negative and as *indefinite* if they are both positive and negative. In the rest of the paper, Q indicates the complete expression $Q(\mathbf{X})$ for the sake of brevity.

3 Positive definite quadratic forms

In Ruben (1962), the following representation of $f_Q(\cdot)$ is proposed

$$f_Q(q) = \sum_{k=0}^{\infty} a_k \frac{q^{\alpha+k-1} \exp\left\{-\frac{q}{2\beta}\right\}}{(2\beta)^{\alpha+k}\Gamma(\alpha+k)} = \sum_{k=0}^{\infty} a_k f_G(q; \alpha+k, 2\beta), \tag{6}$$

where $f_G(q; \alpha+k, 2\beta)$ denotes the p.d.f. of a r.v. G following a gamma distribution with shape parameter $\alpha + k$ and scale 2β evaluated at q. The a_k coefficients are recursively defined as follows:

$$a_0 = \exp\left\{-\frac{\eta}{2}\right\} \prod_{i=1}^{r} \left(\frac{\beta}{\lambda_i}\right)^{\frac{1}{2}}; \qquad a_k = (2k)^{-1} \sum_{l=0}^{k-1} b_{k-l} a_l, \quad k \geq 1;$$

$$b_k = k\beta \sum_{i=1}^{r} \frac{\eta_i}{\lambda_i} c_i^{k-1} + \sum_{i=1}^{r} c_i^k, \quad k \geq 1,$$

where

$$\eta = \sum_{i=1}^{r} \eta_i, \qquad \alpha = \frac{r}{2}, \qquad c_i = 1 - \frac{\beta}{\lambda_i}.$$

To guarantee the absolute convergence of the series, the arbitrary constant β must fulfill the condition $\left|1 - \frac{\beta}{\lambda_i}\right| < 1,\ \forall i$. Ruben (1962) discusses how the choice of β affects the convergence speed of the series. For the following developments, it is worth noting that, if $0 < \beta \leq \lambda_r$, the coefficients a_k have the following properties:

$$a_k \geq 0,\ \forall k; \qquad \sum_{k=0}^{\infty} a_k = 1, \tag{7}$$

i.e. (6) turns out to be a mixture of gamma densities, whose Mellin transform is:

$$\widehat{f}_G(z; \alpha + k, 2\beta) = (2\beta)^{z-1} \frac{\Gamma(z + \alpha + k - 1)}{\Gamma(\alpha + k)}, \quad k \geq 0$$

within the strip of analyticity $\Re(z) > 1 - \alpha - k$. The Mellin transforms of $f_Q(\cdot)$ and $F_Q(\cdot)$ are derived in the following proposition.

Proposition 1. *Considering Ruben's expression for the p.d.f. of positive definite q.f.s (6), its Mellin transform is:*

$$\widehat{f}_Q(z) = (2\beta)^{z-1} \sum_{k=0}^{\infty} a_k \frac{\Gamma(z + \alpha + k - 1)}{\Gamma(\alpha + k)}, \quad \Re(z) > 1 - \alpha. \tag{8}$$

A computationally convenient recursive reformulation is

$$\widehat{f}_Q(z) = (2\beta)^{z-1} \sum_{k=0}^{\infty} a_k P_k(\alpha, z - 1),$$

where

$$P_k(\alpha, z-1) = \begin{cases} \frac{\Gamma(z+\alpha-1)}{\Gamma(\alpha)}, & k = 0; \\ P_{k-1}(\alpha, z-1)\left(1 + \frac{z-1}{\alpha+k+1}\right), & k > 0. \end{cases}$$

The Mellin transform of the c.d.f. is:

$$\widehat{F}_Q(z) = -\frac{(2\beta)^z}{z} \sum_{k=0}^{\infty} a_k \frac{\Gamma(z + \alpha + k)}{\Gamma(\alpha + k)}, \quad -\alpha < \Re(z) < 0.$$

Proof. See Supplementary material. □

We propose to compute $f_Q(\cdot)$ and $F_Q(\cdot)$ numerically inverting the Mellin transform, i.e. by numerical evaluation of the integrals:

$$f_Q(q) = \frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} q^{-z}\widehat{f}_Q(z)\mathrm{d}z,$$

and

$$F_Q(q) = \frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} q^{-(z-1)}\widehat{F}_Q(z-1)\mathrm{d}z = -\frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} q^{-(z-1)}\frac{\widehat{f}_Q(z)}{z-1}\mathrm{d}z,$$

where the last equality follows from equation (4).
Note that both the p.d.f. and the c.d.f. can be recovered from $\widehat{f}_Q(\cdot)$ provided that $1-\alpha < \Re(z) < 1$. Since $\widehat{f}_Q(\cdot)$ does not depend on q, once $\widehat{f}_Q(\cdot)$ is obtained, it can be used for computing the p.d.f., c.d.f., and quantile function of the q.f.. As a matter of fact, numerical inversion of these two integrals requires a small computational effort, the most intensive computiational step being the evaluation of the Mellin transform.

3.1 Numerical algorithm

The algorithm receives as input the vector of positive weights $\lambda_1, \dots, \lambda_r$, the non-centrality parameters $\eta_1, \dots, \eta_r$, the absolute error ε, and a probability level ρ. The latter input determines the range of q values for which it is required to keep the desired error: the algorithm is built to guarantee an absolute error lower than ε for every $f_Q(q)$ and $F_Q(q)$ such that $F_Q^{-1}((1-\rho)/2) < q < F_Q^{-1}((1+\rho)/2)$. Computation of $f_Q(\cdot)$, $F_Q(\cdot)$ and $F_Q^{-1}(\cdot)$ is possible via the functions `dQF`, `pQF`, and `qQF` that are available in the in the R package `QF`. All the R functions take as input the Mellin transform of the p.d.f. that can be computed by means of the function `mellin_QF`. For further details see the package documentation (Gardini et al., 2021). The quantile function `qQF` is based on a Newton-Raphson algorithm exploiting the `dQF` and `pQF` functions.

In order to provide a reliable evaluation of the target functions, the numerical error must be controlled. In particular, three different error sources can be individuated: the *Mellin transform truncation error* e_M, the *inversion integral truncation error* e_T, and the *discretization error* e_D. To isolate the different error sources annexed to the computed value $\tilde{f}_Q(q)$, the Mellin transform is decomposed as follows:

$$\begin{aligned} f_Q(q) &= \frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} q^{-z}\widehat{f}_Q(z)\mathrm{d}z \\ &= \frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} q^{-z}\,(2\beta)^{z-1}\sum_{k=0}^{\infty} a_k P_k(\alpha, z-1)\mathrm{d}z \\ &= \frac{1}{2\pi}\int_{-\infty}^{+\infty} q^{-(h+iy)}\,(2\beta)^{h+iy-1}\sum_{k=0}^{\infty} a_k P_k(\alpha, h+iy-1)\mathrm{d}y \\ &= \frac{\Delta\,(2\beta)^{-1}}{2\pi}\sum_{t=-T}^{T}\left(\frac{2\beta}{q}\right)^{h+i\Delta t}\sum_{k=0}^{K} a_k P_k(\alpha, h-1+i\Delta t) + e_M + e_T + e_D \\ &= \tilde{f}_Q(q) + e_M + e_T + e_D, \end{aligned} \tag{9}$$

where Δ is the numerical integration step size. Thus, the numerical error is

$$\left|f_Q(q) - \tilde{f}_Q(q)\right| = |e_M + e_T + e_D| = |\varepsilon|\,.$$

A similar decomposition and equivalent expression of the errors can be obtained for the c.d.f..

More in detail, the error term e_M arises from truncation at term K of the infinite sum (8), its absolute value is:

$$|e_M| = \left|\frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} q^{-z}\,(2\beta)^{z-1}\sum_{k=K+1}^{\infty} a_k P_k(\alpha, z-1)\mathrm{d}z\right|. \tag{10}$$

The error term e_T is due to truncation of the integration interval from $(h-i\infty; h+i\infty)$ to $(h-i\Delta T; h+i\Delta T)$:

$$\begin{aligned} |e_T| &= \left|\frac{1}{2\pi i}\int_{h-i\infty}^{h-iT\Delta} q^{-z}\widehat{f}_Q(z)\mathrm{d}z\right| + \left|\frac{1}{2\pi i}\int_{h+iT\Delta}^{h+i\infty} q^{-z}\widehat{f}_Q(z)\mathrm{d}z\right| \\ &= \left|\frac{1}{\pi i}\int_{h+iT\Delta}^{h+i\infty} q^{-z}\,(2\beta)^{z-1}\sum_{k=0}^{K} a_k P_k(\alpha, z-1)\mathrm{d}z\right|. \end{aligned}$$

Finally, the discretization error e_D arises from using a Riemann sum with constant integration step Δ for numerical integration.

Without loss of generality, scaled weights $\lambda_i^* = \lambda_i/\lambda_r$ are considered to enhance stability of the algorithm. In order to allow a mixture representation of the p.d.f., $\beta = \lambda_r^* = 1$ is fixed, i.e. the lowest scaled weight. This allows to obtain a bound for e_M.

Proposition 2 (Mellin truncation error bound)**.** *The error e_M can be bounded by*

$$|e_M| \leq \left(1 - \sum_{k=0}^{K} a_k\right) f_G\left(2(\alpha + K); \alpha + K + 1, 2\right)$$

for the p.d.f. and

$$|e_M| \leq \left(1 - \sum_{k=0}^{K} a_k\right) \tag{11}$$

for the c.d.f..

Proof. See Supplementary material. □

This error bound is employed to chose the value K, i.e. the number of a_k coefficients needed to ensure the required precision. A bound for the integral truncation error e_T is provided in the following proposition.

Proposition 3 (Integral truncation error bound)**.** *The error e_T can be bounded by:*

$$|e_T| \leq \left|\widehat{f}_Q(h + i\Delta T)\right| \frac{h^2 + (\Delta T)^2}{\pi h q^h} \left(\frac{\pi}{2} - \arctan\left(\frac{\Delta T}{h}\right)\right)$$

for the p.d.f., and

$$|e_T| \leq \left|\widehat{f}_Q(h + i\Delta T)\right| \frac{h^2 + (\Delta T)^2}{\pi (h-1)^2} q^{1-h} \left(1 - \frac{\Delta T}{\sqrt{(\Delta T)^2 + (h-1)^2}}\right)$$

for the c.d.f..

Proof. See Supplementary material. □

From the practical point of view, truncation value T is selected when both the right-hand sides of the rules derived in Proposition 3 are below the chosen absolute value ε.

With regard to the discretization error e_D, a bound could be sought by observing that the Mellin transform can be cast as a two-sided Laplace transform. For a review on these approaches see Abate and Whitt (1992). However, we found it more practical to implement the following iterative procedure to control e_D:

1. the Mellin transform is computed using an initial integration step Δ_I;

2. quantiles $q_L = F_Q^{-1}\left((1-\rho)/2\right)$ and $q_U = F_Q^{-1}\left((1+\rho)/2\right)$ are computed;

3. both the p.d.f. and the c.d.f. are evaluated at q_L and q_U, and these values are compared to those computed integrating the Mellin with a wider integration step $\Delta_I + \delta$, $\delta > 0$;

4. if the attained precision exceeds the required precision, i.e. if the maximum difference is lower than ε, the integration step is increased. Otherwise, the integration step is halved until the difference is below ε.

4 The ratio of quadratic forms

The distribution of the ratio of q.f.s has been widely studied since it involves the sampling distribution of several test statistics. In this paper, the basic properties of the Mellin transform are employed for studying such distributions. The interest lies in the following r.v.:

$$D = \frac{Q_1}{Q_2},$$

where Q_1 and Q_2 are independent positive q.f.s in the framework of Section 4.1, dependent q.f.s in Section 4.2.

4.1 The ratio of independent positive quadratic forms

In this section, the r.v. D previously defined is studied assuming Q_1 and Q_2 being independent, possibly non-central, q.f.s:

$$Q_1 = \sum_{i=1}^{r_1} \lambda_{i,1}(U_{i,1} + \tilde{\eta}_{i,1})^2, \qquad Q_2 = \sum_{i=1}^{r_2} \lambda_{i,2}(U_{i,2} + \tilde{\eta}_{i,2})^2.$$

Kim et al. (2006) obtained exact results for the distribution of the ratio between a non-central chi-square and a positive q.f., and they also provided an approximation for the distribution of the ratio between independent q.f.s. The exact result is expressed as an infinite weighted sum of Fisher's F p.d.f.s and c.d.f.s.

Starting from the property (2) of the Mellin transform, it is possible to derive the analogous result for the general case in which Q_1 is any positive definite q.f.. Recalling the expression for $\hat{f}_Q(z)$ (8), the Mellin transform of D can be written as:

$$\begin{aligned}\hat{f}_D(z) &= \hat{f}_{Q_1}(z)\hat{f}_{Q_2}(2-z) \\ &= \left(\frac{\beta_1}{\beta_2}\right)^{z-1}\left(\sum_{k=0}^{\infty} a_{k,1}\frac{\Gamma\left(\alpha_1+z+k-1\right)}{\Gamma\left(\alpha_1+k\right)}\right)\left(\sum_{j=0}^{\infty} a_{j,2}\frac{\Gamma\left(\alpha_2+j+1-z\right)}{\Gamma\left(\alpha_2+j\right)}\right),\end{aligned} \tag{12}$$

where the subscripts 1 and 2 denote the quantities related to Q_1 and Q_2 respectively and coefficients $a_{k,1}$ and $a_{j,2}$ are obtained via Ruben's recursion. Note that $\mathcal{S}\left(\widehat{f}_{Q_1}\right) = \{z : \Re(z) > 1-\alpha_1\}$ and $\mathcal{S}\left(\widehat{f}_{Q_2}\right) = \{z : \Re(2-z) > 1-\alpha_2\}$, hence $\mathcal{S}\left(\widehat{f}_D\right) = \{z : 1-\alpha_1 < \Re(z) < 1+\alpha_2\}$. Moreover, the Mellin transform of the c.d.f. can be recovered from $\widehat{f}_D(z)$ using (4), with strip of analyticity $1-\alpha_1 < h < 1$.

The analytical inversion of $\widehat{f}_D(\cdot)$ allows to express the p.d.f. (c.d.f.) of the ratio of independent q.f.s as a weighted sum of p.d.f.s (c.d.f.s) of 2^{nd} kind beta r.v.s, as stated in the following theorem.

Theorem 1 (Distribution of the ratio of independent q.f.s)**.** *Let D be a ratio of two independent positive definite q.f.s Q_1 and Q_2. The p.d.f. $f_D(\cdot)$ can be expressed as:*

$$f_D(d) = \sum_{k=0}^{+\infty}\sum_{j=0}^{+\infty} a_{k,1}a_{j,2} f_{B2}\left(d; \frac{\beta_1}{\beta_2}, \alpha_1+k, \alpha_2+j\right),$$

where $f_{B2}\left(x; b, p, q\right)$ indicates the p.d.f. of a beta r.v. of the 2^{nd} kind (B2) with parameters b, p, q:

$$f_{B2}(x; b, p, q) = \frac{b^{-p}}{B(p,q)} x^{p-1}\left(1+\frac{x}{b}\right)^{-p-q}. \tag{13}$$

and $B(p,q) = \Gamma(p)\Gamma(q)/\Gamma(p+q)$ is the beta function.

Analogously, the c.d.f. $F_D(\cdot)$ can be expressed as

$$F_D(d) = \sum_{k=0}^{+\infty}\sum_{j=0}^{+\infty} a_{k,1}a_{j,2} F_{B2}\left(d; \frac{\beta_1}{\beta_2}, \alpha_1+k, \alpha_2+j\right).$$

Proof. Starting from the analytical inversion of $\widehat{f}_D(\cdot)$, after arranging the terms and switching the integral and sums in virtue of the absolute convergence of the integral and the consequent application of Fubini's theorem, one has:

$$
\begin{aligned}
f_D(d) =& \frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty} d^{-z}\widehat{f}_D(z)\mathrm{d}z\\
=& \left(\frac{\beta_1}{\beta_2}\right)^{-1}\sum_{k=0}^{+\infty}\sum_{j=0}^{+\infty}\frac{a_{k,1}a_{j,2}}{\Gamma\left(\alpha_1+k\right)\Gamma\left(\alpha_2+j\right)}\times\\
&\times\frac{1}{2\pi i}\int_{h-i\infty}^{h+i\infty}\left(\frac{\beta_1}{d\beta_2}\right)^{z}\Gamma\left(\alpha_1+z+k-1\right)\Gamma\left(\alpha_2+j+1-z\right)\mathrm{d}z\\
=& \left(\frac{\beta_1}{\beta_2}\right)^{-1}\sum_{k=0}^{+\infty}\sum_{j=0}^{+\infty}\frac{a_{k,1}a_{j,2}}{B\left(\alpha_1+k,\alpha_2+j\right)}\left(\frac{\beta_1}{d\beta_2}\right)^{\alpha_2+j+1}\left(\frac{\beta_1}{d\beta_2}+1\right)^{-\alpha_1-k-\alpha_2-j}\\
=& \sum_{k=0}^{+\infty}\sum_{j=0}^{+\infty}a_{k,1}a_{j,2}\frac{\left(\beta_1\beta_2^{-1}\right)^{-\alpha_1-k}}{B\left(\alpha_1+k,\alpha_2+j\right)}d^{\alpha_1+k-1}\left(1+\frac{d}{\beta_1\beta_2^{-1}}\right)^{-\alpha_1-k-\alpha_2-j}\\
=& \sum_{k=0}^{+\infty}\sum_{j=0}^{+\infty}a_{k,1}a_{j,2}f_{B2}\left(d;\frac{\beta_1}{\beta_2},\alpha_1+k,\alpha_2+j\right),
\end{aligned}
$$

where the equality involving the integral is due to equation 5.13.1 in Olver et al. (2010). The expression of the c.d.f. can be recovered by integrating $f_D(d)$ switching the order of integration and summations. □

From a computational point of view, the double infinite sum characterizing the density derived in Theorem 1 is difficult to handle since $f_{B2}\left(d;\beta_1/\beta_2,\alpha_1+k,\alpha_2+j\right)$ depends on both k and j. For this reason, the algorithm proposed in the following section takes advantage of the Mellin transform expression (12), which is the product of two infinite sums.

4.1.1 Numerical algorithm

As in Section 3.1, the target quantities $f_D(d)$, $F_D(d)$ and $F_D^{-1}(p)$ are obtained by numerically inverting the Mellin transform, and the associated numerical error has three sources that must be controlled as before. In parallel to decomposition in equation (9), it is worth to remark that the actual computed quantity is:

$$
\begin{aligned}
\tilde{f}_D(d) =& \frac{\Delta}{2\pi}\left(\frac{\beta_2}{\beta_1}\right)\sum_{t=-T}^{T}\left(\frac{\beta_1}{d\beta_2}\right)^{h+i\Delta t}\left(\sum_{k=0}^{K}a_{k,1}P_k(\alpha_1,h-1+i\Delta t)\right)\times\\
&\left(\sum_{j=0}^{J}a_{j,2}P_j(\alpha_2,1-h-i\Delta t)\right).
\end{aligned}
$$

The proposed algorithm considers the scaled eigenvalues of the q.f. at the numerator $\lambda^*_{i,1} = \lambda_{i,1}/\lambda_{r,1}$ and it traces the algorithm presented for the computation of quantities related to the distribution of Q. The number of coefficients used to compute $\tilde{f}_D(d)$ are fixed using criterion (11) as stopping rule both at the numerator and at the denominator. The inequality (3) is used to control e_T. In this case, same expressions of the bounds derived in Proposition 3 are retrieved, plugging in $\hat{f}_D(z)$ instead of $\hat{f}_Q(z)$. Finally, the discretization error e_D is controlled using an iterative procedure, as before.

4.2 The ratio of dependent quadratic forms

In this section, the case in which D is defined in terms of dependent quadratic forms Q_1 and Q_2 is considered. As most of the approaches proposed in the literature, we focus on computation of $F_D(\cdot)$, while computation of $f_D(\cdot)$ is not addressed. To the best of our knowledge, the only attempt to compute $f_D(\cdot)$ is the one by Broda and Paolella (2009). Using the canonical notation (5), D can be expressed as:

$$D = \frac{Q_1}{Q_2} = \frac{\mathbf{Y}^T\mathbf{A}^\star\mathbf{Y}}{\mathbf{Y}^T\mathbf{B}^\star\mathbf{Y}};$$

where $\mathbf{Y} = \mathbf{Z} + \mathbf{L}^{-1}\boldsymbol{\mu}$. Matrices $\mathbf{A}^\star = \mathbf{L}^T\mathbf{A}\mathbf{L}$ and $\mathbf{B}^\star = \mathbf{L}^T\mathbf{B}\mathbf{L}$ are $(p \times p)$-dimensional semi-positive definite matrices. The most common approach to evaluate the c.d.f. at a generic quantile d is transforming the original problem to the evaluation of the c.d.f. of an indefinite q.f.:

$$F_D(d) = \mathbb{P}\left[Q_1/Q_2 \leq d\right] = \mathbb{P}\left[\mathbf{Y}^T\left(\mathbf{A}^\star - d\mathbf{B}^\star\right)\mathbf{Y} \leq 0\right].$$

Then, an algorithm like the one by Imhof (1961) or Davies (1980) is used to compute the value considering the eigenvalues of the matrix $\mathbf{A}^\star - d\mathbf{B}^\star$, i.e. $\boldsymbol{\lambda}(d)$.

We retrieve $F_D(\cdot)$ by leading back the problem to a ratio of independent q.f.s. and then applying the method outlined in Section 4.1. The indefinite q.f. $\mathbf{Y}^T\left(\mathbf{A}^\star - d\mathbf{B}^\star\right)\mathbf{Y}$ is considered as a linear combination of chi-square r.v.s with weights $\boldsymbol{\lambda}(d)$. As highlighted by the notation, the weights depend on d, i.e. the quantile where the c.d.f. is evaluated. Then, positive and negative terms are separated as follows:

$$\sum_{i=1}^{r}\lambda_i(d)(U_i+\tilde{\eta}_i)^2 = \sum_{i:\lambda_i(d)>0}\lambda_i(d)(U_i+\tilde{\eta}_i)^2 - \sum_{i:\lambda_i(d)<0}|\lambda_i(d)|(U_i+\tilde{\eta}_i)^2 = P(d) - N(d);$$

finding a difference of two independent positive definite q.f.s: $P(d)$ and $N(d)$. Hence, $F_D(d)$ can be evaluated as the c.d.f. of a ratio of independent positive definite q.f.s:

$$F_D(d) = \mathbb{P}\left[P(d) - N(d) \leq 0\right] = \mathbb{P}\left[\frac{P(d)}{N(d)} \leq 1\right]; \tag{14}$$

i.e., $F_D(d)$ is retrieved by evaluating the c.d.f. of the r.v. $P(d)/N(d)$ at 1. Clearly, the dependence of the weights on d causes a considerable computational burden, since the Mellin transform needs to be recomputed for each d: this problem is common to all the approaches proposed in the literature. The function `pQF_depratio` is available in the `QF` package for computing the c.d.f. of dependent q.f.s.

5 Numerical evaluations and applications

In Section 5.1, our routine is compared to other popular procedures with respect to both computational time and accuracy. Then, three applications are discussed: Section 5.2 deals with the sample variance distribution of a Gaussian random vector with an autoregressive (AR) covariance matrix, corresponding to the distribution of a positive q.f.. Section 5.3 deals with the computation of size and power of the Behrens-Fisher test statistic, for testing equality of means of two Gaussian populations, under departure from the homoschedasticity hypothesis: this requires dealing with the ratio of independent q.f.s. Finally, Section 5.4 deals with computation of the c.d.f. of the Cronbach's alpha statistic that can be cast as the ratio of dependent q.f.s. The last two examples have also been discussed in Kim et al. (2006).

5.1 Testing the algorithms

The proposed computational algorithm is compared to some popular existing procedures such as algorithms by Davies (1980), Imhof (1961) and Farebrother (1984): all these algorithms are designed to evaluate the c.d.f. of the considered r.v.. Only the latter provides the p.d.f., but cannot be used when the focus is on the ratio of independent q.f.s. The `R` package `CompQuadForm` (De Micheaux, 2017) provides the implementation of these algorithms. In Table 1, the chosen q.f.s and ratios are listed. Some of them are taken from Imhof (1961)

R.v.s	(λ_i, η_i, multiplicity)	R.v.s	(λ_i, η_i, multiplicity)
Q_1	$(6,0,1);\ (3,0,1);\ (1,0,1)$	Q_6	$(30,0,1);\ (1,0,30)$
Q_2	$(6,0,2);\ (3,0,2);\ (1,0,2)$	Q_7	$(i,0,1);\quad i=1,\dots,10$
Q_3	$(6,0,6);\ (3,0,4);\ (1,0,2)$	D_1	Q_1/Q_7
Q_4	$(7,6,6);\ (3,2,2)$	D_2	Q_7/Q_3
Q_5	$(7,6,1);\ (3,2,1)$	D_3	Q_4/Q_1

Table 1: Positive definite q.f.s ($Q_l,\ l = 1,\dots,7$) and ratios ($D_j,\ j = 1,2,3$) studied in Section 5.1.

and they are defined in terms of weights (λ_i), multiplicities, and non-centrality parameters (η_i). An absolute precision level $\varepsilon = 10^{-6}$ is adopted. For our procedure, $\rho = 0.9999$ is fixed: this guarantees the precision ε between quantiles $(1-\rho)/2$ and $(1+\rho)/2$.

The computational times for repeated evaluation ($n = 7$ and $n = 1000$ equally spaced quantiles) of the c.d.f. and p.d.f. have been studied. Note that results from Imhof and Davies algorithms refer to c.d.f. computation only. The complete results consist of the average time obtained with 100 replications on a PC with processor Intel Core i7-9750H (2.60GHz), and are available in Table S3 in the Supplementary material. This analysis allows to highlight that our proposal is competitive in terms of computational times, and a substantial time saving can be noted in case of repeated evaluations ($n = 1000$) since the most computationally expensive step of the algorithm concerns the Mellin transform computation, but this task is accomplished only one time.

Focusing on positive definite q.f.s, the accuracy in computing the c.d.f. at different quantiles is compared. The quantiles related to the following probabilities are chosen: $p = (5\times10^{-6}, 0.25, 0.5, 0.75, 1-5\times10^{-6})$, noting that the first and the last values correspond to quantiles laying outside the range of values (determined by ρ) for which the desired absolute error ε is guaranteed by our algorithm. To measure the accuracy, the absolute relative difference between the value computed by each algorithm (with $\varepsilon = 10^{-6}$) and the benchmark (Farebrother's algorithm with $\varepsilon = 10^{-26}$) is computed. All the registered values are below 1 (further details in Table S4 in the supplementary material), hence the requested accuracy has been achieved in every considered case. Focusing on our algorithm, we remark that this result occurs also for quantiles outside the range of values for which the error level is guaranteed. This finding points out the reliability of the proposed algorithm.

5.2 Sample variance distribution under AR model

In this section, we consider the distribution of the sample variance of an n-dimensional Gaussian random vector $\mathbf{Y} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{\Sigma}(\phi))$, where $\mathbf{\Sigma}(\phi)$ determines an autoregressive process of order 1 with standard normal errors, i.e. the ij-th entry is:

$$[\mathbf{\Sigma}(\phi)]_{ij} = \frac{1}{1-\phi^2}\phi^{|i-j|} \qquad i, j = 1, \ldots, n.$$

The sample variance $V|\phi = n^{-1}\mathbf{Y}^T\mathbf{M}\mathbf{Y}$, where $\mathbf{M}$ is the centering matrix, is a positive definite q.f.. Defining $\mathbf{L}(\phi)$ such that $\mathbf{\Sigma}(\phi) = \mathbf{L}(\phi)^T\mathbf{L}(\phi)$, $V|\phi$ can be expressed as the weighted sum of independent central χ^2 r.v.s with weights $\boldsymbol{\lambda}(\phi)$ corresponding to the eigenvalues of $\mathbf{L}(\phi)^T\mathbf{M}\mathbf{L}(\phi)$. In what follows, we report results obtained for computing $f_Q(\cdot)$ and $F_Q(\cdot)$ fixing $\rho = 0.9999$ and $\varepsilon = 10^{-6}$.

A quantity of major interest for understanding the merit of our approach is the condition number of $\mathbf{L}(\phi)^T\mathbf{M}\mathbf{L}(\phi)$, $\kappa(\phi) = \lambda_1(\phi)/\lambda_r(\phi)$, which is a measure of the skewness of the spectrum and is an increasing function of ϕ, as can be check from Table 2. The same table reports the number of terms required by Farebrother's algorithm (denoted as K) and the number of integration points for numerical inversion of the Mellin transform (denoted as T).

Note that the number of terms required by the Farebrother's algorithm dramatically increases with ϕ: this is expected when using Ruben's formula since the a_k coefficients show slow decay when the skewness of the weights increases. On the other hand, the Mellin transform method is able to store the information required to compute the target quantities using a small number of integration points and the number of terms required for numerical integration decreases with the skewness of the weights. This is coherent with the typical behavior of integral transforms, that *stretching in the time domain* turns out in *squeezing in the frequency domain.* An example of the decay rate of the imaginary part of the Mellin transform for two values of ϕ is shown in the left-panel of Figure 1, while the right-panel

	ϕ	0.00	0.25	0.50	0.75	0.95	0.975
	$\kappa(\phi)$	1.00	2.76	8.74	42.57	667.73	2068.13
	K	1	57	162	563	6159	17838
	T	132	126	80	42	34	34
$F_Q^{-1}(p)$	$p = 0.025$	0.6311	0.6488	0.7268	0.9369	1.6401	2.0116
	$p = 0.5$	0.9667	1.0143	1.2212	1.8870	6.0711	10.8916
	$p = 0.975$	1.4044	1.5129	1.9900	3.8046	23.6796	59.7212

Table 2: Condition number $\kappa(\phi)$, number of a_k coefficients for Farebrother's algorithm (K), number of Mellin transform integration points T, and distribution quantiles $F_Q^{-1}(p)$ for $V|\phi$, with $\phi = 0, 0.25, 0.5, 0.75, 0.95, 0.975$.

shows the imaginary part of the Fourier transform $\bar{f}_Q(y)$. In both cases, non-negligible values of both the integral transforms referred to $\phi = 0.975$ are contained in a shorter range, reflecting the aforementioned *squeezing in the frequency domain* behaviour. Moreover, when ϕ increases, the Fourier transform becomes less smooth, hence more difficult to numerically integrate: this generates the well-known pathological behaviors of the Imhof and Davies algorithms. To evaluate the Mellin transform at integration points $h + i\Delta t,\ t = -T, \ldots, T$, computation of the K coefficients of Ruben's expansion is required. However, once these values are available, they can be used for computing any probabilistic quantity related to the q.f. distribution. In synthesis, the Mellin transform is able to conveniently store all the relevant features of the q.f. in few terms and turns out to be particularly useful when several evaluations of $f_Q(\cdot)$ or $F_Q(\cdot)$ are required, as pointed out in Section 5.1. Moreover, the computational saving is remarkable when the skewness of the weights increases.

In Figure 2, the p.d.f.s of $V|\phi$ are shown for $\phi = 0$ and $\phi = 0.975$: they are computed through our algorithm in the range containing a probability mass equal to ρ. Referring to the `QF` package, the function `compute_MellinQF` can be used to evaluate the Mellin transform. An object of class `MellinQF` is produced and it serves as input of functions `dQF` and `pQF`. The plots confirm that the higher skewness of weights in the case $\phi = 0.975$ induces a probability distribution with a marked positive asymmetry. The function `qQF`

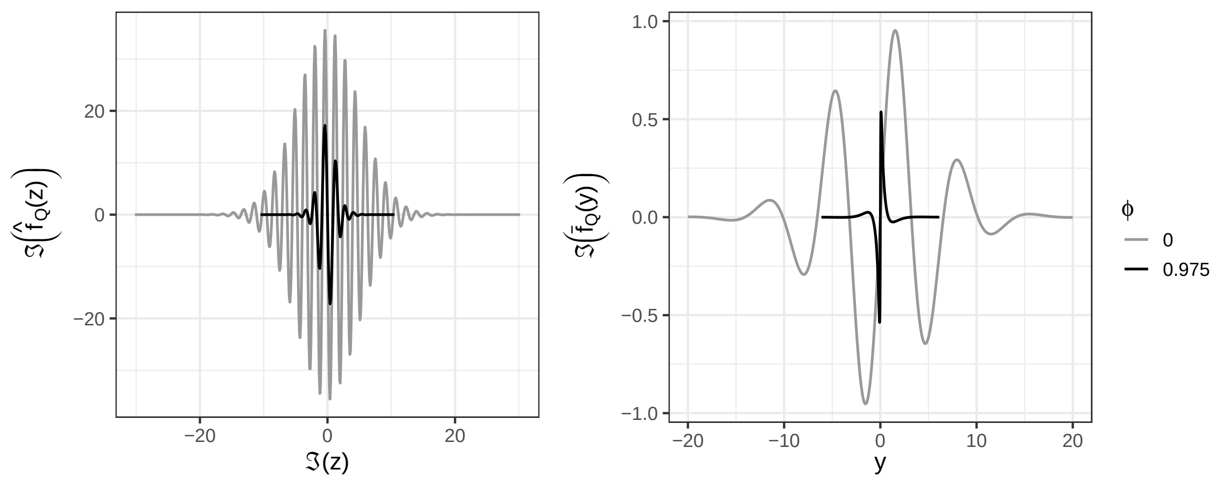

Figure 1: Imaginary part of $\hat{f}_Q(z)$ and $\bar{f}_Q(y)$, with $Q = V|\phi$ and $\phi \in \{0, 0.975\}$.

is an appealing tool provided in the package: it receives as input a vector of probabilities and the `MellinQF` object, delivering the quantiles of the q.f. through a Newton-Raphson algorithm, requiring repeated evaluation of the c.d.f and the p.d.f.. In the second part of Table 2, the outcomes related to quantiles of $V|\phi$ at different values of parameter ϕ are reported.

5.3 Behrens-Fisher test statistic

Consider two Gaussian r.v.s $Y_k \sim N(\mu_k, \sigma_k^2)$, $k = 1, 2$ and the hypothesis system $H_0 : \mu_1 = \mu_2$ vs $H_1 : \mu_1 \neq \mu_2$. In this Section, we study power and size of the Behrens-Fisher test statistic for different values of the ratio σ_1^2/σ_2^2 and different sample sizes n_1 and n_2. The Behrens-Fisher test statistic is

$$t = \frac{\bar{Y}_1 - \bar{Y}_2 - \delta}{\sqrt{s_1^2/n_1 - s_2^2/n_2}}$$

where $\bar{Y}_k$, s_k^2 and n_k are the sample mean, variance and size respectively, $k = 1, 2$, and $\delta = \mu_1 - \mu_2$. Kim et al. (2006) have shown that, for fixed values of the ratio σ_1^2/σ_2^2:

$$t^2 = \frac{\lambda_0 X_0}{\lambda_1 X_1 + \lambda_2 X_2} \tag{15}$$

where $\lambda_0 = \sigma_1^2/n_1 + \sigma_2^2/n_2$ and $\lambda_k = \sigma_k^2 n_1 n_2 [n_1 n_2 (n_1 + n_2 - 2)]$, $k = 1, 2$, while X_0, X_1 and X_2 are independently distributed as chi-squared r.v.s, namely $X_0/\lambda_0 \sim \chi_1^2(\delta^2)$ and

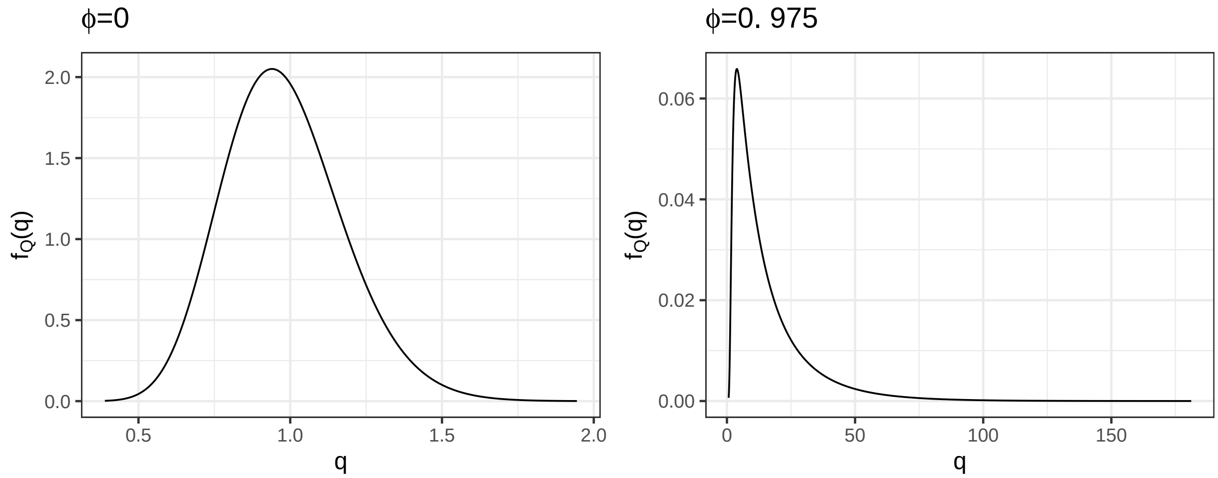

Figure 2: Plot of $f_Q(q)$ for $\phi = 0.05$ and $\phi = 0.975$, with $Q = V|\phi$.

$X_k \sim \chi^2_{n_k-1}$, $k = 1, 2$. Hence, equation (15) corresponds to the ratio of independent q.f.s that can be managed with the algorithm outlined in Section 4.1.

Table S5 in the Supplementary material reproduces the results reported by Kim et al. (2006), comparing the proposed algorithm with Imhof's procedure: numerical Mellin inversion confirms its accuracy for computations involving q.f.s. Figure 3, left panel, shows three power functions with fixed sample sizes $n_1 = 6$, $n_2 = 51$, $\sigma_1^2 = 1$, $\delta \in [0; 4]$ and ratios σ_2^2/σ_1^2 equal to 1, 0.1 and 10. When the ratio is 1, the test is correctly specified and the size of the test is equal to the nominal significance level $\alpha = 0.05$. On the other hand, the size is lower than the nominal level for $\sigma_2^2/\sigma_1^2 = 10$ and higher than the nominal level for $\sigma_2^2/\sigma_1^2 = 0.1$; this ordering characterizes the whole power functions.
In the right panel, we report the Mellin transforms referred to three values of δ, fixing $\sigma_2^2/\sigma_1^2 = 1$. Even if the algorithm can be used to compute the power function with the required accuracy, it is worth noting that this application represents the worst case scenario for methods based on the inversion of the Mellin transform. This is due to divergence at 0 of the density of the χ_1^2 r.v. at the numerator of (15), implying a Mellin transform peaked at $\Im(z) = 0$. As a consequence, numerical inversion requires a dense integration grid. This behaviour is mitigated when the χ^2 non-centrality parameter δ increases.

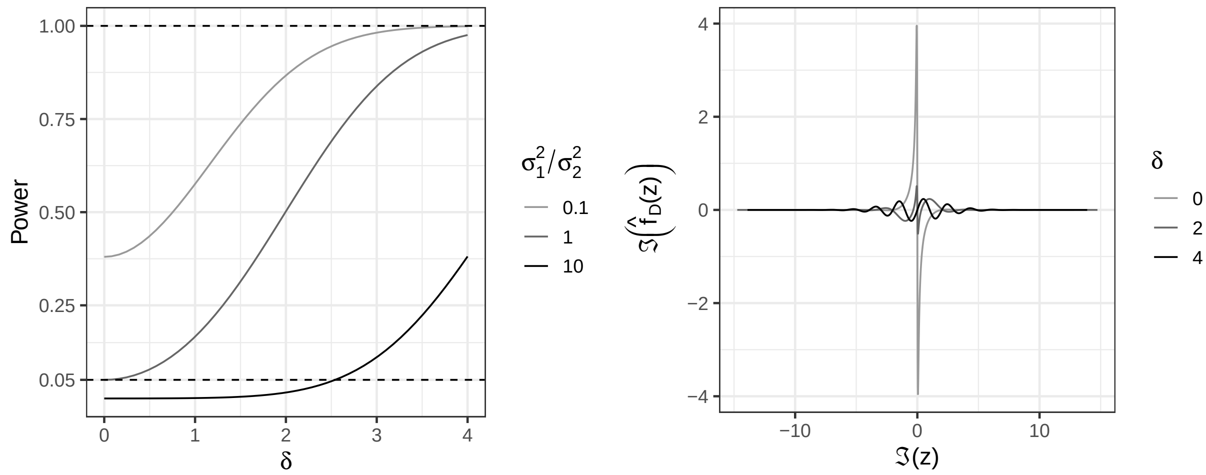

Figure 3: Left panel: power curves of the Behrens-Fisher test. Right-panel: Imaginary part of $\hat{f}_D(z)$ under different values of δ.

5.4 Cronbach's alpha distribution function

In this application, the objective is retrieving the c.d.f. of the Cronbach's alpha statistic (ρ_α), considering the case of N groups each with n observations, for which independent normal distributions $\mathbf{Y}_i \sim \mathcal{N}_n(\mathbf{0}, \mathbf{\Sigma}(\phi))$, $i = 1, \ldots, N$ are assumed, and $\mathbf{\Sigma}(\phi)$ is the covariance matrix of an AR(1) process. We are in the same framework of Kim et al. (2006), and the following probability needs to be evaluated:

$$\mathbb{P}[\rho_\alpha \leq r_\alpha] = \mathbb{P}\left[\sum_{k=1}^{n} \lambda_k(r_\alpha, \phi) W_k \leq 0\right],$$

where $W_k \sim \chi^2_{N-1}$ and $[\lambda_1(r_\alpha, \phi), \ldots, \lambda_n(r_\alpha, \phi)]$ are the eigenvalues of the matrix

$$\mathbf{L}(\phi)^T \left[(n/(n-1) - r_\alpha)\,\mathbf{1}_n\mathbf{1}_n^T - n/(n-1)\mathbf{I}_n\right] \mathbf{L}(\phi).$$

Since it is known that $n-1$ of them are negative and one is positive, an indefinite q.f. is faced (Kistner and Muller, 2004). Our algorithm can be applied leading the problem back to a ratio of independent q.f.s through expression (14).

In the left-panel of Figure 4, the c.d.f. of ρ_α is reported for the case $n = 5$, $N = 10$, and $\phi = 0.5$. To document the efficiency of our algorithm in these situations, we stress that the plot is produced using 1000 points: the requested computational time averaged on

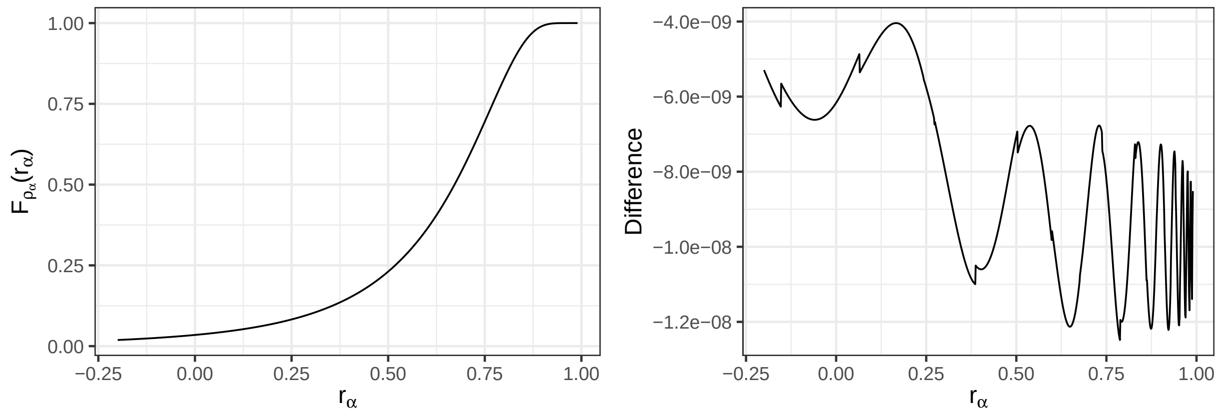

Figure 4: In the left-panel, the c.d.f. of Cronbach's alpha statistic is reported. In the right-panel, the difference between values obtained using our algorithm and Imhof's is displayed.

100 replications is $0.674s$ for the Mellin inversion based algorithm and $0.615s$ for Imhof's procedure. Moreover, in the right-panel, the difference between the computed values under the two procedures is displayed: the absolute value of this difference is far below the required precision $\varepsilon = 10^{-6}$.

6 Concluding remarks

The Mellin transform has been shown to be a useful tool for managing the distribution of q.f.s both from the computational and analytical point of view. From the computational perspective, the proposed algorithms start from the series representation by Ruben (1962) and store the information about the q.f. in a more compact way through the Mellin transform: the convenience of this summarization increases with the skewness of the q.f. weights. In Section 5.2, we show that the number of coefficients of Ruben's expansion needed to guarantee a given precision can be far higher than the integration points required for the Mellin inversion. In other words, the proposed method is particularly convenient when the q.f. is dominated by few weights: in this case, both Davies and Imhof algorithms show some pitfalls, as discussed in Kim et al. (2006). This is expected from methods relying on numerical inversion of the Fourier transform since such transform can show pathological behaviors in terms of smoothness, being less prone to numerical integration. Since the evaluation of the Mellin transform requires computation of the Ruben's coefficients, the proposed algorithms generate appreciable savings in computational time when repeated evaluations are needed: in this case, time-saving can be considerable. However, in Section 5, the algorithms have been shown to be competitive with other approaches in terms of computational time: the approach that outperforms the others depends on the q.f. structure.

From the theoretical point of view, the Mellin transform is particularly suitable for studying the distribution of products and ratios of r.v.s, as pointed out in Epstein (1948): in this paper, the properties of the Mellin transform have been exploited to deliver a representation of the ratio of independent q.f.s as a mixture of beta distributions of the second kind. Analytical treatment of the Mellin transform involves Mellin-Barns integrals that define Meijer's G functions: the theory in this field is rich and well-developed and can be fruitful in providing further insights into the study of estimators or test statistics involving q.f.s.

SUPPLEMENTARY MATERIAL

Supplementary pdf file: it contains some examples of statistical tests that involves q.f.s or ratios of q.f.s, the proofs of the propositions, and additional tables about Section 5.

R-code: a zipped folder containing the code to reproduce the results discussed in the paper.

References

Abate, J. and W. Whitt (1992). The Fourier-series method for inverting transforms of probability distributions. *Queueing systems 10*(1-2), 5–87.

Anderson, T. W. (1948). On the theory of testing serial correlation. *Scandinavian Actuarial Journal 1948*(3-4), 88–116.

Box, G. E. (1954). Some theorems on quadratic forms applied in the study of analysis of variance problems, I. Effect of inequality of variance in the one-way classification. *The Annals of Mathematical Statistics 25*(2), 290–302.

Broda, S. and M. S. Paolella (2009). Evaluating the density of ratios of noncentral quadratic forms in normal variables. *Computational Statistics & Data Analysis 53*(4), 1264–1270.

Davies, R. B. (1980). Algorithm AS 155: The distribution of a linear combination of χ 2 random variables. *Journal of the Royal Statistical Society. Series C (Applied Statistics) 29*(3), 323–333.

De Micheaux, P. L. (2017). CompQuadForm: distribution function of quadratic forms in normal variables. *R package version 1*(3).

Eddelbuettel, D. and R. François (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software 40*(8), 1–18.

Epstein, B. (1948). Some applications of the Mellin transform in statistics. *The Annals of Mathematical Statistics 19*(3), 370–379.

Farebrother, R. (1984). Algorithm AS 204: the distribution of a positive linear combination of χ 2 random variables. *Journal of the Royal Statistical Society. Series C (Applied Statistics) 33*(3), 332–339.

Gardini, A., F. Greco, and C. Trivisano (2021). *QF: Density, Cumulative and Quantile Functions of Quadratic Forms.* R package version 0.0.6.

Imhof, J.-P. (1961). Computing the distribution of quadratic forms in normal variables. *Biometrika 48*(3/4), 419–426.

Kim, H.-Y., M. J. Gribbin, K. E. Muller, and D. J. Taylor (2006). Analytic, computational, and approximate forms for ratios of noncentral and central gaussian quadratic forms. *Journal of Computational and Graphical Statistics 15*(2), 443–459.

Kistner, E. O. and K. E. Muller (2004). Exact distributions of intraclass correlation and cronbach's alpha with gaussian data and general covariance. *Psychometrika 69*(3), 459–474.

Kuonen, D. (1999). Miscellanea. Saddlepoint approximations for distributions of quadratic forms in normal variables. *Biometrika 86*(4), 929–935.

Lieberman, O. (1994). Saddlepoint approximation for the distribution of a ratio of quadratic forms in normal variables. *Journal of the American Statistical Association 89*(427), 924–928.

Olver, F. W., D. W. Lozier, R. F. Boisvert, and C. W. Clark (2010). *NIST handbook of mathematical functions hardback and CD-ROM*. Cambridge university press.

Paris, R. B. and D. Kaminski (2001). *Asymptotics and mellin-barnes integrals*, Volume 85. Cambridge University Press.

Poularikas, A. D. (2018). *Transforms and applications handbook*. CRC press.

Provost, S. B. (1989). On sums of independent gamma random variables. *Communications in Statistics-Theory and Methods 20*(4), 583–591.

Provost, S. B. and A. Mathai (1992). *Quadratic forms in random variables: theory and applications*. M. Dekker.

Provost, S. B. and E. M. Rudiuk (1994). The exact density function of the ratio of two dependent linear combinations of Chi-square variables. *Annals of the Institute of Statistical Mathematics 46*(3), 557–571.

R Core Team (2020). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Robbins, H. and E. Pitman (1949). Application of the method of mixtures to quadratic forms in normal variates. *The Annals of Mathematical Statistics*, 552–560.

Ruben, H. (1962). Probability content of regions under spherical normal distributions, IV: The distribution of homogeneous and non-homogeneous quadratic functions of normal variables. *The Annals of Mathematical Statistics 33*(2), 542–570.