



HAL
open science

Estimation of multivariate generalized gamma convolutions through Laguerre expansions

Oskar Laverny, Esterina Masiello, Véronique Maume-Deschamps, Didier Rullière

► **To cite this version:**

Oskar Laverny, Esterina Masiello, Véronique Maume-Deschamps, Didier Rullière. Estimation of multivariate generalized gamma convolutions through Laguerre expansions. *Electronic Journal of Statistics*, 2021, 15 (2), pp.5158-5202. <10.1214/21-EJS1918>. <hal-03160289v2>

HAL Id: hal-03160289

<https://hal.science/hal-03160289v2>

Submitted on 9 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Estimation of multivariate generalized gamma convolutions through Laguerre expansions.*

Oskar Laverny

*Institut Camille Jordan, UMR 5208,
Université Claude Bernard Lyon 1, Lyon, France
SCOR SE
e-mail: oskar.laverny@univ-lyon1.fr*

Esterina Masiello and Véronique Maume-Deschamps

*Institut Camille Jordan, UMR 5208,
Université Claude Bernard Lyon 1, Lyon, France
e-mail: esterina.masiello@univ-lyon1.fr; veronique.maume@univ-lyon1.fr*

Didier Rullière

*Mines Saint-Étienne,
UMR CNRS 6158, LIMOS, Saint-Étienne, France
e-mail: didier.rulliere@emse.fr*

Abstract: The generalized gamma convolutions class of distributions appeared in Thorin's work while looking for the infinite divisibility of the log-Normal and Pareto distributions. Although these distributions have been extensively studied in the univariate case, the multivariate case and the dependence structures that can arise from it have received little interest in the literature. Furthermore, only one projection procedure for the univariate case was recently constructed, and no estimation procedures are available. By expanding the densities of multivariate generalized gamma convolutions into a tensorized Laguerre basis, we bridge the gap and provide performant estimation procedures for both the univariate and multivariate cases. We provide some insights about performance of these procedures, and a convergent series for the density of multivariate gamma convolutions, which is shown to be more stable than Moschopoulos's and Mathai's univariate series. We furthermore discuss some examples.

MSC2020 subject classifications: Primary 62H12, 60E07; secondary 60E10.

Keywords and phrases: Multivariate Generalized Gamma Convolutions, Estimation, Thorin's measure, Laguerre's basis, Infinite Divisibility.

*This is an original paper

Contents

1 Introduction 2

2 Background 4

2.1 Gamma convolutions classes 5

2.2 Shifted moments and shifted cumulants 10

2.3 Projection from \mathcal{G}_1 to $\mathcal{G}_{1,n}$ 12

3 Construction of an estimator in $\mathcal{G}_{d,n}$ 14

3.1 The tensorized Laguerre basis 15

3.2 An incomplete but however consistent empirical loss . 24

4 Investigation of performance 27

4.1 Projection from a known density 27

4.1.1 Projection from \mathcal{G}_1 densities to $\mathcal{G}_{1,n}$ 28

4.1.2 Projection from densities outside \mathcal{G}_1 to $\mathcal{G}_{1,n}$. . 30

4.2 Estimation from univariate empirical data 31

4.2.1 Heavy-tailed examples 31

4.2.2 An example outside the class 35

4.3 Estimation from multivariate empirical data 37

4.3.1 A Gaussian example 37

4.3.2 A copula with tail dependency 39

5 Conclusion 40

References 41

1. Introduction

Olof Thorin introduced the Generalized Gamma Convolutions (GGC) class of distributions in 1977 as a tool to study the infinite divisibility of Pareto [54] and log-Normal [55] distributions. The quite natural and pragmatic question of identifying if some distributions are infinitely divisible or not turned out to be theoretically fruitful, and notably gave rise to the concept of hyperbolic complete monotone functions. Later, L. Bondesson extended these concepts in a book [5] which is still today a good reference on the subject. An introduction to more recent literature about GGC can be found in the two surveys [24, 7], on the probabilistic side of the problem.

This class contains a lot of known and commonly used distributions, including heavy-tailed distributions such as Pareto, log-Normal and α -stable distributions. This versatility makes the class appealing for the statistical practitioners working in various fields such as life and non-life insurance, reinsurance, anomaly or fraud detection, floods analysis and meteorology, etc. Other classes of general approximating distributions with high-order features can be considered. Mixtures of Erlangs [11, 30, 56], for example, have the property of being dense in L_2 [56] and allow for fast k -MLE estimation algorithms [44]. On the other hand, they lack crucial closure properties and interpretability that the generalized gamma convolutions have.

Surprisingly, on the statistical side, very little work can be found about estimation of generalized gamma convolutions. Only one projection procedure, that is a procedure that projects generalized gamma convolutions onto convolutions of a finite number of gammas, was published recently in [38, 18]. Although the resulting convergence result is stunning, this procedure cannot handle cases where the incoming density is not already inside the class. This typically happens for any empirical distribution, due to sampling noise, even if the sampling density is in the class. Furthermore, we show that the theoretical background of the estimation procedure in [38, 18] is inherently univariate, and no direct extension to a multivariate case is possible.

The multivariate analogue is still an active research field, and, as far as we know, has never been considered as a statistical tool in the literature. In applications, however, statisticians might deal with marginals that are in this class. A way to handle dependence structure in this interpretable framework seems therefore appealing as it would allow studying functions of the random vector, like the sum of components, under dependence assumptions. Therefore, an estimation procedure for multivariate generalized gamma convolutions that takes into account the dependence structure could be useful. The probabilistic background is already developed: based on an original idea from [8], Bondesson introduces a bivariate extension of the Thorin class, and gives some properties about these distributions. Later, [6] extends this concept to the multivariate case. More recently, [45, 46] propose another equivalent framework allowing extension on any cone, with application to matrix gamma convolutions. On the other hand, no estimation procedure nor explicit density exist. We discuss this multivariate class and show that the produced dependence structures are highly flexible: non-exchangeability and tail dependence are possible.

Regarding tools, the literature gives two series expansions for the density of univariate gamma convolutions: Mathai [33] proposes a density based on Kummer's confluent geometric functions and Moschopoulos [40] refined it as a convergent gamma series, with an explicit truncation error. Both these densities have the problem of being based on the smallest scale parameter, and are not well conditioned when the smallest scale is too small. Unfortunately, this is typically the case when approximating a log-Normal or a Pareto distribution by a finite convolution of gammas. Therefore, no stable procedure for the density computation is available for the entirety of the parameter range. We are not aware of any density estimation available in the literature for the multivariate case.

We consider here the problem of estimation from samples of multivariate generalized gamma convolutions. By the introduction of a specific Laguerre basis, which was already considered for non-parametric deconvolution problems by [16, 32, 2], we manage to obtain a convergent series for the density. Classical deconvolution problems usually consider only one source of signal and one source of noise [34, 36, 58, 52]. Although we have here a finite number of signals to be estimated, the Laguerre expansion is still a useful tool. We show that this expansion is quite natural for our problem, as it includes and generalizes

Moschopoulos’s univariate series. Using this basis, we bridge the gaps and provide a new stable algorithm for density computation in both the univariate and multivariate cases, as well as estimation procedures which handle both clean data (the density is given as a formal function) and dirty data (such as an empirical dataset in 64 bits precision). The resulting algorithms are implemented in the Julia package `ThorinDistributions.jl`, available on GitHub¹ under a MIT License and archived in Zenodo [29].

After fixing notations, Section 2 covers some definitions, properties, and algorithms to set up the stage. Section 3 considers the Laguerre expansion of densities in the multivariate Thorin class, provides some error bounds and develops an estimation procedure for these distributions. In Section 4, we investigate the numerical results of our estimation procedure on several examples. Section 5 concludes and gives leads for further work.

2. Background

As we deal with a lot of multivariate objects, we start by fixing some notations.

We use bold letters such as \mathbf{a} to designate generally indexable objects, e.g., vectors or matrices, and corresponding indexed and possibly unbolded letters, such as a_i , designate values in these objects. We use Cartesian indexing. For example, if we consider a row-major matrix \mathbf{r} , then \mathbf{r}_i denotes the i^{th} row of that matrix, $r_{i,j}$ the j^{th} value of that row and if $\mathbf{k} = (i, j)$, then $r_{\mathbf{k}} = r_{i,j}$.

We denote $|\mathbf{x}| = x_1 + \dots + x_d$ the sum of components of a vector \mathbf{x} . The product of factorials of component of an integer vector \mathbf{k} is denoted $\mathbf{k}! = k_1! \dots k_d!$, and we set $\mathbf{x}^{\mathbf{k}} = x_1^{k_1} \dots x_d^{k_d}$. We denote derivatives of a multivariate function f by $f^{(\mathbf{k})}(\mathbf{x}) = \frac{\partial}{\partial^{k_1} x_1} \dots \frac{\partial}{\partial^{k_d} x_d} f(\mathbf{x})$ and scalar products of vectors by $\langle \mathbf{x}, \mathbf{y} \rangle = x_1 y_1 + \dots + x_d y_d$. Finally, we denote $\binom{\mathbf{x}}{\mathbf{y}} = \prod_i \binom{x_i}{y_i}$ the product of binomial coefficients.

Inequalities, sums, and products between vectors of same shape are always intended componentwise, and standard broadcasting between objects of different shapes applies. For example, $1 + 2\mathbf{x}$, $\mathbf{1} + 2\mathbf{x}$ and $(1 + 2x)_{x \in \mathbf{x}}$ all denote the same object, which has same shape as \mathbf{x} .

To characterize the distribution of random vectors, we might use several functions. For a (real-valued) random vector \mathbf{X} of dimension $d \geq 1$, the cumulative distribution function (cdf) of \mathbf{X} is $F(\mathbf{x}) = \mathbb{P}(\mathbf{X} \leq \mathbf{x})$. When they exist, the probability density function (pdf) of \mathbf{X} , $f(\mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} F(\mathbf{x})$, its Laplace transform (lt) $L(\mathbf{t}) = \int e^{-\mathbf{t}\mathbf{x}} f(\mathbf{x}) d\mathbf{x}$, its moment generating function (mgf) $M(\mathbf{t}) = L(-\mathbf{t})$ or its cumulant generating function (cgf) $K(\mathbf{t}) = \ln(M(\mathbf{t}))$ can also be used to characterize the distribution. All these functions might be indexed by the random object to avoid ambiguity if needed.

¹`ThorinDistributions.jl`: <https://github.com/lrnv/ThorinDistributions.jl>.

We denote by $L_2(\mathbb{R}_+^d)$ the set of functions that are Lebesgue square-integrable onto \mathbb{R}_+^d , and by $L_2(\mathbb{R}_+, w)$ the set of square-integrable functions with respect to a weight function w .

The following subsections set the stage of our analysis by reviewing common tools from the literature.

2.1. Gamma convolutions classes

A random variable X is said to be Gamma distributed with shape $\alpha \in \mathbb{R}_+^*$, scale $s \in \mathbb{R}_+^*$ and rate $\frac{1}{s}$, which we denote $X \sim \mathcal{G}_{1,1}(\alpha, s)$, if it admits the cumulant generating function (cgf):

$$K(t) = \ln \mathbb{E} (e^{tX}) = -\alpha \ln (1 - st).$$

We denote by $\mathcal{G}_{1,1}$ this class of distributions. Recall that the cumulant generating function of the sum of independent random variables is the sum of their cumulant generating functions. Therefore, gamma distributions are clearly infinitely divisible with gamma distributed pieces (same scale and smaller shape). We now consider independent convolutions of these distributions:

Definition 2.1.1 ($\mathcal{G}_{1,n}$ and \mathcal{G}_1). *Let X be a random variable with support \mathbb{R}_+ .*

- (i) *X is a gamma convolution with shapes $\alpha \in \mathbb{R}_+^n$ and scales $s \in \mathbb{R}_+^n$, denoted $X \sim \mathcal{G}_{1,n}(\alpha, s)$, if it has cgf:*

$$K(t) = - \sum_{i=1}^n \alpha_i \ln (1 - s_i t).$$

We denote this class of distributions $\mathcal{G}_{1,n}$.

- (ii) *X is an (untranslated) generalized gamma convolutions with Thorin measure ν , denoted $X \sim \mathcal{G}_1(\nu)$, if its cgf writes:*

$$K(t) = - \int_0^\infty \ln (1 - st) \nu(ds),$$

where ν is such that:

$$\int_{[1,+\infty)} |\log(s)| \nu(ds) < \infty \text{ and } \int_{(0,1)} s \nu(ds) < \infty.$$

We denote by \mathcal{G}_1 this class of distributions.

The generalized gamma convolutions class \mathcal{G}_1 , which is commonly called the one dimensional Thorin class, has received a lot of interest in the literature. These limits of convolutions are not dense in the set of positive continuous random

variables, but they contain many interesting distributions, including the log-Normal and Pareto distributions as historically important cases [54, 55]. They have several interesting properties: noteworthy is the fact that the \mathcal{G}_1 class is closed with respect to independent sums of random variables, by construction, but also, as a more recent result shows, by independent products of random variables [7]. In fact, \mathcal{G}_1 can be defined as the smallest class that is closed through convolution and that contains gamma distributions. See [5] for an extensive study of these distributions².

Remark that if the Thorin measure ν of $X \sim \mathcal{G}_1(\nu)$ has a finite discrete support of cardinal n , $X \sim \mathcal{G}_{1,n}(\boldsymbol{\alpha}, \mathbf{s})$ for $\boldsymbol{\alpha}, \mathbf{s}$ vectors such that $\nu = \sum_i \alpha_i \delta_{s_i}$ (where δ_x is the Dirac measure in x).

Mathai [33] provides a series expression for the density of $\mathcal{G}_{1,n}$ random variables based on Kummer’s confluent hypergeometric functions. Later, Moschopoulos [40] refines the result by providing the following gamma series:

Property 2.1.2 (Moschopoulos expansion). *For $X \sim \mathcal{G}_{1,n}(\boldsymbol{\alpha}, \mathbf{s})$, denoting without loss of generality $s_1 = \min \mathbf{s}$, the density f of X is given by the following series:*

$$f(x) = C \sum_{k=0}^{\infty} \delta_k f_{\rho+k, \frac{1}{s_1}}(x)$$

where:

- $f_{a,b}$ is the density of a $\mathcal{G}_{1,1}(a, b)$ distribution,
- $\rho = |\boldsymbol{\alpha}|$ and $C = s_1^\rho \mathbf{s}^{-\boldsymbol{\alpha}}$,
- $\delta_0 = 0$ and for $k \geq 0$,

$$\delta_{k+1} = \frac{1}{k+1} \sum_{i=1}^{k+1} i \gamma_i \delta_{k+1-i}, \text{ where } \gamma_{k+1} = \sum_{i=1}^n \alpha_i \frac{(1 - \frac{s_1}{s_i})^{k+1}}{k+1}.$$

Furthermore, for $S = 1 - \frac{s_1}{\max \mathbf{s}}$, we have that $|\delta_{k+1}| \leq \frac{S^{k+1} \Gamma(\rho+k)}{\Gamma(\rho)(k+1)!}$, ensuring the uniform convergence of the series.

The dependence of this expansion onto the smallest scale parameter s_1 has a major drawback. When the smallest scale is close to zero, the corresponding truncated series is quite unstable. A simple show-case is the simulation of random numbers from a $\mathcal{G}_{1,2}((10, 10^{-3}), (1, 10^{-3}))$ distribution: the Moschopoulos implementation from the R package `coga` [23] gives a density that evaluates to 0 on all simulated random numbers, which is obviously wrong. Mathai’s method, implemented in the same package, produces the same result. Sadly, as later implementation will show, parameters that correspond to approximations of

²Note that in [5], distributions were parametrized through rates $\beta = \frac{1}{s}$ instead. We chose here to parametrize by scales, as it simplifies some of our notations.

useful distributions such as log-Normal, Weibull or Pareto usually have very small scales and trigger the same instability.

A likelihood approach to fit distributions in the Thorin class is therefore not practical, as there is no stable density, which might explain why there are currently no estimation procedures in the literature.

In [5, 6], Bondesson defined a class of multivariate convolutions of gamma distributions based on the following idea from Cherian [8]. Suppose that $\mathbf{Z} = (Z_i)_{i \in \{0,1,2\}}$ are three independent gamma random variables with (respective) shapes α_i and scales s_i such that $s_0 = 1$ w.l.o.g. Then the cgf of the vector $(Z_0 + Z_1, Z_0 + Z_2)$ writes:

$$K(\mathbf{t}) = -\alpha_1 \ln(1 - s_1 t_1) - \alpha_2 \ln(1 - s_2 t_2) - \alpha_0 \ln(1 - s_1 t_1 + s_2 t_2).$$

This construction can be extended to what Bondesson called multivariate generalized gamma convolutions:

Definition 2.1.3 ($\mathcal{G}_{d,n}$ and \mathcal{G}_d). *Let \mathbf{X} be a d -variate random vector with support \mathbb{R}_+^d .*

- (i) *\mathbf{X} is a multivariate gamma convolution with shapes $\boldsymbol{\alpha} \in \mathbb{R}_+^n$ and scale matrix $\mathbf{s} = (s_{i,j}) \in \mathcal{M}_{n,d}(\mathbb{R}_+)$, denoted by $\mathbf{X} \sim \mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ if it has cumulant generating function:*

$$K(\mathbf{t}) = \sum_{j=1}^n -\alpha_j \ln(1 - \langle \mathbf{s}_j, \mathbf{t} \rangle).$$

We denote by $\mathcal{G}_{d,n}$ this class of distributions.

- (ii) *\mathbf{X} is an (untranslated) multivariate generalized gamma convolutions with Thorin measure ν , denoted $\mathbf{X} \sim \mathcal{G}_d(\nu)$, if it is a weak limit of d -variate gamma convolutions. Its cgf writes:*

$$K(\mathbf{t}) = - \int \ln(1 - \langle \mathbf{s}, \mathbf{t} \rangle) \nu(d\mathbf{s}),$$

for a positive measure ν of \mathbb{R}_+^d , the Thorin measure, with suitable integration conditions (see [45]). We denote by \mathcal{G}_d this class of distributions.

Note that $\mathbf{X} \sim \mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ if and only if there exists a vector \mathbf{Y} of n independent unit scale gamma random variables with shapes $\alpha_1, \dots, \alpha_n$ such that $\mathbf{X} = \mathbf{s}'\mathbf{Y}$. The i^{th} marginal X_i has distribution $\mathcal{G}_{1,n}(\boldsymbol{\alpha}, \mathbf{s}_i)$. We can interpret this as a decomposition of the random vector \mathbf{X} into an additive risk factor model with gamma distributed factors. Note that we allow scales to be zero, and some factors might therefore appear only in some marginals. As in the univariate case, we close the class by taking convolutional limits. For the analysis of these distributions, we refer to [5, 6], but also [45] which uses a slightly different but equivalent parametrization, allowing a generalization to other cones than \mathbb{R}_+^d .

We are here interested in the construction of distributions estimators in \mathcal{G}_d , with particular interest for estimators in $\mathcal{G}_{d,n}$. Indeed, $\mathcal{G}_{d,n}$ models have a structure that allows fast simulation, and that provides meaningful insight about the dependence structure, since a $\mathcal{G}_{d,n}$ distribution follows an additive risk factor model. On the other hand, distributions in \mathcal{G}_d with a non-atomic Thorin measure are hard to sample (see [4, 47] for potential solutions to this problem). They have no known closed-form density or distribution function, and even the cumulant generating function requires integration to be evaluated.

Before diving into statistical considerations and estimation of these distributions in Section 3, we give here some properties about the dependence structure that can be achieved in the \mathcal{G}_d class. The following properties give more insights about the relation between the dependence structure of the random vector and the dependence structure of its Thorin measure:

Property 2.1.4 (The support of the Thorin measure). *Let $\mathbf{X} \sim \mathcal{G}_d(\nu)$, and let $Sup(\nu)$ be the support of ν . Then we have:*

(i) *The marginals of \mathbf{X} are mutually independent if and only if*

$$Sup(\nu) \subset S_{\perp} = \bigcup_{j=1}^d \left(\prod_{i=1}^{j-1} \{0\} \right) \times \mathbb{R}_+^* \times \left(\prod_{i=j+1}^d \{0\} \right).$$

In this case, $\nu(\mathbf{s}) = \sum_{j=1}^d \nu_j(s_j) \mathbb{1}_{s_k=0 \forall k \neq j}$ where ν_j denotes the Thorin measure of the j^{th} marginal, and the total masses of ν, ν_1, \dots, ν_d are such that $\nu(\mathbb{R}_+^d) = \nu(S_{\perp}) = \sum_{i=1}^d \nu_i(\mathbb{R}_+)$.

(ii) *The marginals of \mathbf{X} are comonotonous if and only if there exists a constant $\mathbf{c} \in \mathbb{R}_+^d, \|\mathbf{c}\| = 1$ such that*

$$Sup(\nu) \subset S_{\mathbf{c}} = \left\{ \mathbf{r} \in \mathbb{R}_+^d, \frac{\mathbf{r}}{\|\mathbf{r}\|} = \mathbf{c} \right\}.$$

In this case, the multivariate Thorin measure and its marginals have all the same total mass.

(iii) *The r.v. \mathbf{X} has a Lebesgue-square-integrable density if $D \geq d$, where D denotes the minimum integer such that there exists constants $\mathbf{c}_1, \dots, \mathbf{c}_D$ such that $Sup(\nu) \subset \bigcup_{i=1}^D S_{\mathbf{c}_i}$. On the other hand if $D < d$, \mathbf{X} is a singular random vector with a D -dimensional support (this is obviously the case when ν is atomic with less than d atoms).*

Proof. For (i) and (ii), consider that $\mathbf{X} \sim \mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$. Then there exists a matrix \mathbf{Y} of gamma random variables $Y_{i,j} \sim \mathcal{G}_{1,1}(\alpha_j, s_{i,j})$ such that

$$X_i = Y_{i,1} + \dots + Y_{i,n},$$

where row vectors $\mathbf{Y}_{i,\cdot}$ are independent and column vectors $\mathbf{Y}_{\cdot,j}$ have two separate groups of marginals: a first comonotonous group (with gamma marginals) corresponding to scales $s_{i,j}$ that are strictly positive and a second group (with identically 0 marginals) corresponding to null scales. Note that one of these two groups might be empty. From this expansion, we can easily deduce the structure of maximal supports for ν that leads to independent or comonotonous random vectors \mathbf{X} . Consider non-atomic Thorin measures as limits of the previous case.

For (iii), the problem of regularity of \mathbf{X} and the existence of a density in L_2 , simply consider the rank of the (possibly non-square) matrix \mathbf{s} such that $\mathbf{X} = \mathbf{s}\mathbf{Z}$ where $Z_j \sim \mathcal{G}_{1,1}(\alpha_j, 1)$ are independent. If $D = d = n$, then \mathbf{X} is an invertible linear transformation of an independent, d -dimensional random vector, and therefore has a density in L_2 that we can express easily. If $D = d < n$, then \mathbf{X} is the convolution of the previous case with something else, and therefore also has a density in L_2 . Same goes for $D > d$. \square

Hence, the only comonotonous structure that is allowed inside the \mathcal{G}_d class is $X_i = c_i X_0$ for $X_0 \in \mathcal{G}_1$ and $\mathbf{c} > \mathbf{0}$. Between the independent and the comonotonous cases, there is a wide range of achievable dependence structures. Furthermore, since the class \mathcal{G}_d is closed w.r.t. convex convolution, every (positive) value of dependence measures, such as Kendall tau or Spearman rho, are achievable. Some examples in Section 4 also exhibit highly asymmetrical dependence structure and/or tail dependency.

For the study of the relation between the dependence structure of \mathbf{X} and the dependence structure of ν , the Thorin measure proposed in Example 2.1.5 is of interest.

Example 2.1.5 (A curious distribution). *Let ν be a one-dimensional measure defined through $\nu(dx) = \mathbb{1}_{x \in [0,1]} dx$. Then $\mathcal{G}_1(\nu)$ has moment generating function:*

$$M(t) = e^{1 - \frac{(t-1)}{t} \ln(1-t)}$$

Moreover, if for all $n \in \mathbb{N}$, $G_{j,n} \sim \mathcal{G}_{1,1}\left(\frac{1}{n}, \frac{j}{n+1}\right)$ are independent, then:

$$\lim_{n \rightarrow +\infty} \sum_{j=1}^n G_{j,n} \sim \mathcal{G}_1(\nu).$$

Proof. The second result comes from a discretization of the Thorin measure, from continuously uniform on $[0, 1]$ to uniform on $\left\{\frac{j}{n+1}, j \in \{1, \dots, n\}\right\}$. For the first result, note that ν satisfies the integration constraints given by Definition 2.1.1. Then, simply compute the cumulant generating function as:

$$\begin{aligned} K(-t) &= - \int_0^\infty \ln(1+rt) \mathbb{1}_{r \in [0,1]} dr = - \int_0^1 \ln(1+rt) dr \\ &= - \frac{1}{t} \int_1^{1+t} \ln(x) dx = 1 - \frac{1+t}{t} \ln(1+t). \end{aligned}$$

□

This random variable has the nice property that its Thorin measure is uniform on $[0,1]$. It could be simulated more efficiently than by summing gammas through [4, 47] if we wanted. However, its use is limited by the following negative result:

Remark 2.1.6 (The no-bijection result). *Let $\mathbf{X} \sim \mathcal{G}_d(\nu)$ have marginal cdf F_1, \dots, F_d , and denote by G the cdf associated to the distribution of Example 2.1.5. Then, assuming that ν is absolutely continuous, the random vector*

$$\mathbf{Y} = (G^{-1}(F_i(X_i)))_{i \in 1, \dots, d}$$

has uniform marginal Thorin measures, but is not $\mathcal{G}_d(c_\nu)$ -distributed, where c_ν is the copula of ν .

There is in fact a bijection between the copula of the random vector and the copula of its Thorin measure, but conditionally on the marginal distributions. Therefore, an estimation scheme that separates the dependence structure from the marginals is not possible.

Before discussing the estimation of these distributions, we still need to introduce in the next subsection some specific integrals, which correspond to derivatives of the moment generating function and the cumulant generating function of a random vector.

2.2. Shifted moments and shifted cumulants

We formally introduce the shifted moments and the cumulants of a random vector, and present a known and useful result that maps moments to cumulants and vice-versa.

For \mathbf{X} a random vector of dimension d , we recall the following notations for, respectively, the moment generating function M and the cumulant generating function K of \mathbf{X} :

$$\forall \mathbf{t} \in \mathbb{C}^d, M(\mathbf{t}) = \mathbb{E} \left(e^{\langle \mathbf{t}, \mathbf{X} \rangle} \right) \text{ and } K(\mathbf{t}) = \ln M(\mathbf{t}).$$

We now introduce specific notations for the derivatives of these two functions.

Definition 2.2.1 (Shifted moments and cumulants). For \mathbf{X} a random vector of dimension d , for $\mathbf{i} \in \mathbb{N}^d$, define the \mathbf{t} -shifted \mathbf{i}^{th} moments and cumulants as the \mathbf{i}^{th} derivatives of (respectively) M and K , taken in \mathbf{t} :

$$\mu_{\mathbf{i},\mathbf{t}} = M^{(\mathbf{i})}(\mathbf{t}) \text{ and } \kappa_{\mathbf{i},\mathbf{t}} = K^{(\mathbf{i})}(\mathbf{t}).$$

They correspond to the moment and cumulants of an exponentially tilted version of the random vector \mathbf{X} , also called the Escher transform [17] of \mathbf{X} .

Although standard Monte-Carlo estimators for shifted moments are unbiased and easy to compute, shifted cumulants are a little harder to estimate from i.i.d. samples of a random vector. We refer to [51, 50] for some unbiased estimators of multivariate cumulants, and [28] for an application to the estimation of stable laws. To switch from moments to cumulants, Property 2.2.2 gives a bijection between the two:

Property 2.2.2 (Bijection between moments and cumulants). For a given shift $\mathbf{t} \in \mathbb{C}^d$ and $\mathbf{m} \in \mathbb{N}^d$, the sets $(\mu_{\mathbf{i},\mathbf{t}})_{\mathbf{i} \leq \mathbf{m}}$ and $(\kappa_{\mathbf{i},\mathbf{t}})_{\mathbf{i} \leq \mathbf{m}}$ can express one from one another through:

$$\begin{aligned} \mu_{\mathbf{i},\mathbf{t}} &= e^{\kappa_{\mathbf{0},\mathbf{t}}} \sum_{\pi \in \Pi(\mathbf{i})} \prod_{B \in \pi} \kappa_{\mathbf{i}(B),\mathbf{t}} \\ \kappa_{\mathbf{i},\mathbf{t}} &= \sum_{\pi \in \Pi(\mathbf{i})} (|\pi| - 1)! (-1)^{|\pi|-1} (\mu_{\mathbf{0},\mathbf{t}})^{-|\pi|} \prod_{B \in \pi} \mu_{\mathbf{i}(B),\mathbf{t}} \end{aligned}$$

where $\Pi(\mathbf{i}) = \mathcal{P}(B(\mathbf{i}))$ is the set of all partitions of the multi-set $B(\mathbf{i}) = \bigcup_{j=1}^d \bigcup_{p=1}^{i_j} \{j\}$, a partition π in $\Pi(\mathbf{i})$ is constituted of multi-sets B composed of integers, and $\mathbf{i}(B) = (\#\{\text{occurrences of } j \text{ in } B\})_{j \in \{1, \dots, d\}}$.

Proof. This is a direct application of the multivariate Faà Di Bruno's formula [10, 39, 42] on $M = \exp \circ K$ in one hand and $K = \ln \circ M$ in the other hand. \square

Remark that this property can be expressed in several other ways: most analysis in the literature uses multivariate Bell Polynomials [3, 12, 26, 35] for this task. For computational purposes, there also exists some practical functions that give coefficients of these polynomials [57, 19], even for $d > 1$. However, in all generality, the number of coefficients to compute and store is exponential with \mathbf{m} making this method quite unpractical, as the Bell polynomials are combinatorial in nature. To switch from cumulants to moments, the good property of the exponential function of being its own derivative allows a recursive approach, which is described in full details in [37]. We adapted the notation and simplified the indexes conventions from [37], which allowed, after adding a slack variable to handle the dimensionality index, for an even faster implementation than the one from [37]. The expression in Algorithm 1 below uses derivatives of the cgf

and mgf as it is the case that we need, but the derivative of the exponential of any function could be computed through exactly the same algorithm, knowing the derivatives of the exponent.

Algorithm 1: Recursive computation of $\boldsymbol{\mu}$ from $\boldsymbol{\kappa}$.

Input: $\mathbf{m} \in \mathbb{N}^d$, Shifted Cumulants $(\kappa_{\mathbf{k},t})_{\mathbf{k} \leq \mathbf{m}}$
Result: Shifted Moments $(\mu_{\mathbf{k},t})_{\mathbf{k} \leq \mathbf{m}}$
Set $\mu_{\mathbf{0},t} = \exp(\kappa_{\mathbf{0},t})$
Set $\mu_{\mathbf{k},t} = 0$ for all $\mathbf{k} \neq \mathbf{0}$
foreach $\mathbf{k} : \mathbf{0} \neq \mathbf{k} \leq \mathbf{m}$ **do**
 Set d as the index of the first k_i that is non-zero.
 Set $\mathbf{p} = \mathbf{k}$
 Set $p_d = p_d - 1$
 foreach $\mathbf{l} : \mathbf{l} \leq \mathbf{p}$ **do**
 | $\mu_{\mathbf{k},t} += (\mu_{\mathbf{l},t}) (\kappa_{\mathbf{k}-\mathbf{l},t}) \binom{\mathbf{p}}{\mathbf{l}}$
 end
end
Return $\boldsymbol{\mu}$

Note that the main loop of Algorithm 1 must be done in the right order such that each $\mu_{\mathbf{k}}$ is computed before being used. This recursive formulation has stunning results and is at the moment the fastest way this computation can be done³. Unfortunately, no equivalently fast procedure can be found the other way around, as the derivatives of the log function are more complicated.

We now dive into the estimation of multivariate gamma convolutions. We highlight in the next subsection the univariate projection procedure from [18, 38].

2.3. Projection from \mathcal{G}_1 to $\mathcal{G}_{1,n}$

Suppose that we have a density in \mathcal{G}_1 . Then, for a certain $n \in \mathbb{N}$, the procedure provided by Miles, Furman & Kuznetsov in [38, 18], hereafter denoted MFK, gives expression for the shapes $\boldsymbol{\alpha}$ and scales \mathbf{s} of a $\mathcal{G}_{1,n}(\boldsymbol{\alpha}, \mathbf{s})$ that is exponentially close to the desired density. We will not describe the algorithm here, we refer to the original article for the details as the exposition is quite long, but we will comment on the results and caveats and use it as a comparison basis for later experiments.

A deep result about Padé approximation of Stieltjes functions supports the MFK procedure. Indeed, the first derivative of the cumulant generating function of a random variable in \mathcal{G}_1 writes (using the rate notation $r = \frac{1}{s}$ and the corresponding push-forward ν_r of the Thorin measure ν):

$$K^{(1)}(-t) = \int \frac{1}{r+t} \nu_r(dr).$$

³The complexity is not exponentially increasing in the size of the arrays, on the contrary to naive implementations of Faà Di Bruno's formula, see [37].

When ν_r is discrete with n atoms, this expression is a one-dimensional Stieltjes function. This Stieltjes function can be approximated from $(\mu_i)_{i \leq 2n}$ by $[m-1/m]$ standard Padé approximants, (see [13]), whose denominator has zeros that are all simple, real, and negative. Applying a partial fraction decomposition to the Padé approximation is a direct way to estimate the rates and shapes of the gammas, since the obtained form for $K^{(1)}$, and therefore $K(t)$, matches the form given in Definition 2.1.1 for a $\mathcal{G}_{1,n}$ model.

The MFK algorithm has two major drawbacks. The first one is that the partial fraction decomposition of the Padé approximation, central in the procedure, is typical to univariate Stieltjes functions. Since the first derivative of the cgf is a Stieltjes function only in \mathcal{G}_1 , the method cannot be used to approximate a random variable that is not in the \mathcal{G}_1 class. Since equivalent results do not exist for multivariate Stieltjes functions, the method is restricted to one-dimensional problems. The second one is that precision of computations needs to be at 1024 bits or higher for the algorithm to produce a result at all.

We do have an explanation for each of these drawbacks. In all generality, the Padé approximation of the Stieltjes function is a special case of a more general problem, a moment problem for the Thorin measure. Indeed, remark that:

$$\kappa_{k,t} = (k-1)! \int \frac{s^k}{(1+st)^k} \nu(ds),$$

and use the push-forward $x = \frac{s}{1+st}$, mapping \mathbb{R}_+ to the one-dimensional t -simplex $\Delta_1(t) = [0, \frac{1}{t}]$. You obtain a moment problem of the form:

$$\xi_k = \int_{\Delta_1(t)} x^k \xi(dx), \forall k \leq m,$$

where $\xi = (\xi_k)_{k \leq m} = \left(\frac{\kappa_{k,t}}{(k-1)!} \right)_{k \leq m}$ is the vector of moments that the measure ξ needs to fit.

We can then obtain a solution with n atoms for the measures ξ and ν by solving this moment problem using classical tools, see, e.g., [48] for an extensive discussion of potential algorithms. When the vector of cumulants κ comes from a density in the \mathcal{G}_1 class, the vector of moments ξ of the measure ξ is inside the moment cone of positive measure on $\Delta_1(t)$. We therefore obtain a proper measure ξ matching these moments. On the other hand, if ξ is *not* a sequence of moments of some measure on the simplex, which happens as soon as the random variable is *not* in \mathcal{G}_1 , then the moment problem becomes unsolvable. In this case, the approximation of the Stieltjes function fails to provide correct (within boundaries) atoms and weights for the measure ξ , and moreover for the Thorin measure.

Experiments from [38, 18] suggests that, if the density used to compute the shifted moments μ is in \mathcal{G}_1 , the integrations of moments μ and the derivation

of $\boldsymbol{\kappa}$ and $\boldsymbol{\xi}$ from $\boldsymbol{\mu}$ need to be carried out with at least 300 digits of accuracy (about 1024 bits) so that $\boldsymbol{\xi}$ is in the desired cone of moments. From an empirical dataset, even if the true distribution lies in \mathcal{G}_1 , evaluation of $\boldsymbol{\mu}$ through Monte-Carlo with enough precision is therefore impossible.

Last but not least, the results about the complete factorization of Padé approximants denominator of a Stieltjes function, which simplifies the moment problem, are not true for $d > 1$. The approach could nevertheless be carried out, as the corresponding moment problem still has a solution if the used cumulants are true cumulants from a \mathcal{G}_d distribution, which is an even stronger precision requirement as the dimension increases. This possibility is discussed in the next section.

3. Construction of an estimator in $\mathcal{G}_{d,n}$

We attempt a multivariate extension of the MFK algorithm. In a multivariate setting, the Padé approximants of the cumulant generating function exists but has a denominator that is not known to be completely factorizable. Thus, we cannot take a partial fraction expansion to obtain the cumulant generating function in its canonical form to match the coefficients. Nevertheless, through the moment problem interpretation of MFK algorithm developed in Subsection 2.3, we can still express and try to solve the corresponding multivariate moment problem.

For $\mathbf{X} \sim \mathcal{G}_d(\nu)$, consider K the cumulant generating function of \mathbf{X} , given in Definition 2.1.3 as:

$$K(\mathbf{t}) = - \int \ln(1 - \langle \mathbf{s}, \mathbf{t} \rangle) \nu(d\mathbf{s}).$$

Assuming $|\mathbf{i}| \geq 1$ and $\mathbf{t} \leq \mathbf{0}$, we have that

$$\frac{\partial^{\mathbf{i}}}{\partial^{\mathbf{i}} \mathbf{t}} \ln(1 - \langle \mathbf{s}, \mathbf{t} \rangle) = -\mathbf{s}^{\mathbf{i}} (|\mathbf{i}| - 1)! (1 - \langle \mathbf{s}, \mathbf{t} \rangle)^{-|\mathbf{i}|},$$

and, therefore,

$$\begin{aligned} \kappa_{\mathbf{i}, \mathbf{t}} = K^{(\mathbf{i})}(\mathbf{t}) &= - \frac{\partial^{\mathbf{i}}}{\partial^{\mathbf{i}} \mathbf{t}} \int \ln(1 - \langle \mathbf{s}, \mathbf{t} \rangle) \nu(d\mathbf{s}) \\ &= \int \mathbf{s}^{\mathbf{i}} (|\mathbf{i}| - 1)! (-1)^{|\mathbf{i}|-1} (1 - \langle \mathbf{s}, \mathbf{t} \rangle)^{-|\mathbf{i}|} \nu(d\mathbf{s}) \\ &= (|\mathbf{i}| - 1)! \int \frac{\mathbf{s}^{\mathbf{i}}}{(1 - \langle \mathbf{s}, \mathbf{t} \rangle)^{|\mathbf{i}|}} \nu(d\mathbf{s}). \end{aligned} \tag{1}$$

Note that $\mathbf{t} \leq \mathbf{0}$ is inside the region of convergence of the function, and use now the (bijective, continuous) push-forward:

$$\mathbf{x} = \frac{\mathbf{s}}{1 - \langle \mathbf{s}, \mathbf{t} \rangle} \iff \mathbf{s} = \frac{\mathbf{x}}{1 + \langle \mathbf{x}, \mathbf{t} \rangle},$$

which goes from \mathbb{R}_+^d to the simplex $\Delta_d(-\mathbf{t}) = \{\mathbf{x} \in \mathbb{R}_+^d, \langle \mathbf{x}, -\mathbf{t} \rangle \leq 1\}$, yielding an equivalent problem of finding the measure ξ that solves the more standard moment problem:

$$\forall \mathbf{i} \leq \mathbf{m}, \xi_{\mathbf{i}} = \int_{\Delta_d(-\mathbf{t})} \mathbf{x}^{\mathbf{i}} \xi(d\mathbf{x}), \quad (2)$$

where we denoted $\xi = \left(\frac{\kappa_{\mathbf{i}, \mathbf{t}}}{(|\mathbf{i}|-1)!} \right)_{\mathbf{i} \leq \mathbf{m}}$ the moments that ξ needs to match according to Equation (1).

Finding an n -atomic measure ξ solution to these equations is equivalent to finding the parametrization of a $\mathcal{G}_{d,n}$ distribution that fulfills the cumulant constraints. The corresponding Generalized Moment Problem is known as a hard problem in the literature, but can be solved through semi-definite positive relaxations, following, e.g., [20, 43, 21]. Sadly, these algorithms and the corresponding literature focus on the first case, where ξ are, indeed, moments of a measure supported on the simplex $\Delta_d(-\mathbf{t})$, which is the case only if $\mathbf{X} \in \mathcal{G}_{d,n}$ and if ξ are computed with enough precision.

If the knowledge about the distribution of \mathbf{X} is given through empirical observations, with sampling noise, or if the true distribution is *not* in the $\mathcal{G}_{d,n}$ class, these moments will not belong to the moment cone, and the moment problem will have no solution. By projecting the moments onto the moment cone, we could force the known algorithms to provide a solution, but the projection onto this cone is not an easy task (see [14, 15] for details on moment cones).

The optimization of an n -atomic Thorin measure from samples of the random vector \mathbf{X} therefore rises several questions, both in univariate and multivariate cases. If we empirically compute cumulants, we will stand outside the moment cone and the moment problem will have no solution. Then, should we consider Equation (2) as a multi-objective optimization problem? Do we seek a Pareto front, or can we consider that some moments are less important than others? If we do, how do we weight the several objectives?

To produce a meaningful loss to minimize, we will weight these moment conditions. For this purpose, we leverage a certain Laguerre basis of $L_2(\mathbb{R}_+^d)$ to expand the density of random vectors in $\mathcal{G}_{d,n}$, and to construct a coherent loss for this, potentially impossible to solve, moment problem.

3.1. The tensorized Laguerre basis

The standard Laguerre polynomials [41] form an orthogonal basis of the set $L_2(\mathbb{R}_+, e^{-x})$ of square integrable functions with respect to the weight function $x \mapsto e^{-x}$. From these polynomials, we can extract the following orthonormal basis of the set $L_2(\mathbb{R}_+^d)$ of functions that are Lebesgue square-integrable on \mathbb{R}_+^d :

Definition 3.1.1 (Laguerre function). [9, 32, 16] For all $\mathbf{k} \in \mathbb{N}^d$, define the Laguerre functions $\varphi_{\mathbf{k}}$ with support \mathbb{R}_+^d by:

$$\varphi_{\mathbf{k}}(\mathbf{x}) = \prod_{i=1}^d \varphi_{k_i}(x_i) \text{ where for } k \in \mathbb{N}, \varphi_k(x) = \sqrt{2}e^{-x} \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-2x)^\ell}{\ell!}.$$

For every function $f \in L^2(\mathbb{R}_+^d)$, for every $\mathbf{k} \in \mathbb{N}^d$, coefficients of f in the basis are denoted by:

$$a_{\mathbf{k}}(f) = \int_{\mathbb{R}_+^d} f(\mathbf{x})\varphi_{\mathbf{k}}(\mathbf{x}) d\mathbf{x},$$

and we have, since the basis is orthonormal, that $f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}}(f)\varphi_{\mathbf{k}}(\mathbf{x})$ and that $\|f\|_2^2 = \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}}^2(f)$.

Furthermore, for \mathbf{X} a random vector with density $f \in L_2(\mathbb{R}_+^d)$ and shifted moments $\boldsymbol{\mu}$, we have by simple plug-in of f and $\varphi_{\mathbf{k}}$ into the expression of $a_{\mathbf{k}}$ that:

$$a_{\mathbf{k}}(f) = \sqrt{2}^d \sum_{\boldsymbol{\ell} \leq \mathbf{k}} \binom{\mathbf{k}}{\boldsymbol{\ell}} \frac{(-2)^{|\boldsymbol{\ell}|}}{\boldsymbol{\ell}!} \boldsymbol{\mu}_{\boldsymbol{\ell}, -1}. \quad (3)$$

We now denote $a_{\mathbf{k}}(\boldsymbol{\alpha}, \mathbf{s})$ the \mathbf{k}^{th} Laguerre coefficient of the $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ distribution.

Property 3.1.2 ($\mathbf{X} \sim \mathcal{G}_{d,1}(\boldsymbol{\alpha}, \mathbf{s})$). For $\mathbf{X} \sim \mathcal{G}_{d,1}(\boldsymbol{\alpha}, \mathbf{s})$, we have a bijection between the $d + 1$ first Laguerre coefficients and the $d + 1$ parameters, given through the following equations:

$$a_{\mathbf{k}}(\boldsymbol{\alpha}, \mathbf{s}) = \sqrt{2}^d \sum_{\boldsymbol{\ell} \leq \mathbf{k}} \binom{\mathbf{k}}{\boldsymbol{\ell}} \frac{(-2\mathbf{s})^\ell}{\boldsymbol{\ell}!} \frac{\Gamma(\boldsymbol{\alpha} + |\boldsymbol{\ell}|)}{\Gamma(\boldsymbol{\alpha})} (1 + |\mathbf{s}|)^{-\boldsymbol{\alpha} - |\boldsymbol{\ell}|},$$

on one hand, and

$$\alpha = \left\{ \frac{1}{c_2} - \frac{1}{\ln(c_1)} W_0 \left(\frac{\ln(c_1)}{c_2} e^{\frac{\ln(c_1)}{c_2}} \right) \right\}$$

$$s_i = \frac{a_{\mathbf{0}} - a_{\mathbf{1}(i)}}{2\alpha a_0} \text{ for all } i \in 1, \dots, d$$

reciprocally, where:

- $\mathbf{1}(i)$ is the d -variate vector with j^{th} component equal to $\mathbb{1}_{i=j}$, for all j .
- $c_1 = a_{\mathbf{0}}\sqrt{2}^{-d}$ and $c_2 = \frac{d}{2} - \frac{1}{2a_0} \sum_{i=1}^d a_{\mathbf{1}(i)}$
- W_0 is the zeroth branch of the Lambert function W defined by:

$$y = W(x) \iff x = ye^y.$$

Proof. For the first part, note that there exists $X_0 \sim \mathcal{G}_{1,1}(\alpha, 1)$ such that $\mathbf{X} = \langle \mathbf{s}, (X_0, \dots, X_0) \rangle$. Therefore, denoting by M the mgf of X_0 ,

$$\begin{aligned} \mu_{\mathbf{k}, -\mathbf{1}} &= \mathbb{E} \left(\mathbf{X}^{\mathbf{k}} e^{-\langle \mathbf{1}, \mathbf{X} \rangle} \right) = \mathbf{s}^{\mathbf{k}} \mathbb{E} \left(X_0^{|\mathbf{k}|} e^{-|\mathbf{s}| X_0} \right) \\ &= \mathbf{s}^{\mathbf{k}} M^{(|\mathbf{k}|)}(-|\mathbf{s}|) \\ &= \mathbf{s}^{\mathbf{k}} \frac{\Gamma(\alpha + |\mathbf{k}|)}{\Gamma(\alpha)} (1 + |\mathbf{s}|)^{-\alpha - |\mathbf{k}|}. \end{aligned}$$

Therefore, the Laguerre coefficients $a_{\mathbf{k}}(\alpha, \mathbf{s})$ of \mathbf{X} are given by:

$$a_{\mathbf{k}}(\alpha, \mathbf{s}) = \sqrt{2^d} \sum_{\ell \leq \mathbf{k}} \binom{\mathbf{k}}{\ell} \frac{(-2\mathbf{s})^\ell}{\ell!} \frac{\Gamma(\alpha + |\ell|)}{\Gamma(\alpha)} (1 + |\mathbf{s}|)^{-\alpha - |\ell|}.$$

Then, for the reverse assertion, consider the expression of the first coefficients:

$$a_{\mathbf{0}} = \sqrt{2^d} (1 - |\mathbf{s}|) \text{ and } a_{\mathbf{1}(i)} = a_{\mathbf{0}} \left(1 - 2\alpha \frac{s_i}{1 + |\mathbf{s}|} \right).$$

Use the simplex transformation $\mathbf{x} = \frac{\mathbf{s}}{1 + |\mathbf{s}|}$ to solve for \mathbf{s} as a function of α in the second equation, and inject in the first. We obtain an equation of the form $ae^x + bx + c$ which is solved by the zeroth branch of the Lambert W function. \square

This example is quite peculiar since by Property 2.1.4 the random vector is singular. Note that the edge case $\exists i \in 1, \dots, d, r_i = 0$, in which the i^{th} marginal is identically zero, is included in the previous result.

We do not know if the same kind of inversion could be carried out analytically for the coefficients of $\mathcal{G}_{d,n}$ densities, since Laguerre coefficients are not as trivially expressed in all generality. Furthermore, this kind of parametric inversion would not work for empirical datasets or densities outside the class, and is therefore not our focus here. However, we can efficiently compute coefficients from parameters

of a $\mathcal{G}_{d,n}$ random vector through Algorithm 2, based on Algorithm 1:

Algorithm 2: Laguerre coefficients of $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ random vectors.

Input: Shapes $\boldsymbol{\alpha} \in \mathbb{R}_+^n$, scales $\mathbf{s} \in \mathcal{M}_{n,d}(\mathbb{R}_+)$, and truncation threshold $\mathbf{m} \in \mathbb{N}^d$

Result: Laguerre coefficients $\mathbf{a} = (a_{\mathbf{k}})_{\mathbf{k} \leq \mathbf{m}}$ of the $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ density

Side-product: -1 -shifted cumulants $\boldsymbol{\kappa}$ and moments $\boldsymbol{\mu}$ up to order \mathbf{m} of the $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ distribution.

Compute the simplex version of the scales $\mathbf{x}_i = \frac{s_i}{1+|s_i|}$ for all $i \in 1, \dots, n$.

Let $\kappa_{\mathbf{0}} = -\sum_{i=1}^n \alpha_i \ln(1 - |\mathbf{x}_i|)$

Let $a_{\mathbf{0}} = \mu_{\mathbf{0}} = \exp(\kappa_{\mathbf{0}})$

foreach $\mathbf{0} \neq \mathbf{k} \leq \mathbf{m}$ **do**

Let $a_{\mathbf{k}} = \mu_{\mathbf{k}} = 0$

Let d be the index of the first k_i that is non-zero.

Let $\mathbf{p} = \mathbf{k}$

Set $p_d = p_d - 1$

Let $\kappa_{\mathbf{k}} = (|\mathbf{k}| - 1)! \sum_{i=1}^n \alpha_i \mathbf{x}_i^{\mathbf{k}}$ according to Eq. (1)

foreach $l \leq \mathbf{p}$ **do**

Set $\mu_{\mathbf{k}} += (\mu_l) (\kappa_{\mathbf{k}-l}) \binom{\mathbf{p}}{\mathbf{l}}$ according to Algorithm 1

Set $a_{\mathbf{k}} += \mu_l \binom{\mathbf{k}}{\mathbf{l}} \frac{(-2)^{|\mathbf{l}|}}{\mathbf{l}!}$ according to Eq. (3)

end

Set $a_{\mathbf{k}} += \mu_{\mathbf{k}} \frac{(-2)^{|\mathbf{k}|}}{\mathbf{k}!}$

end

$\mathbf{a} = \sqrt{2}^d \mathbf{a}$

Return \mathbf{a}

The complexity of Algorithm 2 is quadratic in \mathbf{m} . Note that, as in Algorithm 1, computations need to be performed in the right order. Sometimes, the Laguerre coefficients $\mathbf{a} = (a_{\mathbf{k}})_{\mathbf{k} \leq \mathbf{m}}$ overflow the Float64 limits, but the implementation we provide in the Julia package `ThorinDistributions.jl`⁴, archived in Zenodo [29], ensures that the computations do not overflow by using multiple precision arithmetic when needed. Furthermore, as in Algorithm 1, the use of Miatto’s fast recursion gives this algorithm a good efficiency, even for reasonably large d, n, \mathbf{m} .

By reordering the coefficients and leveraging generalizations of Laguerre polynomials, we have the following link with Moschopoulos’s density:

Remark 3.1.3 (Generalized Laguerre basis and Moschopoulos density). *For a random vector $\mathbf{X} \sim \mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ with Laguerre coefficients \mathbf{a} , the density f could*

⁴`ThorinDistributions.jl`: <https://github.com/lrnv/ThorinDistributions.jl>.

be rewritten as a gamma series:

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{N}^d} b_{\mathbf{k}} f_{\mathbf{k}}(x), \text{ where } f_{\mathbf{k}}(x) = \frac{\mathbf{x}^{\mathbf{k}} e^{|\mathbf{x}|}}{\mathbf{k}!}$$

$$\text{and } b_{\mathbf{k}} = \sqrt{2}^d (-2)^{|\mathbf{k}|} \sum_{\ell \geq \mathbf{k}} \binom{\ell}{\mathbf{k}} a_{\ell}.$$

We note that $f_{\mathbf{k}}$ is the density of a vector with independent $\mathcal{G}_{1,1}(k_i, 1)$ marginals. Laguerre's polynomials can be generalized by adding a parameter $\rho \geq 0$, which we took here to be 0. This alteration would reflect in the series by yielding densities of independent random vectors with $\mathcal{G}_{1,1}(k_i + \rho, 1)$ marginals instead. This hyperparameter could be varied to produce different concurrent models, and could even be chosen dimension wise. Noteworthy is the fact that, in the one-dimensional case, if we chose ρ to be the total mass of the Thorin measure, we retrieve Moschopoulos series from Property 2.1.2 as a Laguerre expansion.

Sadly, even if mixtures of gammas are easier to fit by k -MLE procedure [44, 49], we have no way of identifying the subspace of coefficients that would match a true convolution of gammas: as [5] noted in the univariate case, generalized gamma convolutions can be expressed as mixtures of gammas, but there is no simple reverse mapping. Last but not least, if the expansion through generalized Laguerre polynomials with a parameter corresponding to the total mass of the Thorin measure would converge faster, we would have no way of estimating this parameter beforehand from data, and for many useful cases it would be infinite (log-Normals among others).

However, under simple technical assumptions on the parameters of a $\mathcal{G}_{d,n}$ random vector, Laguerre coefficients are exponentially decreasing. Property 3.1.8 will provide such a bound. Before exposing the main result, we start by a few technical statements. Lemma 3.1.4 gathers some useful properties about a particular univariate Möbius transform.

Lemma 3.1.4 (Property of a Möbius transform). *Let h be the Möbius transform:*

$$h(t) = \frac{t + 1}{t - 1}.$$

For $a \in \mathbb{C}$ and $b \in \mathbb{R}_+$, we denote the disc with center a and radius b by:

$$D(a, b) = \{t \in \mathbb{C}, |t - a| < b\}.$$

Also denote $c(t) = \frac{t^2 + 1}{t^2 - 1}$ and $r(t) = \left| \frac{2t}{t^2 - 1} \right|$, and finally denote \mathbb{C}_- and \mathbb{C}_+ the left half and right half of the complex plane. The following holds:

- (i) h is its own inverse, i.e., $h(h(t)) = t$, and h has a simple pole at $t = 1$ which has limit $+\infty$ if $\text{Re}(t) > 1$ and $-\infty$ if $\text{Re}(t) < 1$, $t \rightarrow 1$.
- (ii) $h(1/t) = -h(t)$ for all $t \in \mathbb{C}$ and $|h(t)| > 1$ for all $\text{Re}(t) > 0$, $\text{Im}(t) < +\infty$

- (iii) $h(D(0, 1)) = \mathbb{C}_-$
- (iv) If $b < 1$, $h(D(0, b)) = D(c(b), r(b)) \subset \mathbb{C}_-$
- (v) If $b > 1$, $\mathbb{C} \setminus h(D(0, b)) = D(c(b), r(b)) \subset \mathbb{C}_+$

Proof. Point (i) and (ii) are obvious. To prove the 3 remaining points, we use the fact that h is a Möbius transform, see [1], and hence maps generalized circles to generalized circles. Furthermore, h maps reals to reals, and we obtain $c(b)$ and $r(b)$ from the expression of $h(b)$ and $h(-b)$. Last, by checking if the disk with center $c(b)$ and radius $r(b)$ contains $h(0) = -1$, we conclude. \square

Definition 3.1.5 provides a specific regularity assumption on a parametrization of generalized gamma convolution that is a little cumbersome to describe but will be helpful later.

Definition 3.1.5 (ε -well-behaved gamma convolution). *A gamma convolution $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ is said to be ε -well-behaved, in short ε -w.b., if $|\boldsymbol{\alpha}| > 1$ and $\forall I \subseteq \{1, \dots, n\}$ such that $\sum_{i \in I} \alpha_i > \sum_{i \notin I} \alpha_i$, denoting $\mathbf{s}_I = (\mathbf{s}_i)_{i \in I}$, for any solution \mathbf{t}^* of the linear system of equations $\mathbf{s}_I \mathbf{t} = \mathbf{1}$, at least for one $j \in 1, \dots, d$, $|h(t_j^*)| > 1 + \varepsilon$.*

An equivalent statement when $d = 1$ is simply that $|\alpha| > 1$ and $\mathbf{s} \in]\frac{\varepsilon}{2+\varepsilon}, \frac{2+\varepsilon}{\varepsilon}[^n$

If there exists $\varepsilon > 0$ such that the model is ε -w.b., we say that the model is well-behaved, in short w.b.

Of course, any ε_1 -w.b. gamma convolution is ε_2 -w.b. for any $0 < \varepsilon_2 < \varepsilon_1$. Property 3.1.6 gives equivalent statements for the well-behavior of a gamma convolution, hopefully clarifying the requirements from Definition 3.1.5.

Property 3.1.6 (Well-behaved gamma convolution). *For a gamma convolution $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ with Thorin measure ν , the following statements are equivalent:*

- (i) $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ is well-behaved.
- (ii) $\exists \varepsilon > 0$ such that $\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s})$ is ε -well-behaved.
- (iii) $|\boldsymbol{\alpha}| > 1$ and $\forall I \subseteq \{1, \dots, n\}$ such that $\sum_{i \in I} \alpha_i > \sum_{i \notin I} \alpha_i$, $\text{Ker}(\mathbf{s}_I) = \{\mathbf{0}\}$.
- (iv) $\nu(\mathbb{R}_+^d) > 1$ and there exists no $\mathbf{c}_1, \dots, \mathbf{c}_p$ linearly independent, $p < d$, such that

$$\nu \left(\bigcup_{i=1}^p S_{\mathbf{c}_i} \right) > \frac{\nu(\mathbb{R}_+^d)}{2},$$

where $S_{\mathbf{c}}$ was defined in Property 2.1.4.

Proof. First, (i) \iff (ii) is just the definition of well-behavior. Furthermore, (iv) is equivalent to (iii): a subset of scales \mathbf{s}_I such that $\sum_{i \in I} \alpha_i > \sum_{i \notin I} \alpha_i$ is a subset of scale that has the majority of the weight of the Thorin measure

assigned to it, and $\text{Ker}(\mathbf{s}_I) = 0$ if and only if at least d linearly independent directions are inside \mathbf{s}_I . We therefore show only (ii) \iff (iii).

By Rouché-Capelli's Theorem, see [53], any complex solution (if there are some) of the linear system $\mathbf{s}_I \mathbf{t} = \mathbf{1}$ is unique if and only if $\text{Ker}(\mathbf{s}_I) = \{\mathbf{0}\}$. If it exists, call this solution $\mathbf{t}^* \in \mathbb{C}^d$.

Now, since for a given $i \in I$, $s_{i,j}$ are all non-negative reals and one of them must not be zero, \mathbf{t}^* must have at least one component t_j^* that has a positive real part. Furthermore, this component is clearly of bounded modulus since the solution is unique and all $s_{i,j}$ are finite.

Finally, the requirement $|h(t_j^*)| > 1 + \varepsilon$ is equivalent, through Lemma 3.1.4, point (v), to $|t_j^* - c(1 + \varepsilon)| < r(1 + \varepsilon)$ (where c and r were defined in Lemma 3.1.4).

Since the disk $D(c(1 + \varepsilon), r(1 + \varepsilon))$ is in \mathbb{C}_+ , and tends to the whole \mathbb{C}_+ as ε goes to 0 by Lemma 3.1.4, point (iv), there always exist a $\varepsilon > 0$ so that a bounded t_j^* with positive real part is inside this disk. \square

The proof of Property 3.1.6 gives a way of finding the greatest ε that makes a well-behaved gamma convolution ε -w.b. for the general multivariate case. In a certain sense, the ε -w.b. property is a quantification of the statements from Property 2.1.4: the Thorin measure must not assign half of its mass or more to any subset that would correspond to a non absolutely continuous gamma convolution. A surprise comes in the next examples, where we show the unexpected fact that the w.b. subclass of gamma convolutions is closed w.r.t finite convolutions.

Example 3.1.7 (Simple w.b. examples). *As soon as the total mass of their Thorin measure is greater than one, the following gamma convolutions are well-behaved:*

- (i) *Any univariate gamma convolution (the best ε is given in Definition 3.1.5).*
- (ii) *Any independent random vector in $\mathcal{G}_{d,n}$.*
- (iii) *Any invertible linear transformation of an independent random vector in $\mathcal{G}_{d,n}$.*
- (iv) *Any finite convolution of well-behaved gamma convolutions.*

Proof. The three first cases are deduced directly from Definition 3.1.5 and Property 3.1.6. For point (iv), recall that for two gamma convolutions with Thorin measures ν_1 and ν_2 , the Thorin measure of the sum is:

$$\nu(\mathbf{s}) = \nu_1(\mathbf{s}) + \nu_2(\mathbf{s}).$$

Now, according to the reading of Property 3.1.6, point (iv), no set of $p < d$ rays S_{e_1}, \dots, S_{e_p} concentrate more than half of ν_1 weights and more than half of ν_2 weights. Hence, no set of $p < d$ rays can concentrate more than half of ν

weights. Therefore, the class of w.b. gamma convolutions is closed w.r.t finite convolutions. \square

An interesting fact is that any well-behaved gamma convolution has exponentially decreasing Laguerre coefficients, as the following property shows:

Property 3.1.8 (Exponential decay of Laguerre coefficients). *For any $\varepsilon, \varepsilon'$ s.t. $\varepsilon > \varepsilon' > 0$ and any dimension d , there exists a finite positive constant $B(d, \varepsilon')$ such that the Laguerre coefficients $(a_{\mathbf{k}})_{\mathbf{k} \in \mathbb{N}^d}$ of any d -variate ε -well-behaved gamma convolution verify:*

$$|a_{\mathbf{k}}| \leq B(d, \varepsilon')(1 + \varepsilon')^{-|\mathbf{k}|}.$$

Proof. The sketch of the proof is as follows: we construct a multivariate complex function R that has the Laguerre coefficients as its Taylor coefficients, and we show by a singularity analysis that this function is analytic around the origin, in a polydisc with polyradius $1 + \varepsilon$. We conclude by applying Cauchy's inequality. As the one-dimensional case allows for more detailed computations, we start by working for all d , and split later into the two cases $d = 1$ and $d > 1$.

We follow the path of [25] to express the Laguerre coefficients $(a_{\mathbf{k}})_{\mathbf{k} \in \mathbb{N}^d}$ as the Taylor coefficients of a certain (d -variate, complex) function R . From the computation of the moment generating function $M(\mathbf{t})$ of the Laguerre expanded density $f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \varphi_{\mathbf{k}}(\mathbf{x})$, we have:

$$\begin{aligned} M(\mathbf{t}) &= \int_{\mathbb{R}_+^d} e^{\langle \mathbf{t}, \mathbf{x} \rangle} f(\mathbf{x}) d\mathbf{x} = \sqrt{2}^d \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \prod_{i=1}^d \sum_{j=0}^{k_i} \binom{k_i}{j} \frac{(-2)^j}{j!} \int_{\mathbb{R}_+} e^{(t_i-1)x_i} x_i^j dx_i \\ &= \sqrt{2}^d \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \prod_{i=1}^d \sum_{j=0}^{k_i} \binom{k_i}{j} \frac{(-2)^j}{(1-t_i)^{j+1}} \\ &= \sqrt{2}^d (1-t)^{-1} \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \left(1 - \frac{2}{1-t}\right)^{\mathbf{k}} \\ &= \sqrt{2}^d (1-t)^{-1} \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \left(\frac{\mathbf{t}+1}{\mathbf{t}-1}\right)^{\mathbf{k}}, \end{aligned}$$

which implies that:

$$\sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \left(\frac{\mathbf{t}+1}{\mathbf{t}-1}\right)^{\mathbf{k}} = \sqrt{2}^{-d} (1-t)^1 M(\mathbf{t}).$$

Now denote $\mathbf{h}(\mathbf{t}) = (h(t))_{t \in \mathbf{t}}$, so that \mathbf{h} applies the Möbius transform $h(t) = \frac{t+1}{t-1}$ from Lemma 3.1.4 componentwise. Using the substitutions $\mathbf{y} = \mathbf{h}(\mathbf{t}) \iff \mathbf{t} =$

$\mathbf{h}(\mathbf{y})$ (since h is its own inverse), which implies that $(1 - \mathbf{t})^{\mathbf{1}} = 2^d (1 - \mathbf{y})^{-\mathbf{1}}$, we define the function R as:

$$\begin{aligned} R(\mathbf{y}) &:= \sum_{\mathbf{k} \in \mathbb{N}^d} a_{\mathbf{k}} \mathbf{y}^{\mathbf{k}} = \sqrt{2}^{-d} M(\mathbf{h}(\mathbf{y})) (1 - \mathbf{h}(\mathbf{y}))^{\mathbf{1}} \\ &= \sqrt{2}^d M(\mathbf{h}(\mathbf{y})) (1 - \mathbf{y})^{-\mathbf{1}}. \end{aligned}$$

We now study the regularity of the function R , which is tightly related to the regularity of M . Recall first the expression of the moment generating function for our random vector:

$$M(\mathbf{t}) = \prod_{i=1}^n (1 - \langle \mathbf{s}_i, \mathbf{t} \rangle)^{-\alpha_i}$$

Denote by \mathcal{V}_R the singular variety of R , i.e., the set of points where R is not analytic. From the singularities of $\mathbf{y} \mapsto (1 - \mathbf{y})^{-\mathbf{1}}$ and from those of M , we have:

$$\mathcal{V}_R \subseteq \{\mathbf{y} : \exists j, y_j = 1\} \cup \bigcup_{i=1}^n \{\mathbf{y} : \langle \mathbf{s}_i, \mathbf{h}(\mathbf{y}) \rangle = 1\}.$$

Consider first the fact that, since the model is ε -w.b., $|\alpha| > 1$. Hence, $M(\mathbf{h}(\mathbf{y}))$ dominates $(1 - \mathbf{y})^{-\mathbf{1}}$ when \mathbf{y} goes to $\mathbf{1}$ (as $\mathbf{h}(\mathbf{y})$ goes to infinity) and therefore R is analytic at $\mathbf{1}$. We therefore only need to consider the singularities of $M(\mathbf{h}(\mathbf{y}))$, and we can further restrict \mathcal{V}_R :

$$\mathcal{V}_R \subseteq \bigcup_{i=1}^n \{\mathbf{y} : \langle \mathbf{s}_i, \mathbf{h}(\mathbf{y}) \rangle = 1\}.$$

We now split the argument into the cases $d = 1$ and $d > 1$, to show that the ε -w.b. condition is equivalent to the analyticity of R on the polydisc with polyradius $1 + \varepsilon$.

When $d = 1$, all singularities of R are such that $sh(y) = 1$ for a scale $s \in \mathbf{s}$. In other words, the singularities are at the points $y = h(\frac{1}{s})$ for each scale $s \in \mathbf{s}$. Recall that $-h(x) = h(x^{-1})$, and that h is (strictly) decreasing before and after its vertical asymptote at 1. Three cases arise:

- If $s \in]0, 1[$, then $|y| = |h(\frac{1}{s})| = h(\frac{1}{s}) > 1 + \varepsilon \iff \frac{1}{s} < h(1 + \varepsilon) \iff s > h(1 + \varepsilon)^{-1} = \frac{\varepsilon}{2 + \varepsilon}$.
- If $s \in]1, +\infty[$, then $|y| = |h(\frac{1}{s})| = h(s) > 1 + \varepsilon \iff s < h(1 + \varepsilon) = \frac{2 + \varepsilon}{\varepsilon}$.
- If $s = 1$, $h(\frac{1}{s}) = \pm\infty$ and there is no singularity.

Hence,

$$\min_{y \in \mathcal{V}_R} |y| > 1 + \varepsilon \iff \frac{\varepsilon}{2 + \varepsilon} < \mathbf{s} < \frac{2 + \varepsilon}{\varepsilon},$$

which corresponds (with $|\alpha| > 1$) to the ε -w.b. condition.

We now turn ourselves to the multivariate case. When $d > 1$, the singularities of $\mathbf{t} \mapsto M(\mathbf{t})$ are the zeros of the function:

$$G : \mathbf{t} \mapsto \prod_{i=1}^n (1 - \langle \mathbf{s}_i, \mathbf{t} \rangle)^{\alpha_i},$$

By the definition of a ε -w.b model, for any \mathbf{t}^* a zero of G , we have that:

$$\exists i \in 1, \dots, d, |h(t_i^*)| > 1 + \varepsilon.$$

Hence, the singularities of R , which are images by \mathbf{h} of the zeros of G , must have at least one dimension that is outside the centered polydisc with radius $1 + \varepsilon$. Hence, when $d > 1$, R is also analytic on the centered polydisc with polyradius $1 + \varepsilon$.

Now, for any d , by Cauchy's inequality, see Theorem 2.2.7 in [22], taking a ε' strictly between 0 and ε we have the wanted finite upper bound on the coefficients:

$$B(d, \varepsilon') = \sup_{\substack{(\alpha, \mathbf{s}) : \mathcal{G}_{d,n}(\alpha, \mathbf{s}) \text{ is } \varepsilon\text{-w.b.} \\ \forall i, |y_i| < 1 + \varepsilon'}} R(y) < \infty.$$

□

Since we can now compute efficiently the Laguerre coefficients of densities in $\mathcal{G}_{d,n}$, with an insurance that they would decrease fast if the model is well-behaved, we propose in the next subsection to finally discuss an estimation algorithm based on an Integrated Square Error loss that leverages these coefficients.

3.2. An incomplete but however consistent empirical loss

Expressing the density of $\mathcal{G}_{d,n}$ random vectors into the basis from Definition 3.1.1 allows us to truncate the basis and effectively compute an approximated density. From this, by evaluating empirical Laguerre coefficients from data, we will fit the parameters to data by minimizing the L_2 distance between coefficients (since the basis is orthonormal). For parameters α, \mathbf{s} , denote by $f_{\alpha, \mathbf{s}}$ the density of the $\mathcal{G}_{d,n}(\alpha, \mathbf{s})$ distribution, and f the density underlying the observations $\mathbf{x}_1, \dots, \mathbf{x}_N$.

We would like to minimize the integrated square error:

$$\|f - f_{\alpha, \mathbf{s}}\|_2^2 = \sum_{\mathbf{k} \in \mathbb{N}^d} (a_{\mathbf{k}}(f) - a_{\mathbf{k}}(\alpha, \mathbf{s}))^2.$$

Estimating $a_{\mathbf{k}}(f) = \int f(\mathbf{x})\varphi_{\mathbf{k}}(\mathbf{x}) d\mathbf{x}$ by a simple Monte-Carlo plug-in

$$\widehat{a}_{\mathbf{k}} = \frac{1}{N} \sum_{i=1}^N \varphi_{\mathbf{k}}(\mathbf{x}_i),$$

we could compute an approximation of this loss. However, we cannot compute all couples of coefficients $\widehat{a}_{\mathbf{k}}$ and $a_{\mathbf{k}}(\boldsymbol{\alpha}, \mathbf{s})$ for all $\mathbf{k} \in \mathbb{N}^d$, so that we are forced to truncate the basis.

Definition 3.2.1 ($\mathcal{G}_{d,n}$ parameters estimator). *For a dataset $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$, we define parameters estimator of a well-behaved $\mathcal{G}_{d,n}$ distribution that matches the observations as:*

$$(\widehat{\boldsymbol{\alpha}}, \widehat{\mathbf{s}}) = \arg \min_{\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s}) \text{ w.b.}} L_{\mathbf{m}}(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{s}),$$

where the approximated truncated loss $L_{\mathbf{m}}$ is given by:

$$L_{\mathbf{m}}(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{s}) = \sum_{\mathbf{k} \leq \mathbf{m}} (\widehat{a}_{\mathbf{k}} - a_{\mathbf{k}}(\boldsymbol{\alpha}, \mathbf{s}))^2,$$

and where the empirical Laguerre coefficients are given by:

$$\widehat{a}_{\mathbf{k}} = \frac{1}{N} \sum_{i=1}^N \varphi_{\mathbf{k}}(\mathbf{x}_i).$$

The loss $L_{\mathbf{m}}(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{s})$ can be efficiently computed through Algorithm 2. It will be 0 if the first coefficients of the Laguerre expansion of f and of our estimator match perfectly, i.e., assuming $f \in \mathcal{G}_{d,n}$ and \mathbf{m} big enough, if we found the right Thorin measure. Note that if the goal was given theoretically, like in a projection from \mathcal{G}_d to $\mathcal{G}_{d,n}$, we could compute $\widehat{a}_{\mathbf{k}}$ by formal integration with high precision instead of Monte-Carlo. We will discuss this option in the next section. However, even in this case, the error that comes from the truncation of the basis cannot be computed.

To show the consistency of this loss, we use in the following the results from Property 3.1.8.

Property 3.2.2 (Consistency). *Consider that $f \sim \mathcal{G}_{d,n}(\boldsymbol{\alpha}_0, \mathbf{s}_0)$ is well-behaved, and denote \mathbf{x} a dataset generated from this density with N observations. Fix $\mathbf{m} \in \mathbb{N}^d$. Suppose that we have a global minimizer*

$$(\widehat{\boldsymbol{\alpha}}, \widehat{\mathbf{s}}) = \arg \min_{\mathcal{G}_{d,n}(\boldsymbol{\alpha}, \mathbf{s}) \text{ w.b.}} L_{\mathbf{m}}(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{s}),$$

which depends on the threshold \mathbf{m} and on the N observations \mathbf{x} . Then:

$$\|f - f_{\widehat{\boldsymbol{\alpha}}, \widehat{\mathbf{s}}}\|_2^2 \xrightarrow[\mathbf{m} \rightarrow \infty]{\substack{a.s \\ N \rightarrow \infty}} 0.$$

Proof. To show the result, we start by expressing the loss in the Laguerre basis, and we split the basis on indices smaller and greater than \mathbf{m} . For the sake of simplicity, we will omit f in $a_{\mathbf{k}}(f)$ and denote $a_{\mathbf{k}} = a_{\mathbf{k}}(f)$.

$$\begin{aligned} \|f - f_{\hat{\alpha}, \hat{\mathbf{s}}}\|_2^2 &= \sum_{\mathbf{k} \in \mathbb{N}^d} (a_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}}))^2 \\ &= \underbrace{\sum_{\mathbf{k} \leq \mathbf{m}} (a_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}}))^2}_A + \underbrace{\sum_{\mathbf{k} > \mathbf{m}} (a_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}}))^2}_B \end{aligned}$$

We first show that $A \xrightarrow[N \rightarrow \infty]{a.s.} 0$ for any fixed \mathbf{m} . For A , we can treat each term in the summation independently as the sum is finite. A can be then expanded further by plugging-in the Monte-Carlo estimator $\widehat{a}_{\mathbf{k}} = \frac{1}{N} \sum_{i=1}^N \varphi_{\mathbf{k}}(\mathbf{x}_i)$ and expanding the square:

$$A = \underbrace{\sum_{\mathbf{k} \leq \mathbf{m}} (a_{\mathbf{k}} - \widehat{a}_{\mathbf{k}})^2}_{A_1} + 2 \underbrace{\sum_{\mathbf{k} \leq \mathbf{m}} (a_{\mathbf{k}} - \widehat{a}_{\mathbf{k}}) (\widehat{a}_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}}))}_{A_2} + \underbrace{\sum_{\mathbf{k} \leq \mathbf{m}} (\widehat{a}_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}}))^2}_{L_{\mathbf{m}}(\mathbf{x}, \hat{\alpha}, \hat{\mathbf{s}})}.$$

Now, since $\widehat{a}_{\mathbf{k}} \xrightarrow[N \rightarrow \infty]{a.s.} a_{\mathbf{k}}$ by the strong law of large numbers, $A_1 \xrightarrow[N \rightarrow \infty]{a.s.} 0$.

Furthermore, since we restricted the optimization to well-behaved estimators, each $a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}})$ is bounded by Property 3.1.8. Now $|\widehat{a}_{\mathbf{k}}| \leq \sqrt{2^d}$ because $|\varphi_{\mathbf{k}}(\mathbf{x})| \leq \sqrt{2^d}$ whatever $\mathbf{x} \in \mathbb{R}_+^d$. Therefore, for any \mathbf{k} , $(\widehat{a}_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}}))$ is bounded and $(a_{\mathbf{k}} - \widehat{a}_{\mathbf{k}}) \xrightarrow[N \rightarrow \infty]{} 0$, which makes $A_2 \xrightarrow[N \rightarrow \infty]{a.s.} 0$.

Last but not least, by definition of $(\hat{\alpha}, \hat{\mathbf{s}})$ as global minimizers,

$$L_{\mathbf{m}}(\mathbf{x}, \hat{\alpha}, \hat{\mathbf{s}}) \leq L_{\mathbf{m}}(\mathbf{x}, \alpha_0, \mathbf{s}_0) \xrightarrow[N \rightarrow \infty]{a.s.} 0,$$

since Monte-Carlo estimators are not biased, and (α_0, \mathbf{s}_0) is inside the range of optimizations. Therefore, we can conclude that $A \xrightarrow[N \rightarrow \infty]{a.s.} 0$ for any \mathbf{m} fixed.

Now consider the remainder B . We show that $B \xrightarrow[\mathbf{m} \rightarrow \infty]{a.s.} 0$ uniformly in N . Since both the true model and the approximation are well-behaved, there exists $\varepsilon > \varepsilon' > 0$ so that they are both ε -w.b. By Property 3.1.8 we have then,

$$\begin{aligned} |a_{\mathbf{k}}| &\leq B(d, \varepsilon')(1 + \varepsilon')^{-|\mathbf{k}|} \\ \text{and } |a_{\mathbf{k}}(\hat{\alpha}, \hat{\mathbf{s}})| &\leq B(d, \varepsilon')(1 + \varepsilon')^{-|\mathbf{k}|}. \end{aligned}$$

Therefore:

$$\begin{aligned}
 B &= \sum_{\mathbf{k} > \mathbf{m}} (a_{\mathbf{k}} - a_{\mathbf{k}}(\hat{\boldsymbol{\alpha}}, \hat{\mathbf{s}}))^2 \\
 &\leq 4B(d, \varepsilon')^2 \sum_{\mathbf{k} > \mathbf{m}} ((1 + \varepsilon')^2)^{-|\mathbf{k}|} \xrightarrow[m \rightarrow \infty]{a.s.} 0,
 \end{aligned}$$

uniformly in N , which concludes the argument. □

Although the consistency result restricts the parameter space to only well-behaved Thorin measures, in practice we run our optimization procedures without this restriction, and simply check at the end that the resulting estimators are well-behaved. When $d = 1$, we can always fix ε to make the estimated gamma convolution ε -w.b. if $|\hat{\boldsymbol{\alpha}}| > 1$. In practice, numerical computations have troubles producing collinear scales, so every estimated multivariate gamma convolution is usually well-behaved as soon as the total mass of the produced Thorin measure is greater than 1. Recall that univariate distributions like log-Normal and Pareto have an infinite total mass for the Thorin measure.

We now have a reliable loss to minimize and an efficient algorithm to compute it, allowing us to estimate multivariate gamma convolutions from data. To test the approach, we propose in the next section to discuss some potential applications.

4. Investigation of performance

In this section, we show the results of our algorithm on several examples, both in univariate and multivariate settings.

Note that, through the recursive Faà di Bruno’s formula in Algorithm 2, we produce Laguerre coefficients as very high degree multinomials into $\boldsymbol{\alpha}, \mu_{\mathbf{0}}$ and $\boldsymbol{\xi}$ (the simplex projection of the scales). The loss will therefore have a myriad of local minima, making it highly non-convex and forcing us to use global minimization routines. We found the Particle Swarm [27, 31] global optimization routine to work particularly well on our test sets. All the implementation was gathered in the provided Julia package, along with the code corresponding to this investigation.

We start by the projection of known densities onto the $\mathcal{G}_{1,n}$ class, differentiating the cases when the known density is inside \mathcal{G}_1 or not. Then we discuss the estimation of $\mathcal{G}_{1,n}$ models from empirical data, with heavy tailed simulated data, and we finally look at the estimation of $\mathcal{G}_{d,n}$ models from multivariate empirical data, with a particular emphasis on the properties of the dependence structure underlying the data.

4.1. Projection from a known density

The two first examples are a log-Normal and a Weibull, that have the particularity of being respectively inside and outside \mathcal{G}_1 . We will also take a look at

the Pareto case, which is inside \mathcal{G}_1 , but might be outside $\mathcal{L}_2(\mathbb{R}_+)$.

Only one projection algorithm is available in the literature that projects \mathcal{G}_1 densities onto $\mathcal{G}_{1,n}$, finding a set of scales and shapes that matches the input distribution: the MFK procedure that we already discussed in Section 2.3.

We want to compare the two algorithm, MFK and our, on fair grounds. Therefore, we slightly modify our algorithm to use Laguerre coefficients based on the *theoretical* integrals $\mu_{k,-1}$, as MFK. The choice of the input distributions is also guided by [38, 18]. The experiment is as follows:

- We compute shifted moments $(\mu_{k,-1})_{k \leq 2n}$ of the distribution through direct integration of the density, at 1024 bits of precision (to match the 300 digits that MFK requires).
- We use MFK to compute an approximation in $\mathcal{G}_{1,n}$ based on these moments.
- Through Eq 3, we compute Laguerre coefficients based on the same shifted moments $\mu_{k,-1}$, and we minimize the loss given through Algorithm 2 to obtain our approximation.
- We compute Kolmogorov-Smirnov (one-sample, exact) distances between $N = 10\,000$ simulations from the estimated $\mathcal{G}_{1,n}$ model to the theoretical distribution, on $B = 100$ different simulations.

4.1.1. Projection from \mathcal{G}_1 densities to $\mathcal{G}_{1,n}$

The log-Normal distribution is a fundamental distribution in the \mathcal{G}_1 class. The proof that it belongs to \mathcal{G}_1 is actually what historically initiated the study of the class by Thorin. To match MFK's experiment, we conduct our first comparison with a $LN(\mu = 0, \sigma = 0.83)$ density. The resampled Kolmogorov-Smirnov distances are summarized, for several n and for both algorithms, in Figure 1.

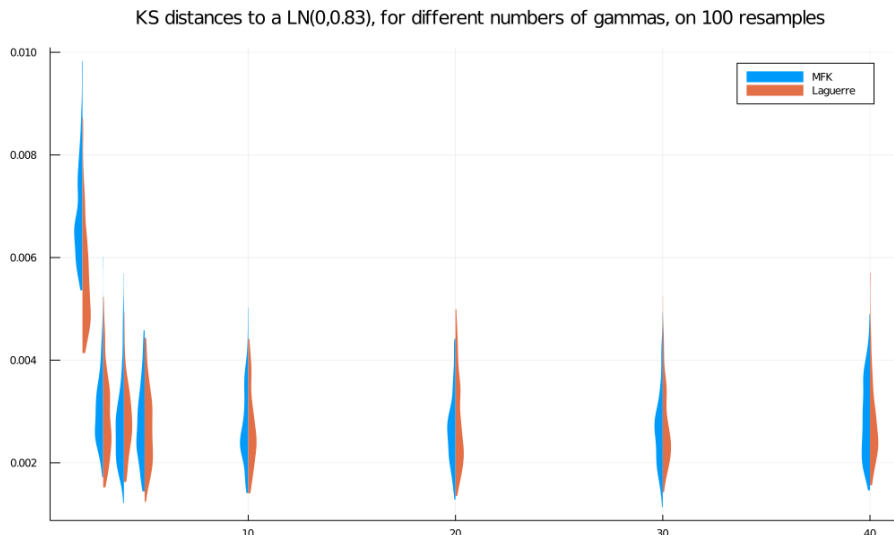


FIG 1. Violin densities of resampled KS distances (smaller is better). Blue: MFK projection, Orange: Laguerre Projection. In abscissa is $n \in \{2, 3, 4, 5, 10, 20, 30, 40\}$, the number of gammas used in the approximation.

Although MFK has a very strong convergence result, it seems that for $n = 2$ gammas, our estimator performs better with the same information (theoretical shifted moments given with 1024 bits precision), and that for $n > 2$, the performance is overall the same. Indeed, if we look at estimated parameters, we see that in most cases the two algorithms found approximations that are close to each other. We reported the shapes and scales that both algorithms produced for $n = 2, 3, 4$ and 5 in Table 1.

TABLE 1
Estimated parameters from both algorithms on the log-Normal example

	MFK		Laguerre	
	$\hat{\alpha}$	\hat{s}	$\hat{\alpha}$	\hat{s}
n=2	0.5688	1.5622	0.5458	1.6283
	2.4596	0.1941	2.4539	0.1999
n=3	0.2195	2.4902	0.2070	2.5781
	0.8934	0.6718	0.8919	0.6875
n=4	2.6081	0.0972	2.6071	0.0987
	0.0919	3.4076	0.0844	3.5307
n=5	0.4596	1.2208	0.4555	1.2513
	1.0042	0.3737	1.0063	0.3792
	2.6956	0.0588	2.6957	0.0594
	0.0398	4.3466	0.0346	4.5447
	0.2569	1.7963	0.2492	1.8576
	0.5708	0.7234	0.5721	0.7394
	1.0576	0.2399	1.0609	0.2428
	2.7574	0.0396	2.7582	0.0399

Although the produced estimations are close to each other on this example, our method has several clear benefits: the full density does not need to belong to \mathcal{G}_1 , and we do not require shifted moments up to 300 digits precision. In the next subsection, we describe the same experiment on a Weibull density that is known to make MFK simply fail, even with accurate enough computations.

4.1.2. Projection from densities outside \mathcal{G}_1 to $\mathcal{G}_{1,n}$

Let X be a Weibull r.v. with shape $k > 0$. Its pdf can then be written as:

$$f(x) = kx^{k-1}e^{-x^k} \mathbf{1}_{x \geq 0}.$$

When $k > 1$, this distribution does not belong to \mathcal{G}_1 . In [38, 18], authors report negative shapes for the MFK approximation in $\mathcal{G}_{1,2}$ of the Weibull with shape $k = \frac{3}{2}$, which is clearly a failure. This failure is occurring because the truncated moment problem that MFK is solving has no solution, as we detailed in Section 2.3, since the (normalized) cumulants ξ are not inside the moment cone.

Instead, the Laguerre basis is an orthonormal basis of $L_2(\mathbb{R}_+)$, which contains this density. The projection procedure detailed at the top of Subsection 4.1 still works correctly: it produces meaningful shapes and scales that are positive, whatever the number of gammas we ask for. Figure 2 displays a Violin plot of KS distances for several numbers of gammas:

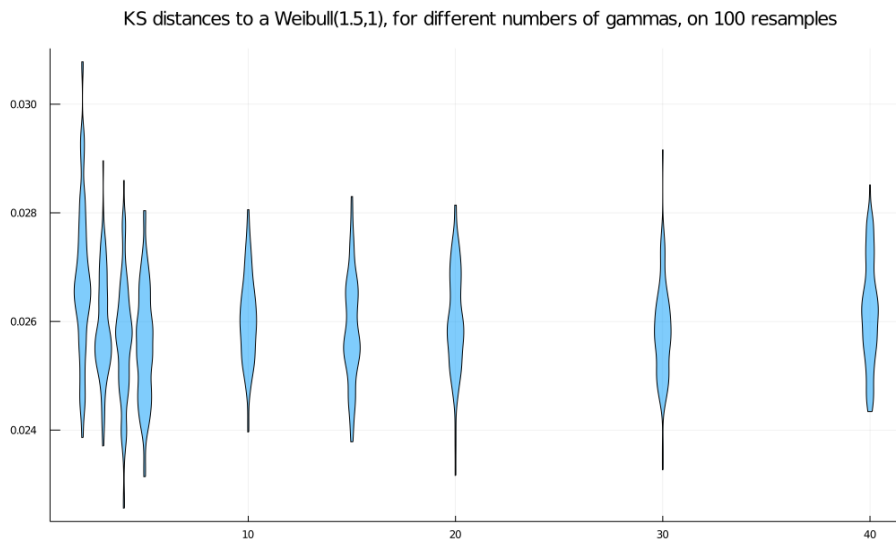


FIG 2. Violin densities of resampled KS distances for approximation of a Weibull($k = 1.5$). In abscissa is $n \in \{2, 3, 4, 5, 10, 20, 30, 40\}$, the number of gammas used in the approximation. Only the Laguerre approximation results are presented.

On Figure 2, we see that the KS distance is decreasing for an increasing number of gammas n , when n is small, but reaches a minimum quite quickly. Indeed, the distance is lower-bounded since the target is not in the class. Although none of the proposed approximations passes the KS test, we will revisit this example later from an estimation point of view to show that the projection is still an acceptable approximation of the input Weibull distribution.

We also note that the KS distance is not stable for increasing n : the algorithm has troubles to fix parameters to 0 if he does not need more precision. An additional penalty term in this direction might be a good solution to overcome this behavior.

Through these two first examples, we showed the projection possibilities of our algorithm, and compared them with MFK, the only existing method in the literature. We saw that for input densities inside the \mathcal{G}_1 class our method performs as good as MFK. Moreover, our method still works properly and produces meaningful results for examples outside the class, contrary to MFK which does not give any results. We will now talk about estimation from data, a problem that has never been solved for gamma convolutions, and showcase how our algorithm can find a good representation for a given dataset, whatever the dimension of the dataset.

In the following, we will treat empirical datasets given in standard IEE 745 float64 precision.

4.2. Estimation from univariate empirical data

Before coming back in more details to the Weibull case, we wanted to perform some tests with empirical data coming from simulations of two heavy-tailed distributions: log-Normal and Pareto.

4.2.1. Heavy-tailed examples

We consider first a dataset of 100'000 samples from a log-Normal distribution with the same parameter values as in Section 4.1, $LN(\mu = 0, \sigma = 0.83)$, in 64 bits precision.

From this sample, we compute empirical Laguerre coefficients, and then we optimize the set of shapes and scales of the estimated density to minimize the truncated L_2 loss between them and the ones produced through Algorithm 2. Following [18, 38], we choose the size of the basis as being $m = 2n + 1$, such that if we remove the first moment that is irrelevant, we have a number of moments equal to the number of parameters⁵. We ran the experiment for different number of gammas $n = 10, 20$ and 40. The corresponding results are respectively available in Figures 3, 4 and 5.

⁵Since Julia starts counting at 1 and not 0, we actually have m moments and not $m + 1$ as we did in the previous theoretical analysis.

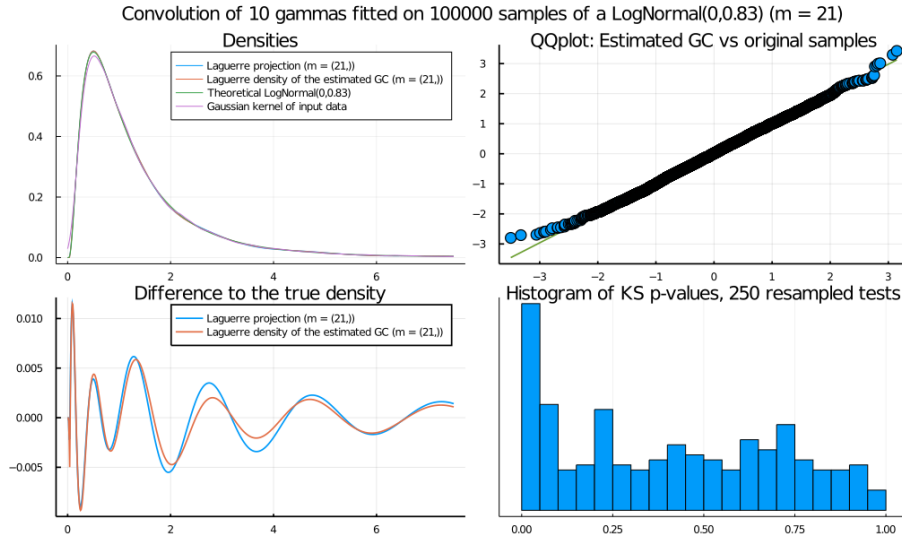


FIG 3. *Log-Normal results with 10 gammas: Upper left: the comparison of the density approximations. Lower left: the difference of to the true density. Upper right: a quantile-quantile plot of the approximation against the true distribution. Lower right: p-values of one-sample KS tests of simulations from the estimator against the true distribution.*

We see on these graphs that the approximation is good enough to fool the Kolmogorov-Smirnov test. The absolute pointwise error between the Laguerre approximation of the density of the gamma convolution and the true log-Normal density is lower than 0.005 for $n = 10, 20$, and twice smaller for $n = 40$. A positive but surprising result is that, although the distance between the Laguerre approximation and the true density does not reduce with increasing m , the distance between the gamma approximation and the true density does.

The KS test setup is as follows: we simulate 250 resamples of 100'000-sized dataset from the estimated density, and compare it with a one-sample exact KS test to the true density (and not to the samples that the estimator was computed from). For all values of n , we observe close-to-uniform distribution of the p-values, which is a good result.

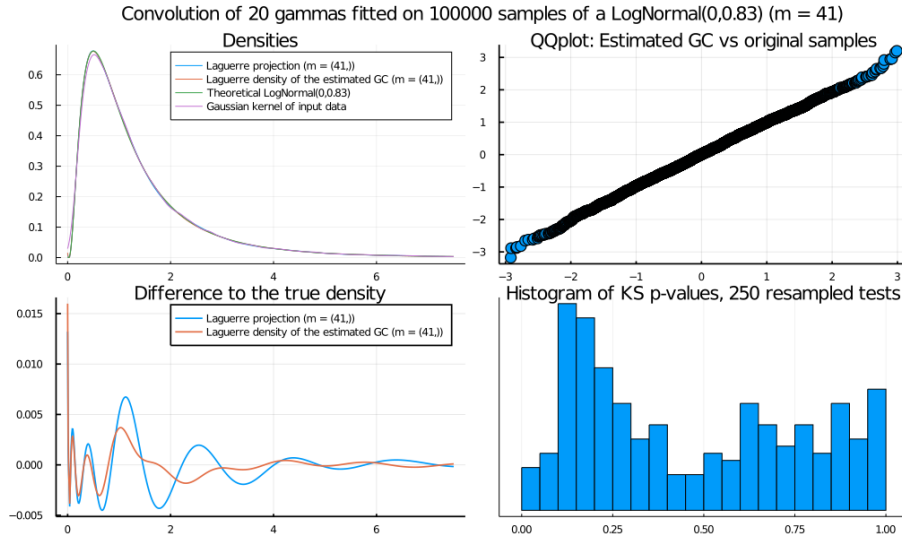


FIG 4. Log-Normal results with 20 gammas. Same legend as Figure 3.

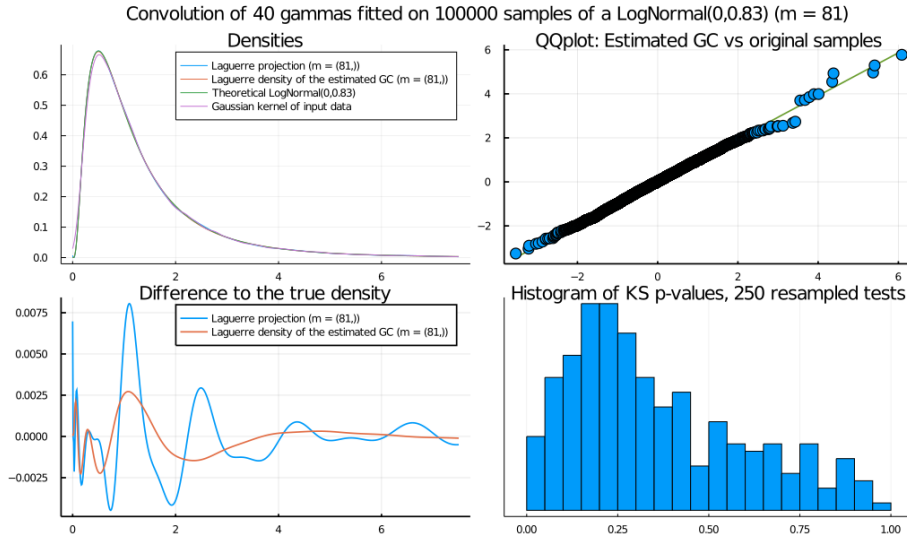


FIG 5. Log-Normal results with 40 gammas. Same legend as Figure 3.

We also treated several datasets simulated from Pareto’s distribution, with exactly the same experiment. Pareto’s distribution has a shape parameter $k > 0$ that influences the thickness of the tail: the distribution has a variance if and

only if $k > 2$, an expectation if and only if $k > 1$ and the density belongs to L_2 if and only if $k > 1/2$. Note that these inequalities are strict.

The fact that the distribution has no variance or no expectation is not really a problem for our procedure, since we use only shifted moments that all Pareto distributions have, but our convergence result requires a square-integrable density. Therefore, we propose to check several values of the shape parameter k , covering all cases, while varying the number n of gammas. Figure 6 summarizes the results by only showing quantile-quantile plots of the estimated distributions against the starting Pareto distribution. All the models follow the same procedure as for the previous log-Normal estimators.




FIG 6. Quantile-quantile plots for Pareto experiments, with $N = 1000$ samples (log-scaled). We only show the 995 first points: 5 points are excluded in the tail for clarity. Each row corresponds to a shape parameter for the Pareto, and each column to a number of gammas.

Note that even quantile-quantile plots between two simulations from the same Pareto distribution have a tendency to deviate in the extremes (due to the fat tail), which is why we excluded the last points. In each case, we also ran the same Kolmogorov-Smirnov testing procedure as for the log-Normal, and obtained very good uniformity of the p-values, except for $k = 0.25$ when n is too small. We see on the graph that increasing the number of gammas usually produces better results, but also that the heavier the tail (smaller shapes) the harder it is to get a good fit. Note that the KS distance is not the best distance to take into account the tails, which could explain why the distances from Figure 1 did not decrease for high values of n .

We see on these graphs that the method is accurate enough for reasonably high

quantiles, and very good in the body of the distribution. The $k = 0.25$ case is not in the theoretical boundaries for the convergence result to take place, and from sample to sample we have to wait until at least 20 gammas to have an acceptable fit.

4.2.2. An example outside the class

We revisit the Weibull($k = 1.5$) case of Section 4.1, a distribution that is not inside the class.

We ran our algorithm on 100'000 samples from a Weibull with shape $k = \frac{3}{2}$, with successive numbers of gammas of 2, 10, 20 and 40. The corresponding graphs are shown in Figure 7. Here, we show only the difference on the density estimation and the quantile-quantile plot as one notable thing is that the Laguerre expansion of the density gets better and better as we increase the size of the basis (jointly with the number of gammas), but the estimator does not: the difference to the density is always the same, and always makes Kolmogorov-Smirnov reject the proposed estimator, although quantile-quantile plots are all OK.

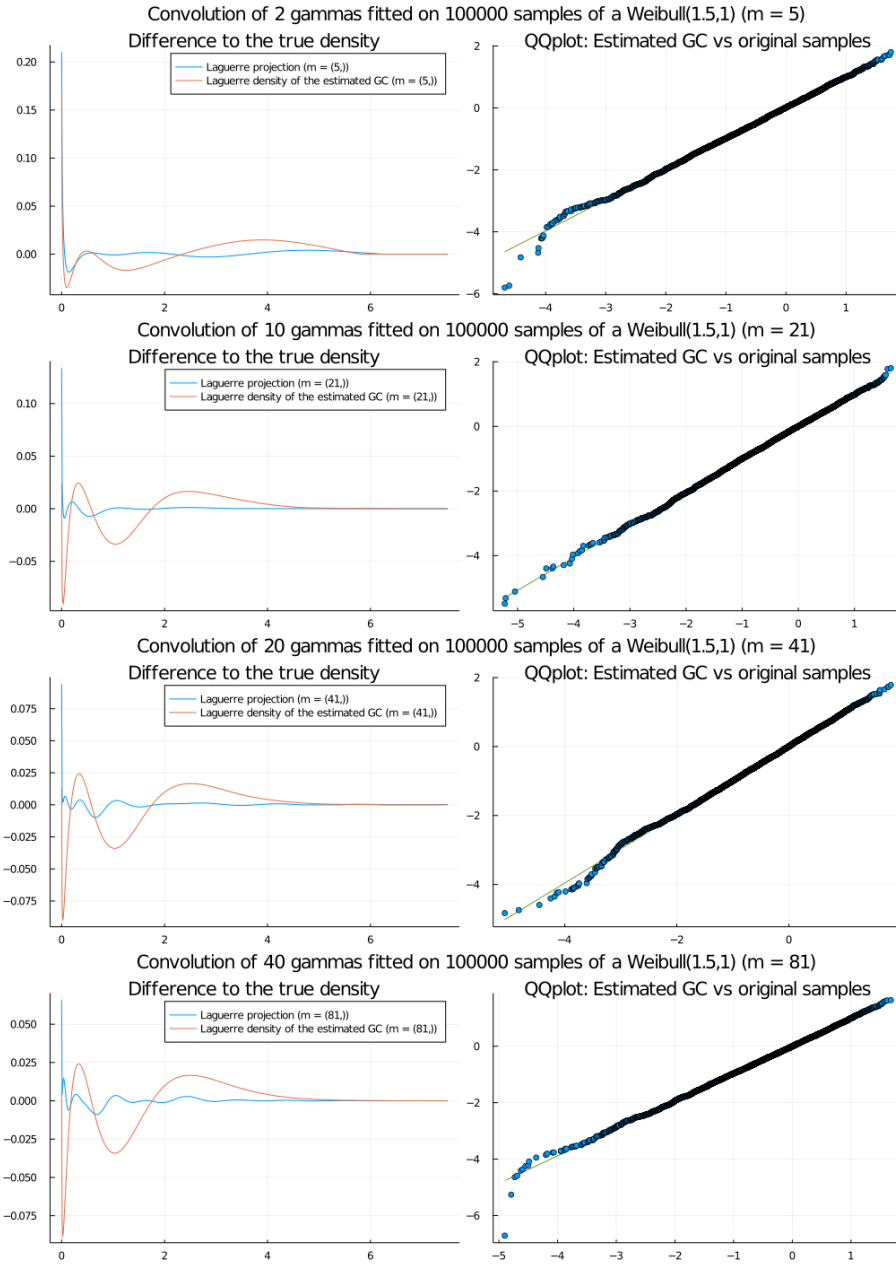


FIG 7. Weibull results with 2, 10, 20 and 40 gammas, respectively on each row. Same legend as Figure 3.

On these estimations, several things should be noted:

- The KS tests reject the hypothesis that the distributions are the same whatever the number of gammas, which was expected.
- The pointwise difference on the density of the estimator does not match the error of the Laguerre expansion (which is getting smaller as the basis grows) but stays the same whatever the size of the basis and the number of gammas. This error represents some kind of distance between the Weibull distribution and the class.
- The quantile-quantile plot is on the other hand getting better and better as we add more gammas.

These three points show that our procedure works even for data that are far from the \mathcal{G}_1 class, finding some orthogonal projection, relative to the basis, onto the class. This is a very good point for the procedure.

In the next examples, we will showcase some multivariate uses of the procedure, with a particular emphasis on the dependence structure of the simulated data and of the produced estimators.

4.3. Estimation from multivariate empirical data

Now that we extensively discussed the univariate case, we turn ourselves to the multivariate capabilities of our algorithm. As no known algorithm estimates multivariate gamma convolutions, we are not able to compare our results to other procedures. However, we can still check quantile-quantile plots for marginals and verify that the shape of the dependence structure matches the inputted one.

As use-cases of our models might vary, we propose two different examples: a multivariate log-Normal, with a Gaussian copula that induces no tail dependency, and a (survival) Clayton copula, including a strong tail dependency. Note that we reduced the exposition to estimation from empirical data, but formal integration to obtain $(\mu_{k,-1})_{k \leq m}$ from a formal density at any desired computational precision is also a perfectly fine use case. However, we do not have a clever way of choosing the parameter \mathbf{m} yet. We found that setting all m_i 's equal to n provide enough evaluation points to obtain a good fit on our simple examples.

4.3.1. A Gaussian example

We do not know if the d -variate log-Normal distribution, defined as the componentwise exponentiation of a d -variate Gaussian random vector, is inside \mathcal{G}_d . However, we can still run our procedure and observe the resulting approximations.

A first result is given in Figure 8, that corresponds to a $\mathcal{G}_{2,10}$ estimation taken from 100'000 data points simulated from a $MLN(\mu = \mathbf{0}, \sigma_1 = \sigma_2 = 1, \rho = 0.5)$ with two standard log-Normal margins and a Gaussian copula with $\rho =$

0.5. Figure 8 proposes quantile-quantile plots for the marginals, and a simple Gaussian kernel density estimator for the pseudo-observations on the original and estimated models.

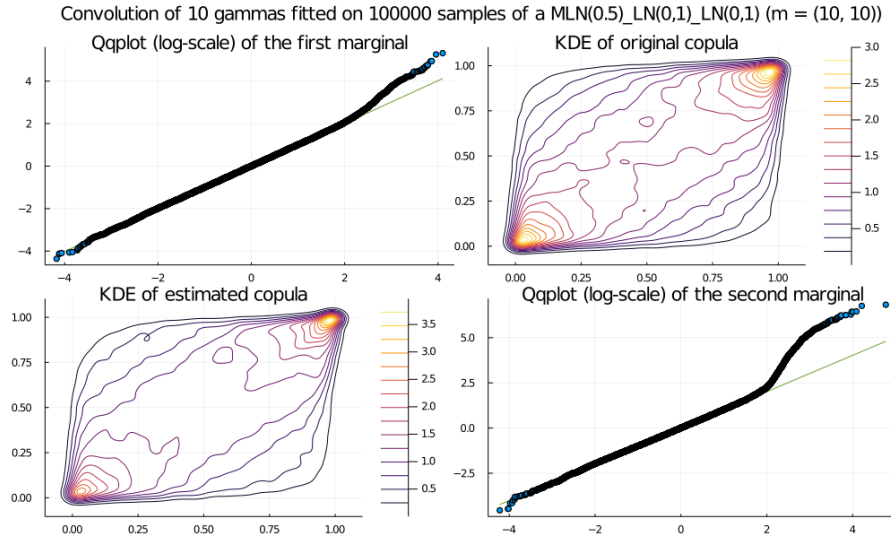


FIG 8. Multivariate log-Normal results with 10 gammas. Top left and bottom right: quantile-quantile plots for, respectively, the first and second marginals. Bottom left and top right: levels plots of pseudo-observations from, respectively, simulations from the estimated model and the dataset.

We see on Figure 8 that the dependence structure was correctly reproduced by the approximation as a convolution of 10 (comonotonous) gamma random vectors, but marginal tails are not estimated correctly. Thankfully, pushing the number of multivariate gammas to $n = 20$ solves this issue, as shown in Figure 9:

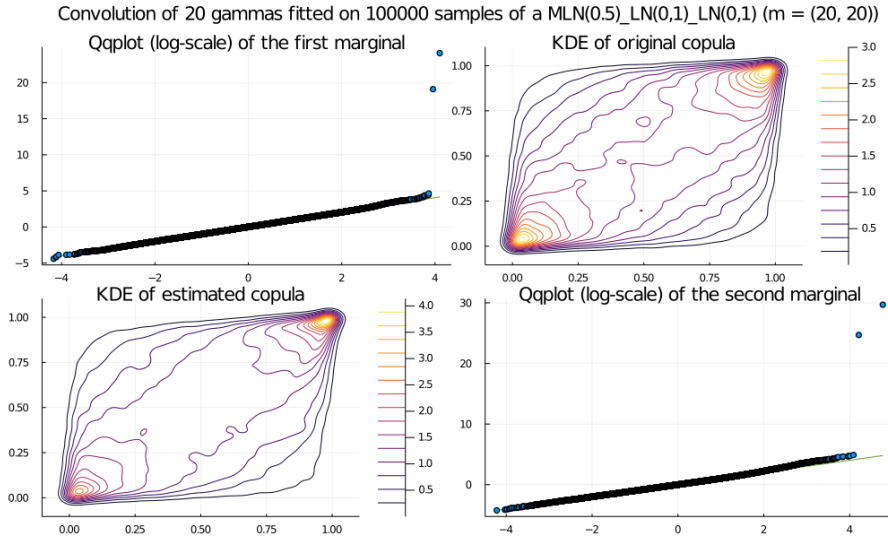


FIG 9. Multivariate log-Normal results with 20 gammas. Same legend as Figure 8.

On Figure 9, we see that the dependence structure is still very good, and that the problem of the marginal tails has been solved by adding a few more gammas. Note that since we have 100'000 points in the quantile-quantile plots, one point out of the line corresponds to a deviation at the 0.00001-quantiles, which are of little importance in many use cases.

Overall, marginal quantile-quantile plots are correct when the number of gammas is big enough, and the Gaussian dependence structure is well fitted. The Gaussian copula is a copula that exhibits no tail dependency. We already know that our model can exhibit asymmetric behaviors, but the possibility of tail dependency property is also an important feature to have. Therefore, in the next example we run the model on datasets that have such a feature.

4.3.2. A copula with tail dependency

To convince ourselves that tail dependency is a property that our estimator can exhibit, we will try to fit data whose dependence structure is given through a (survival) Clayton copula. This copula is a symmetric copula that exhibits upper tail dependency. We took here a parameter $\theta = 7$, yielding a very high upper tail dependence.

In Figure 10, you can see the results corresponding to data sampled from a Clayton copula, a Pareto($k = 2.5$) marginal and a $LN(0, 0.83)$ marginal. The conditions of the experiments are identical to the previous subsection ones.

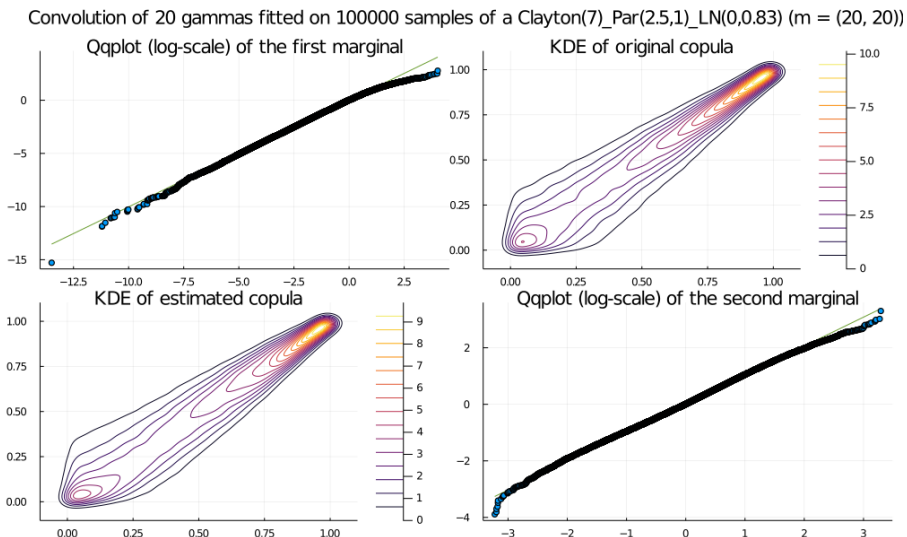


FIG 10. Results with 20 gammas. Same legend as Figure 8.

We note that marginal upper tails are not perfectly fitted: this could probably be overcome by taking a greater number of gammas. However, we see that the shape of the estimated copula is not perfectly symmetric: there is no constraint on the model to produce a symmetric dependence structure, and therefore producing one is hard. The small misfit around the (0.25, 0.02) quantile is probably due to this: marginals are not on the same scale, and therefore the Laguerre coefficients with $k_1 \gg k_2$ are probably on a different scale than those with $k_2 \gg k_1$. The goodness of fit is therefore not symmetric.

5. Conclusion

Although the class of gamma convolutions has very appealing properties for the practitioner, the estimation of gamma convolutions from empirical data had not been considered in the literature. The closest algorithm we could find is the recent MFK algorithm, that we described in Subsection 2.3, which projects densities from the \mathcal{G}_1 class to the $\mathcal{G}_{1,n}$ class. Unfortunately, extending this algorithm for general estimation purposes is hard if not impossible.

Using a well-chosen Laguerre basis, we constructed a series expansion for the density of a distribution in $\mathcal{G}_{d,n}$. Coupled with a simple L_2 loss for the density, we were able to project densities from \mathcal{G}_d , $d > 1$ onto $\mathcal{G}_{d,n}$, which was not possible, but also to project any density onto $\mathcal{G}_{d,n}$, which was not possible either. Finally, through the same loss, we can estimate $\mathcal{G}_{d,n}$ models for any reasonable d, n , which was also not treated in the literature.

Although the convergence result is not as strong as MFK's one, the precision needs and computational speeds are far better. We saw that given the same information, shifted moments of the random variable with high precision, the performance of the two algorithms is overall the same, except maybe for small numbers of gammas where we beat MFK. For both algorithms, Figure 1 also shows that the model's goodness of fit seems to be decreasing for large numbers of gammas, and an adaptive estimator might be a suitable way out.

Maybe a better loss could be found for the estimation of these gamma convolutions. Furthermore, automatic penalization of the model could be done through the number of gammas, but also through sparsity of the Thorin measure, according to Property 2.1.4, yielding more independent factors when possible. Such a modification will also provide a clearer view at the additive risk factor structure.

Last but not least, many features are still to be discovered about multivariate Thorin classes. The estimation of multiplicative structures could yield interesting results. Bondesson showed that the product of random variables in \mathcal{G}_1 is also in \mathcal{G}_1 , but we still do not have a corresponding result for the multivariate case.

References

- [1] ARNOLD, D. N. and ROGNESS, J. P. (2008). Möbius Transformations Revealed. *Notices of the American Mathematical Society* **55** 1226–1231.
- [2] BENHADDOU, R., PENSKY, M. and RAJAPAKSHAGE, R. (2019). Anisotropic Functional Laplace Deconvolution. *Journal of Statistical Planning and Inference* **199** 271–285.
- [3] BERNARDINI, A., NATALINI, P. and RICCI, P. E. (2005). Multidimensional Bell Polynomials of Higher Order. *Computers & Mathematics with Applications* **50** 1697–1708.
- [4] BONDESSON, L. (1982). On Simulation from Infinitely Divisible Distributions. *Advances in Applied Probability* **14** 855–869.
- [5] BONDESSON, L. (1992). *Generalized Gamma Convolutions and Related Classes of Distributions and Densities. Lecture Notes in Statistics* **76**. Springer New York, New York, NY.
- [6] BONDESSON, L. (2009). On Univariate and Bivariate Generalized Gamma Convolutions. *Journal of Statistical Planning and Inference* **139** 3759–3765.
- [7] BONDESSON, L. (2015). A Class of Probability Distributions That Is Closed with Respect to Addition as Well as Multiplication of Independent Random Variables. *Journal of Theoretical Probability* **28** 1063–1081.
- [8] CHERIAN, K. (1941). A Bivariate Correlated Gamma-Type Distribution Function. *Journal of the Indian Mathematical Society* **5** 133–144.
- [9] COMTE, F. and GENON-CATALOT, V. (2015). Adaptive Laguerre Density Estimation for Mixed Poisson Models. *Electronic Journal of Statistics* **9** 1113–1149.

- [10] CONSTANTINE, G. M. and SAVITS, T. H. (1996). A Multivariate Faa Di Bruno Formula with Applications. *Transactions of the American Mathematical Society* **348** 503–520.
- [11] COSSETTE, H., LANDRIault, D., MARCEAU, E. and MOUTANABBIR, K. (2016). Moment-Based Approximation with Mixed Erlang Distributions. *Variance* **10** 161–182.
- [12] CVIJOVIĆ, D. (2011). New Identities for the Partial Bell Polynomials. *Applied Mathematics Letters* **24** 1544–1547.
- [13] DEREVYAGIN, M. S. and DERKACH, V. A. (2007). On the Convergence of Padé Approximations for Generalized Nevanlinna Functions. *Transactions of the Moscow Mathematical Society* **68** 119–162.
- [14] DI DIO, P. J. and SCHMÜDGEN, K. (2018). The Multidimensional Truncated Moment Problem: The Moment Cone. *arXiv preprint arXiv:1809.00584*.
- [15] DI DIO, P. J. and SCHMÜDGEN, K. (2018). The Multidimensional Truncated Moment Problem: Atoms, Determinacy, and Core Variety. *Journal of Functional Analysis* **274** 3124–3148.
- [16] DUSSAP, F. (2020). Anisotropic Multivariate Deconvolution Using Projection on the Laguerre Basis.
- [17] ESCHER, F. (1932). On the Probability Function in the Collective Theory of Risk. *Skand. Aktuarie Tidskr.* **15** 175–195.
- [18] FURMAN, E., HACKMANN, D. and KUZNETSOV, A. (2017). On Log-Normal Convolutions: An Analytical-Numerical Method With Applications to Economic Capital Determination. *SSRN Electronic Journal*.
- [19] HARDY, M. (2006). Combinatorics of Partial Derivatives. *The Electronic Journal of Combinatorics* **13** R1.
- [20] HELTON, J. W. and NIE, J. (2012). A Semidefinite Approach for Truncated K-Moment Problems. *arXiv:1105.0410 [math]*.
- [21] HENRION, D. and MALICK, J. (2012). Projection Methods in Conic Optimization. In *Handbook on Semidefinite, Conic and Polynomial Optimization*, (M. F. Anjos and J. B. Lasserre, eds.) **166** 565–600. Springer US, Boston, MA.
- [22] HORMANDER, L. (1973). *An Introduction to Complex Analysis in Several Variables*. Elsevier.
- [23] HU, C., POZDNYAKOV, V. and YAN, J. (2020). Density and Distribution Evaluation for Convolution of Independent Gamma Variables. *Computational Statistics* **35** 327–342.
- [24] JAMES, L. F., ROYNETTE, B. and YOR, M. (2008). Generalized Gamma Convolutions, Dirichlet Means, Thorin Measures, with Explicit Examples. *Probability Surveys* **5** 346–415.
- [25] KABARDOV, M. M. and RYABOV, V. M. (2009). Acceleration of the Convergence of the Laguerre Series in the Problem of Inverting the Laplace Transform. *Computational Mathematics and Mathematical Physics* **49** 579–588.
- [26] KIM, T., KIM, D. and JANG, G.-W. (2019). On Central Complete and Incomplete Bell Polynomials I. *Symmetry* **11** 288.

- [27] KOYUNCU, H. and CEYLAN, R. (2019). A PSO Based Approach: Scout Particle Swarm Algorithm for Continuous Global Optimization Problems. *Journal of Computational Design and Engineering* **6** 129–142.
- [28] KRUTTO, A. (2019). Empirical Cumulant Function Based Parameter Estimation in Stable Laws. *Acta et Commentationes Universitatis Tartuensis de Mathematica* **22** 311–338.
- [29] LAVERNY, O. (2021). <https://www.github.com/lrnv/ThorinDistributions.jl>: ThorinDistributions.jl v0.1.
- [30] LEE, S. and LIN, X. S. (2012). Modeling Dependent Risks with Multivariate Erlang Mixtures. *Astin Bulletin* **42** 153–180.
- [31] ZONG-JUN LI (2016). Experiment of Standard Particle Swarm Optimization 2011 in Julia on the CEC2013 Benchmark Functions.
- [32] MABON, G. (2017). Adaptive Deconvolution on the Non-Negative Real Line: Adaptive Deconvolution on \mathbb{R}^+ . *Scandinavian Journal of Statistics* **44** 707–740.
- [33] MATHAI, A. M. (1982). Storage Capacity of a Dam with Gamma Type Inputs. *Annals of the Institute of Statistical Mathematics* **34** 591–597.
- [34] MEISTER, A. (2009). *Deconvolution Problems in Nonparametric Statistics. Lecture Notes in Statistics* **193**. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [35] MELMAN, V. S., SHABLYA, Y. V., KRUCHININ, D. V. and SHELUPANOV, A. A. (2018). Realization of a Method for Calculating Bell Polynomials Based on Compositae of Generating Functions. **10** 14.
- [36] MENDEL, J. M. (1990). *Maximum-Likelihood Deconvolution. Signal Processing and Digital Filtering*. Springer New York, New York, NY.
- [37] MIATTO, F. M. (2019). Recursive Multivariate Derivatives of $e^{f(X_1, \dots, X_n)}$ of Arbitrary Order. *arXiv:1911.11722 [cs, math]*.
- [38] MILES, J., FURMAN, E. and KUZNETSOV, A. (2019). Risk Aggregation: A General Approach via the Class of Generalized Gamma Convolutions. *Variance*.
- [39] MISHKOV, R. L. (2000). Generalization of the Formula of Faà Di Bruno for a Composite Function with a Vector Argument. *International Journal of Mathematics and Mathematical Sciences* **24** 481–491.
- [40] MOSCHOPOULOS, P. G. (1985). The Distribution of the Sum of Independent Gamma Random Variables. *Annals of the Institute of Statistical Mathematics* **37** 541–544.
- [41] MUSTAPHA, H. and DIMITRAKOPOULOS, R. (2010). Generalized Laguerre Expansions of Multivariate Probability Densities with Moments. *Computers & Mathematics with Applications* **60** 2178–2189.
- [42] NARDO, E. D., GUARINO, G. and SENATO, D. (2011). A New Algorithm for Computing the Multivariate Faà Di Bruno’s Formula. *Applied Mathematics and Computation* **217** 6286–6295.
- [43] NIE, J. (2012). The A-Truncated K-Moment Problem. *arXiv:1210.6930 [math]*.
- [44] NIELSEN, F. (2012). k -MLE: A Fast Algorithm for Learning Statistical Mixture Models. *2012 IEEE International Conference on Acoustics, Speech*

- and *Signal Processing (ICASSP)* 869–872.
- [45] PÉREZ-ABREU, V. and STELZER, R. (2012). A Class of Infinitely Divisible Multivariate and Matrix Gamma Distributions and Cone-Valued Generalised Gamma Convolutions. *arXiv:1201.1461 [math, stat]*.
 - [46] PÉREZ-ABREU, V. and STELZER, R. (2014). Infinitely Divisible Multivariate and Matrix Gamma Distributions. *Journal of Multivariate Analysis* **130** 155–175.
 - [47] RIDOUT, M. S. (2009). Generating Random Numbers from a Distribution Specified by Its Laplace Transform. *Statistics and Computing* **19** 439–450.
 - [48] SCHMUDGEN, K. (2017). *The Moment Problem. Graduate Texts in Mathematics* **277**. Springer International Publishing, Cham.
 - [49] SCHWANDER, O. and NIELSEN, F. (2013). Fast Learning of Gamma Mixture Models with K-MLE. In *Similarity-Based Pattern Recognition*, (D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, E. Hancock and M. Pelillo, eds.) **7953** 235–249. Springer Berlin Heidelberg, Berlin, Heidelberg.
 - [50] SMITH, K. D. (2020). A Tutorial on Multivariate k -Statistics and Their Computation. *arXiv:2005.08373 [math, stat]*.
 - [51] SMITH, P. J. (2020). A Recursive Formulation of the Old Problem of Obtaining Moments from Cumulants and Vice Versa. **3**.
 - [52] SPOTT, D. A. (1983). Estimating the Parameters of a Convolution by Maximum Likelihood. *Journal of the American Statistical Association* **78** 457–460.
 - [53] SUETIN, P., KOSTRIKIN, A. I. and MANIN, Y. I. (1989). *Linear algebra and geometry*. CRC Press.
 - [54] THORIN, O. (1977). On the Infinite Divisibility of the Pareto Distribution. *Scandinavian Actuarial Journal* **1977** 31–40.
 - [55] THORIN, O. (1977). On the Infinite Divisibility of the Lognormal Distribution. *Scandinavian Actuarial Journal* **1977** 121–148.
 - [56] TIJMS, H. C. (2003). *A First Course in Stochastic Models*. John Wiley and sons.
 - [57] WITHERS, C. (1994). Expressions for Moments and Cumulants in Terms of Each Other Using Bell Polynomials. In *Proceedings of the 1994 New Zealand Statistical Association Conference, Massey University* 226–229.
 - [58] ZHANG, Y., CHEN, Y. and LI, D. D.-U. (2016). Optimizing Laguerre Expansion Based Deconvolution Methods for Analysing Bi-Exponential Fluorescence Lifetime Images. *Optics Express* **24** 13894.