



Unsupervised mixture estimation via approximate maximum likelihood based on the Cramér - von Mises distance

Marco Bee

Department of Economics and Management, University of Trento, via Inama, 5, Trento, 38122, Italy



ARTICLE INFO

Article history:

Received 5 November 2022
 Received in revised form 5 April 2023
 Accepted 10 April 2023
 Available online 17 April 2023

Keywords:

Dynamic mixtures
 Simulation
 Classification
 Tail events
 Statistical distance

ABSTRACT

Mixture distributions with dynamic weights are an efficient way of modeling loss data characterized by heavy tails. However, maximum likelihood estimation of this family of models is difficult, mostly because of the need to evaluate numerically an intractable normalizing constant. In such a setup, simulation-based estimation methods are an appealing alternative. The approximate maximum likelihood estimation (AMLE) approach is employed. It is a general method that can be applied to mixtures with any component densities, as long as simulation is feasible. The focus is on the dynamic lognormal-generalized Pareto distribution, and the Cramér - von Mises distance is used to measure the discrepancy between observed and simulated samples. After deriving the theoretical properties of the estimators, a hybrid procedure is developed, where standard maximum likelihood is first employed to determine the bounds of the uniform priors required as input for AMLE. Simulation experiments and two real-data applications suggest that this approach yields a major improvement with respect to standard maximum likelihood estimation.

© 2023 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Estimating the right tail of a non-negative probability distribution is very important in many fields such as hydrology, economics and finance. At the same time, the problem may be difficult from the modeling point of view, since the data-generating processes of the tail and the body are often different, so that no single stochastic model guarantees a precise description of the entire distribution. On top of that, statistical inference for the tail is challenging, because samples of extreme observations are typically small.

A fundamental set of tools for estimating the tail is known as Extreme Value Theory (EVT), which focuses on the distribution of the largest observations, or of the excesses above a high threshold. However, EVT only fits the tail; if interest is in the whole distribution, it is advisable to resort to a spliced distribution combining two different models for the body and the tail. A possible drawback is the need of continuity and differentiability constraints that reduce the number of free parameters (Scollnik, 2007).

An even more flexible alternative is the dynamic model developed by Frigessi et al. (2002). This approach employs a non-negative distribution with a size distribution supported on $[0, \infty)$ for the body and the Generalized Pareto distribution (GPD) for the tail. Unlike classical finite mixtures, the density $f(x)$ of the resulting random variable X is based on a dynamic weight, which is a monotonically increasing function of x . In this way, the GPD is given more and more weight as we move

E-mail address: marco.bee@unitn.it.

farther into the right tail, and the tail behavior is governed by the GPD. If the weight function is continuous, the density f is continuous as well. Finally, this is an unsupervised method, because no threshold needs to be chosen or estimated.

Despite its interesting properties, to the best of our knowledge this model has not been used in practice. The likely reason is that maximum likelihood estimation (MLE) is rather complex. On one hand, the mixing weight and the component densities share all parameters, so that, as pointed out by Frigessi et al. (2002), the use of the EM algorithm is precluded. On the other hand, direct maximization of the likelihood is feasible, but the normalizing constant is given by an integral that cannot be computed explicitly, and whose numerical approximation is not trivial.

Given the difficulties related to MLE, simulation-based methods are an appealing alternative that avoids the evaluation of the normalizing constant, since there is no optimization of the likelihood function. Hence, the possible error caused by an inaccurate approximation of the normalizing constant is eliminated, at the price of the introduction of simulation error.

In this paper we explore the second strategy, and use the Approximate Maximum Likelihood Estimation (AMLE) method proposed by Rubio and Johansen (2013) and applied to similar setups by Bee et al. (2015), Bee et al. (2017a) and Tafakori et al. (2022). The approach is essentially a frequentist version of the family of Approximate Bayesian Computation (ABC) techniques, aimed at approximate estimation of the posterior distributions of model parameters; see, e.g., Sunnåker et al. (2013), Beaumont (2019), Drovandi and Frazier (2022). In the early stages of ABC, the crucial issue was the choice of the summary statistics used to assess the proximity of true and simulated data. This issue is shared by AMLE and is not easy to address in general, except for uncommon setups where sufficient statistics are available.

To overcome this problem, in recent years several papers have developed ABC versions that compare empirical distributions of observed and simulated data, thus bypassing the need of choosing summary statistics. This way of proceeding was first proposed by Park et al. (2016), who employed a non-parametric approach; a more comprehensive treatment was given by Bernton et al. (2019). See Drovandi and Frazier (2022) and the references therein for a thorough review. In the following, we will term these ways of proceeding *full-data approaches*, to emphasize the use of all the data, with no dimensionality reduction transformation. Among the various proposals in the literature, we base our implementation on the Cramér-von Mises (CvM) distance, which is a simple and effective measure.

AMLE and MLE are compared via simulation: our outcomes suggest that AMLE has smaller root-mean-squared-error (RMSE) in all setups considered in the experiments. Two empirical analyses confirm the suitability of the approach for modeling and estimating skewed and fat-tailed distributions. Furthermore, the estimated weight allows one to identify the approximate number of observations generated by the heavy-tailed component.

On the theoretical side, full-data AMLE yields a likelihood approximation that converges pointwise to the posterior distribution. Under slightly stronger conditions, the mode of the approximation converges pointwise to the mode of the likelihood.

The rest of the paper is organized as follows. In Section 2 we describe the dynamic mixture distribution. Section 3 contains a detailed account of the AMLE versions with and without summary statistics. In sections 4 and 5 we report the outcomes of the simulation experiments and of the empirical analysis, respectively. Finally, Section 6 concludes the paper and outlines open problems and possible further developments.

2. The dynamic mixture model

The density of a dynamic mixture model is given by

$$f(x; \theta) = \frac{(1 - p(x; \theta_0))f_1(x; \theta_1) + p(x; \theta_0)f_2(x; \theta_2)}{Z}, \quad x \in \mathbb{R}^+, \tag{1}$$

where $\theta = (\theta_0, \theta_1, \theta_2)$, Z is a normalizing constant and θ_i , $i = 0, 1, 2$, are the parameter vectors of the weights $p(x; \theta_0)$, the body $f_1(x; \theta_1)$ and the tail $f_2(x; \theta_2)$, respectively. For the weight function it is convenient to employ the cumulative distribution function (cdf) of some continuous random variable, whereas f_1 and f_2 are continuous densities with positive support. Analogously to Frigessi et al. (2002), here $p(x; \theta_0)$ is taken to be equal to the Cauchy cdf:

$$p(x; \mu_c, \tau) = \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{x - \mu_c}{\tau}\right),$$

so that $\theta_0 = (\mu_c, \tau)$, where $\mu_c \in \mathbb{R}$ and $\tau \in \mathbb{R}^+$ are the location and scale parameter, respectively.

As for f_1 and f_2 , we use the lognormal and the zero-mean generalized Pareto densities, respectively. The use of the latter for the tail is related to its importance in EVT, since it is the asymptotic distribution of the excesses (see, e.g., Embrechts et al., 1997). As of the lognormal, it is chosen for two reasons. First, its flexibility can accommodate various possible shapes. Second, in economics there is a long-standing debate about the proper data-generating process of various relevant variables, with theoretical reasons suggesting a lognormal distribution, possibly with a Pareto-type tail. See, e.g., D’Acci (2019) and the references therein for city size, Axtell (2001), Di Giovanni et al. (2011), Tang (2015), Bee et al. (2017b) and Kondo et al. (2021) for firm size.

Accordingly, in the lognormal-GPD case the density (1) is given by

$$f(x; \theta) = \frac{(1 - p(x; \mu_c, \tau))f_1(x; \mu, \sigma^2) + p(x; \mu_c, \tau)f_2(x; \beta, \xi)}{Z}, \tag{2}$$

where $\mu_c, \mu, \xi \in \mathbb{R}$, $\tau, \sigma^2, \beta \in \mathbb{R}^+$, $\theta = (\mu_c, \tau, \mu, \sigma^2, \beta, \xi)'$, $\theta_1 = (\mu, \sigma^2)$, $\theta_2 = (\beta, \xi)$, $f_1(x; \mu, \sigma^2)$ is the lognormal density with parameters μ and σ^2 , $f_2(x; \beta, \xi)$ is the GPD pdf centered at 0 with scale and shape parameters equal to β and ξ , respectively. Finally, the normalizing constant Z is equal to

$$Z = Z(\theta) = 1 + \frac{1}{\pi} I,$$

where

$$I = \int_0^\infty \left[\frac{1}{\beta} \left(1 + \frac{\xi x}{\beta} \right)^{-1/\xi - 1} - \frac{1}{\sqrt{2\pi}\sigma x} e^{-\frac{1}{2} \left(\frac{\log x - \mu}{\sigma} \right)^2} \right] \arctan \left(\frac{x - \mu_c}{\tau} \right) dx. \tag{3}$$

See Frigessi et al. (2002) for details.

From a modeling point of view, in the literature there are at least two possible alternatives to (2), both based on the Pareto instead of the GPD. Scollnik (2007) proposes a lognormal-Pareto mixture constrained to have a continuous and differentiable density. The constraints reduce the number of free parameters by two; in particular, the mixing weight becomes a function of the lognormal variance and of the parameters of the Pareto distribution (Scollnik, 2007; Bee, 2015). Another possibility is the use of (2) with a mixing weight $p(x)$ equal to the Heaviside function. In this setup, the density is no longer continuous; see Bee (2022a) for the analysis of this case.

3. Estimation

In the rest of the paper, we focus on the lognormal-GPD dynamic mixture, but other models corresponding to different component densities can be estimated via the same method. The only requirement for implementing the proposed AMLE approach is indeed the ability to simulate from the two distributions of the mixture.

3.1. Maximum likelihood

As usual, maximum likelihood estimators can be found by maximizing the log-likelihood function obtained by taking the natural logarithm of (2):

$$l(\theta; \mathbf{x}) = \sum_{i=1}^n \log \left\{ \frac{(1 - p(x_i; \mu_c, \tau)) f_1(x_i; \mu, \sigma^2) + p(x_i; \mu_c, \tau) f_2(x_i; \beta, \xi)}{Z} \right\}. \tag{4}$$

Since the normalizing constant Z depends on all parameters, the evaluation of (3) is crucial for MLE. To achieve a sufficiently high level of precision in the approximation, Frigessi et al. (2002) suggest to split the integral on $[0, \infty)$ in a sum of integrals on finite intervals, such as $[0, 1]$, $[1, 2]$, and so on.

Numerical maximization of (4) is implemented in the `OpVaR` R package, where the `dmixing` function carries out a numerical approximation of the whole improper integral from 0 to ∞ . On the contrary, we have tackled the evaluation via the `quadinf` function of the `pracma` R package on non-overlapping intervals $[n - 1, n]$, $n \in \mathbb{N}$. The stopping criterion is $I_{n-1,n} < \epsilon_I$ with $\epsilon_I = 10^{-4}$, where $I_{n-1,n}$ is the value of the integral on $[n - 1, n]$. An unreported numerical comparison between the `OpVaR` and our procedure produced considerably better outcomes, especially in terms of variance, for the latter, which will always be used in the following.

The tolerance ϵ_I that defines the stopping criterion affects both the accuracy of the numerical evaluation of the normalizing constant and the computational burden of MLE. The actual value $\epsilon_I = 10^{-4}$ employed in the following has been determined via simulation (see Sect. 4.1).

3.2. Approximate Maximum Likelihood: some background

The Approximate Maximum Likelihood method is a simulation-based estimation procedure (Rubio and Johansen, 2013). We give a general description in this section, and then detail the full-data algorithm (Section 3.2.1) and the approach based on summary statistics (Section 3.2.2). Finally, Section 3.3 describes the last stage of the procedure, which is common to both cases.

Given a sample $\mathbf{x} = (x_1, \dots, x_n)' \in \mathbb{R}^{q \times n}$ from a distribution with density $f(\mathbf{x}; \theta)$, let $L(\theta; \mathbf{x})$ be the likelihood function, where $\theta \in \Theta \subset \mathbb{R}^p$ is a vector of parameters. We formally introduce the method by assuming a Bayesian setup with a prior distribution $\pi(\theta)$. This is done only for mathematical convenience: as discussed below, by using a uniform prior, in the following we carry out a likelihood analysis. The posterior $\pi(\theta|\mathbf{x})$ is given by

$$\pi(\theta|\mathbf{x}) = \frac{f(\mathbf{x}|\theta)\pi(\theta)}{\int_{\Theta} f(\mathbf{x}|\mathbf{t})\pi(\mathbf{t})d\mathbf{t}}. \tag{5}$$

Let's now introduce an approximation of the likelihood defined as follows:

$$\hat{f}_\epsilon(\mathbf{x}|\boldsymbol{\theta}) = \int_{\mathbb{R}^n} K_\epsilon(\mathbf{x}|\mathbf{z}) f(\mathbf{z}|\boldsymbol{\theta}) d\mathbf{z}, \tag{6}$$

where $K_\epsilon(\mathbf{x}|\mathbf{z})$ is a normalized Markov kernel (Casella and Robert, 2004, Sect. 6.2; Heyer, 1982; for the specific form used here, see (7) below) depending on a scale parameter ϵ . Now we plug (6) into (5) to obtain the following approximation of the posterior:

$$\hat{\pi}_\epsilon(\boldsymbol{\theta}|\mathbf{x}) = \frac{\hat{f}_\epsilon(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\int_{\Theta} \hat{f}_\epsilon(\mathbf{x}|\mathbf{t})\pi(\mathbf{t})d\mathbf{t}}.$$

If the prior $\pi(\boldsymbol{\theta})$ is uniform, the maximization of the likelihood is equivalent to the maximization of the posterior. Now we exploit the quantities defined above to give a pseudo-code of the algorithm.

Algorithm 1. (AMLE)

1. Obtain a sample $\boldsymbol{\theta}_\epsilon^* = (\theta_{\epsilon,1}^*, \dots, \theta_{\epsilon,\ell}^*)'$ from the approximate posterior $\hat{\pi}_\epsilon(\boldsymbol{\theta}|\mathbf{x})$; ℓ is commonly called ABC sample size;
2. Use this sample to construct a non-parametric estimator $\hat{\phi}$ of the density $\hat{\pi}_\epsilon(\boldsymbol{\theta}|\mathbf{x})$;
3. Compute the maximum of $\hat{\phi}$, $\tilde{\boldsymbol{\theta}}_{\ell,\epsilon}$. This is an approximation of the MLE $\hat{\boldsymbol{\theta}}$.

Algorithmically, AMLE is slightly different in the full-data approach and in the setup based on summary statistics. In particular, Step 1 is usually based on the ABC rejection algorithm (Beaumont, 2010), whose implementation is not the same in the two frameworks. Accordingly, we provide the details of the two cases in sections 3.2.1 and 3.2.2.

3.2.1. Stage 1a: AMLE without summary statistics

In a full-data approach, the ABC sample mentioned at Step 1 of Algorithm 1 is obtained as follows.

Algorithm 2. (ABC rejection algorithm)

1. Simulate $\boldsymbol{\theta}^*$ from the uniform prior $\pi(\cdot)$;
2. Generate $\mathbf{z} = (z_1, \dots, z_n)'$ from $f(\cdot|\boldsymbol{\theta}^*)$;
3. Accept $\boldsymbol{\theta}^*$ with probability $\propto K_\epsilon(\mathbf{x}|\mathbf{z})$, otherwise return to Step 1.

Let $P_\theta^{(n)} \in \mathcal{P}$ denote the distribution of $\mathbf{z}|\boldsymbol{\theta}$ and $P_0^{(n)} \in \mathcal{P}$ the distribution of \mathbf{x} , where $\mathcal{P} \stackrel{\text{def}}{=} \{P_\theta^{(n)}, \boldsymbol{\theta} \in \Theta \subset \mathbb{R}^p\}$ is the class of probability measures that is assumed to have generated the data. Furthermore, let $\rho : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}^+$ be a statistical distance on \mathcal{P} . The kernel K_ϵ is defined on the space of the true and simulated data:

$$K_\epsilon(\mathbf{x}|\mathbf{z}) \propto \begin{cases} 1 & \text{if } \rho(P_0^{(n)}, P_\theta^{(n)}) < \epsilon, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

Since $P_0^{(n)}$ is unknown and $P_\theta^{(n)}$ is intractable, (7) cannot be computed. The common workaround (Drovandi and Frazier, 2022) replaces the two distributions with the corresponding empirical counterparts, so that (7) becomes:

$$K_\epsilon(\mathbf{x}|\mathbf{z}) \propto \begin{cases} 1 & \text{if } \rho(\hat{F}(\mathbf{x}), \hat{F}_\theta(\mathbf{z})) < \epsilon, \\ 0 & \text{otherwise,} \end{cases}$$

where $\hat{F}(\cdot)$ and $\hat{F}_\theta(\cdot)$ are the empirical cdfs based on the observed data \mathbf{x} and the simulated data \mathbf{z} , respectively. As of ρ , there are various possibilities; in the following, we take it to be the Cramér-von Mises distance, given by

$$\hat{C}(\hat{F}, \hat{F}_\theta) = \int_{\mathbb{R}^q} [\hat{F}(t) - \hat{F}_\theta(t)]^2 d\hat{H}(t), \tag{8}$$

where $\hat{H}(t) = (\hat{F}(t) + \hat{F}_\theta(t))/2$. Note that (8) can be easily computed in terms of the ranks of the observed and simulated samples. Moreover, it is robust to fat tails and outliers, which is a very desirable property in our application and is therefore the main reason for preferring it; see Drovandi and Frazier (2022, p. 7) for details and comparisons to other distances.

The full-data approach may give the false impression that ABC (of which AMLE is a reinterpretation) is a method that compares statistical distances between simulated and empirical distributions. However, the main goal of ABC is the approximation of properties of the posterior distribution of parameters. Historically, it has been first employed in population genetics to estimate parameters in setups where the likelihoods of the underlying probability models were intractable. Moreover, such early applications typically used summary statistics. See, e.g., Beaumont et al. (2002) or Tavaré (2018) for details.

3.2.2. Stage 1b: AMLE with summary statistics

In the version of AMLE based on summary statistics, steps 1 and 2 of Algorithm 2 remain unchanged, whereas Step 3 is modified as follows.

3. Use \mathbf{x} to compute an m -dimensional summary statistics $\boldsymbol{\eta}(\mathbf{x})$; accept $\boldsymbol{\theta}^*$ with probability $\propto K_\epsilon^\rho(\boldsymbol{\eta}(\mathbf{x})|\boldsymbol{\eta}(\mathbf{z}))$, otherwise return to Step 1.

The kernel is now defined on the space of the summary statistics:

$$K_\epsilon^\rho(\boldsymbol{\eta}(\mathbf{x})|\boldsymbol{\eta}(\mathbf{z})) \propto \begin{cases} 1 & \rho(\boldsymbol{\eta}(\mathbf{x}), \boldsymbol{\eta}(\mathbf{z})) < \epsilon, \\ 0 & \text{otherwise.} \end{cases}$$

Here $\rho: \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}^+$ is a metric; usually, but not necessarily, ρ is the Euclidean distance.

Rubio and Johansen (2013) show that the replacement of the observed sample with a (vector of) summary statistics implies no loss of information if and only if $\boldsymbol{\eta}$ is a jointly sufficient statistic for the unknown parameters of the model: in this case, conditioning upon the sufficient statistics is the same as conditioning upon the sample.

3.3. Stage 2: computing the estimator

The two approaches outlined in sections 3.2.1 and 3.2.2 yield a sample $\boldsymbol{\theta}_\epsilon^*$ of size ℓ . The AMLE approach now exploits this sample to compute the estimate in the same way, regardless of the method used for obtaining it. Step 2 of Algorithm 1 requires to find the nonparametric estimator $\hat{\phi}$ and its maximum. For this task, we consider the following four techniques (Bee et al., 2017a):

- (i) the sample mean vector (“M”);
- (ii) the vector of the maxima of the univariate estimated kernel densities (“UK”);
- (iii) the maximum of the multivariate kernel density (“MK”);
- (iv) the maximum of the product of the univariate kernel densities, estimated using the marginal data (“PUK”);

In general, “M” is easy to compute but is only appropriate when the distribution of the simulated values is approximately symmetric. “MK” would in principle be the best approach, but it requires a large sample size, especially when the parameter space is high-dimensional. Hence, in practice, one should rather employ “UK” or “PUK”. While it is difficult to give a criterion for choosing between them, it is often the case that they yield similar estimates (see, e.g., Bee, 2022b).

Finally, the maximum of $\hat{\phi}$ computed via any of the methods (i)-(iv), denoted by $\tilde{\boldsymbol{\theta}}_{l,\epsilon} \stackrel{\text{def}}{=} \max_{\boldsymbol{\theta}} \hat{\phi}$, is the AMLE estimator, i.e. the AMLE approximation of the MLE $\hat{\boldsymbol{\theta}}$.

The first step of Algorithm 2 is based on uniform priors whose support must contain the true parameter value. Hence, it cannot be too narrow, or it may not contain the true parameter value. However, if it is too wide, the computational burden of the algorithm becomes unnecessarily large. In the present setup, since we can approximate and maximize numerically the likelihood, we exploit the MLEs to find the proper support of the uniform priors as follows.

- Compute the MLEs and their standard errors by means of non-parametric bootstrap;
- Discard possible outliers in each bootstrap distribution;
- For the lognormal and GPD parameters, the support of the uniform prior is set equal to the 99% confidence interval of the bootstrap distribution after discarding the outliers;
- For the Cauchy parameters, the support is given by the range of the bootstrap distribution after discarding the outliers.

The reason why the entire range is used for the Cauchy parameters is that, according to our Monte Carlo experiments, MLEs of these two parameters are less precise.

In the procedure outlined above, outliers are identified via the classical box-plot approach (Tukey, 1977). Since the histograms of the bootstrap distributions of the MLEs of μ_c and of τ (not shown here to save space) are skewed, we also resorted to the adjusted box-plot proposed by Hubert and Vandervieren (2008). With respect to the classical box-plot, the results only change for the MLEs of the Cauchy distribution parameters: in this case, the supports of the uniform priors are wider, since in presence of skewness the adjusted box-plot finds fewer outliers than the classical one (Hubert and Vandervieren, 2008). However, in our simulation experiments (see Section 4), the supports yielded by the classical box-plot are wide enough (i.e., they contain the true parameter value with a rather large margin), so that wider supports only increase the computational burden. Nevertheless, it may not be generally true that the classical box-plot is preferable, hence it is worth considering the use of the adjusted box-plot, especially when the distributions of the MLEs are very skewed.

3.4. Asymptotic properties

The limiting theory for the setup where AMLE is based on a summary statistic $\boldsymbol{\eta}$ is covered in Rubio and Johansen (2013): when $\boldsymbol{\eta}$ is non-sufficient, the AMLE approximation converges pointwise, under regularity conditions, to the posterior distribution. See Rubio and Johansen (2013, Proposition 2) for details.

In the full-data approach, stronger results hold true. Rubio and Johansen (2013) prove that, when sufficient statistics are available, under an additional assumption, the mode of the approximation $\hat{\pi}_\epsilon(\boldsymbol{\theta}|\mathbf{x})$ converges pointwise to the mode of the likelihood $\pi(\boldsymbol{\theta}|\mathbf{x})$. Noting that the original data or its empirical distribution constitute a sufficient statistic, the present setup turns out to be a special case of the framework where sufficient statistics exist. For the sake of completeness, we state explicitly the main theorem and a corollary; the proofs are omitted, since they are almost identical to those in Rubio and Johansen (2013). Before stating the results, we need to introduce the following condition.

Condition 1. (Concentration Condition) A family of symmetric Markov kernels with densities K_ϵ indexed by $\epsilon > 0$ satisfies the concentration condition if the member densities become increasingly concentrated as ϵ decreases:

$$\int_{\mathcal{B}_\epsilon(\mathbf{x})} K_\epsilon(\mathbf{x}|\mathbf{y})d\mathbf{y} = \int_{\mathcal{B}_\epsilon(\mathbf{x})} K_\epsilon(\mathbf{y}|\mathbf{x})d\mathbf{x} = 1, \quad \forall \epsilon > 0,$$

where $\mathcal{B}_\epsilon(\mathbf{x}) := \mathbf{z} : |\mathbf{z} - \mathbf{x}| \leq \epsilon$.

Proposition 1 below shows that the AMLE approximation converges pointwise to the posterior distribution.

Proposition 1. Let $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)' \in \mathbb{R}^{q \times n}$ be a sample from $f(\cdot|\boldsymbol{\theta})$, $\boldsymbol{\theta} \in \boldsymbol{\Theta} \subset \mathbb{R}^d$, and let $\rho : \mathbb{R}^{q \times n} \times \mathbb{R}^{q \times n} \rightarrow \mathbb{R}$ be the CvM distance (8). Suppose further that $f(\cdot|\boldsymbol{\theta})$ is ρ -continuous $\forall \boldsymbol{\theta} \in \mathbf{D}$, where $\mathbf{D} \subset \mathbb{R}^d$ is compact. Assume also that

$$\sup_{(\mathbf{t}, \boldsymbol{\theta}) \in \mathcal{B}_\epsilon \times \mathbf{D}} f(\cdot|\boldsymbol{\theta}) < \infty,$$

and let

$$K_\epsilon(\mathbf{x}|\mathbf{y}) = \begin{cases} 1 & \text{if } \rho(\mathbf{x}, \mathbf{y}) < \epsilon, \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

Then, $\forall \boldsymbol{\theta} \in \mathbf{D}$ and the kernel (9),

$$\lim_{\epsilon \rightarrow 0} \hat{\pi}_\epsilon(\boldsymbol{\theta}|\mathbf{x}) = \pi(\boldsymbol{\theta}|\mathbf{x}).$$

Under the additional condition of equicontinuity of $\hat{\pi}_\epsilon(\cdot|\mathbf{x})$, Proposition 1 can be used to show that the mode of the approximation $\hat{\pi}_\epsilon(\cdot|\mathbf{x})$ converges to the mode of the likelihood $\pi(\cdot|\mathbf{x})$.

Proposition 2. Let $\tilde{\boldsymbol{\theta}}_\epsilon$ be the unique maximum of $\hat{\pi}_\epsilon(\cdot|\mathbf{x})$, $\forall \epsilon > 0$, and assume that $\pi(\cdot|\mathbf{x})$ has a unique maximizer $\tilde{\boldsymbol{\theta}}$. Under the conditions in Proposition 1, and if $\hat{\pi}_\epsilon(\cdot|\mathbf{x})$ is equicontinuous on \mathbf{D} , then

$$\lim_{\epsilon \rightarrow 0} \hat{\pi}_\epsilon(\tilde{\boldsymbol{\theta}}_\epsilon|\mathbf{x}) = \pi(\tilde{\boldsymbol{\theta}}|\mathbf{x}).$$

4. Monte Carlo experiments

4.1. Simulation design and numerical details

Four experiments are run: we draw samples of size $n \in \{100, 500\}$ from the dynamic mixture with parameters $\mu_c = 1$, $\tau = 2$, $\mu = 0$, $\sigma = 0.5$, $\xi \in \{0.25, 0.5\}$, $\beta = 3.5$. The densities of the two distributions corresponding to $\xi = 0.25$ and $\xi = 0.5$ are displayed in Fig. 1.

MLEs result from the maximization of the log-likelihood function by means of the `optim` R command, where the normalizing constant is approximated as described in Section 3.1. Starting values for μ and σ are the lognormal MLEs computed with the observations below the median. Similarly, initial values for ξ and τ are the GPD MLEs obtained with the observations above the median. As for μ_c and τ , finding a reasonable initialization is more complicated. We use $\tau^0 = \log(\text{sd}(\mathbf{x})/2)$, where sd is the sample standard deviation, and $\mu_c^0 = q_{0.25}(\mathbf{x})$, where $q_{0.25}$ is the first quartile of the data. The former is identical to the initialization employed in the `OpVaR` package, whereas the latter is different: the package uses the third quartile, but according to a Monte Carlo study, this proposal strongly overestimates the true parameter value when the two distributions overlap considerably.




Fig. 1. Densities of the two simulated distributions. The bodies ($x < 15$) are represented in the top panel, the tails ($x \geq 15$) in the bottom panels.

Table 1

MLE. Computing times (in seconds) and RMSE of the MLEs in 100 replications of the simulation experiment for threshold $\epsilon_l \in \{10^{-s}\}_{s=2,\dots,8}$. The true value of ξ is 0.25 and the sample size is $n = 100$.

ϵ_l	Time	μ_c	τ	μ	σ	β	ξ
10^{-2}	5.26	5260.488	7771.369	0.260	0.235	1.902	0.271
10^{-3}	11.43	4498.541	7423.595	0.215	0.241	1.777	0.255
10^{-4}	17.44	3850.602	7006.747	0.182	0.187	1.334	0.232
10^{-5}	52.09	4227.770	6312.171	0.173	0.201	1.389	0.242
10^{-6}	124.61	4489.436	7123.037	0.180	0.186	1.330	0.239
10^{-7}	233.14	3917.092	7107.936	0.176	0.192	1.332	0.241
10^{-8}	385.21	3945.839	7216.598	0.175	0.191	1.330	0.234

The choice of the numerical value of the tolerance ϵ_l is based on a Monte Carlo simulation: we simulate $n = 100$ observations from (2) in the setup with $\xi = 0.25$ and compute the MLEs with $\epsilon_l \in \{10^{-s}\}_{s=2,\dots,8}$. After repeating the experiment 100 times, we calculate the average computing time and the RMSE of all parameters (see Table 1). The RMSEs suggest no significant improvement with $\epsilon_l < 10^{-4}$, hence in the following we set $\epsilon_l = 10^{-4}$. As a side remark, it is worth noting that the outcomes in Table 1 seem to denote a difficulty in identifying the Cauchy parameters, whose RMSEs remain quite large for all values of ϵ_l .

We have implemented AMLE with and without summary statistics, as in sections 3.2.1 and 3.2.2, respectively. In the former approach, the summary statistic is the empirical characteristic function, which has proved to be an effective choice in other cases (see, e.g., Bee and Trapin, 2018; Bee, 2022b; Tafakori et al., 2022). However, a small simulation experiment suggested a better performance of the full-data approach, both in statistical and computational terms (smaller root-mean-squared-error and shorter computing time). Hence, only the outcomes of the full-data approach based on the CvM distance are shown in this section.

The ABC sample size is $\ell = 100$, with $k = 5 \cdot 10^5$, since a small pilot simulation with $k = 2 \cdot 10^6$ did not suggest any significant improvement in the statistical properties of the estimators with respect to $k = 5 \cdot 10^5$. The supports of the uniform priors are determined via the procedure outlined in Section 3.2.2. Non-parametric bootstrap is replaced in this section by parametric bootstrap, since here we repeatedly simulate the true distribution.

Table 2

Case $\xi = 0.25$. Bias, standard deviation and RMSE of the estimates obtained via AMLE and MLE in $B = 100$ replications of the simulation experiment. The true value of ξ is 0.25.

n			μ_c	τ	μ	σ	β	ξ
100	Bias	AMLE	0.554	0.879	0.090	0.114	0.375	0.008
		MLE	-428.537	686.732	0.053	0.034	0.300	-0.028
	Sd	AMLE	0.554	0.221	0.123	0.088	0.571	0.045
		MLE	4273.488	6823.124	0.190	0.200	1.132	0.197
	RMSE	AMLE	0.784	0.907	0.152	0.144	0.683	0.046
		MLE	4294.921	6857.596	0.197	0.203	1.171	0.199
500	Bias	AMLE	0.201	0.283	0.0179	0.035	0.184	-0.013
		MLE	-0.018	0.429	-0.000	-0.004	0.029	0.007
	Sd	AMLE	0.225	0.348	0.049	0.046	0.319	0.054
		MLE	0.839	3.198	0.069	0.080	0.526	0.097
	RMSE	AMLE	0.302	0.448	0.052	0.056	0.368	0.056
		MLE	0.839	3.227	0.069	0.080	0.527	0.097

Table 3

Case $\xi = 0.5$. Bias, standard deviation and RMSE of the estimates obtained via AMLE and MLE in $B = 100$ replications of the simulation experiment. The true value of ξ is 0.5.

n			μ_c	τ	μ	σ	β	ξ
100	Bias	AMLE	0.881	2.084	0.049	0.087	0.780	-0.094
		MLE	0.677	0.447	0.018	0.019	0.652	-0.066
	Sd	AMLE	1.217	1.542	0.121	0.114	0.923	0.146
		MLE	2.093	2.953	0.158	0.168	1.093	0.193
	RMSE	AMLE	1.502	2.592	0.131	0.143	1.208	0.174
		MLE	2.199	2.980	0.159	0.169	1.271	0.204
500	Bias	AMLE	0.403	0.062	0.027	0.035	0.158	-0.005
		MLE	0.048	0.330	0.009	0.004	0.085	-0.001
	Sd	AMLE	0.443	0.605	0.059	0.053	0.413	0.075
		MLE	0.708	1.251	0.067	0.063	0.496	0.083
	RMSE	AMLE	0.599	0.608	0.065	0.062	0.442	0.075
		MLE	0.709	1.293	0.068	0.063	0.503	0.083

In each experiment, we compute the bias, the standard deviation and the RMSE of all parameters, with $B = 100$ replications. The AMLE outcomes shown in this section are based on the sample mean, since the results of the other approaches (“UK”, “MK” and “PUK”) are nearly identical.

Computer-intensive methods are usually characterized by a much larger computational burden, and AMLE makes no exception. When $n = 500$, regardless of the true value of ξ , AMLE needs about 55 minutes, whereas MLE takes 32 seconds (see Table 1). When $n = 100$, these times reduce to 12.5 minutes for AMLE and 17 seconds for MLE.¹

4.2. Simulation results

Tables 2 and 3 display bias, standard deviation and root-mean-squared-error (RMSE) of both AMLEs and MLEs of the parameters. The tables suggest that AMLE is always preferable to MLE in terms of RMSE, by a larger amount when $\xi = 0.25$; in most cases, the gain in RMSE results from a smaller standard deviation of AMLE, which more than compensates a slightly larger bias. Moreover, the advantage is larger for small sample sizes.

For the case $\xi = 0.25$, the simulated distributions of the AMLEs of the parameters are displayed in Fig. 2 (for $n = 100$) and Fig. 3 (for $n = 500$), along with the supports of the uniform distributions used in the AMLE algorithm; the corresponding plots when $\xi = 0.5$ are very similar and therefore omitted. The histograms are approximately bell-shaped for both sample sizes, which explains why the four AMLE approaches yield very similar outcomes.

Similarly to Frigessi et al. (2002), our analysis suggests that estimation of the dynamic weight parameters is difficult, in particular with the MLE method: as can be seen from the large bias and standard deviation of $\hat{\mu}_c^{MLE}$ and $\hat{\tau}^{MLE}$ in Table 2, there are cases where the numerical optimization of the log-likelihood function does not yield sensible outcomes. It is also worth noting that in a few cases (approximately 2-5% of the replications, with the largest values when $\xi = 0.25$ and $n = 100$) the numerical maximization of the log-likelihood broke down without producing a result. In such cases we have discarded the sample and simulated the observations again.

¹ These computing times are based on R codes run on a Windows machine with an i7-6700 CPU @ 3.40 GHz. The simulation experiments have been run in parallel on the HPC computer cluster of the University of Trento, with a total time of about 22 hours for 100 replications when $n = 500$.




Fig. 2. Simulated distributions of the AMLE estimators with $n = 100$. The number of replications is $B = 100$ and the true parameter values are $(\mu_c, \tau, \mu, \sigma, \xi, \beta) = (1, 2, 0, 0.5, 0.25, 3.5)$. The dashed vertical lines are the lower and upper bounds of the uniform priors.

5. Empirical analysis

5.1. Operational risk

An important application of the model considered in this paper is the analysis of loss distributions, since they are typically skewed and heavy-tailed. In many cases, the final goal is the estimation of risk measures such as the Value-at-Risk (VaR) or the Expected Shortfall (ES). For this purpose, the tail must be estimated with a high level of precision, so that accurate models of the tail are of paramount importance.

In this section, we use operational risk losses collected at the Italian bank Unicredit between 2005 and 2014. In particular, we model the monetary amounts of the losses in the *Business Disruption and System Failure* (BDSM) business class. The sample size is $n = 152$. The data are displayed in Fig. 4, along with the densities estimated via AMLE and MLE. Table 4 shows the MLEs and AMLEs; standard errors are computed via non-parametric bootstrap with $B = 100$ replications.

There is a considerable difference between AMLE and MLE when estimating τ, β and, to a lesser extent, μ_c . The large standard errors imply that, for these parameters, the precision is low, but in terms of variability AMLEs are definitely preferable. Notwithstanding the differences in the estimated values of some parameters, the densities in Fig. 4 are rather similar: this suggests that the errors in the estimates of individual parameters are likely to offset each other (Frigessi et al., 2002). Notice also that the observations above 50 in Fig. 4 are very likely to come from the GPD, since the expected values of the lognormal and GPD (computed using the PUK estimators) are as different as $\exp\{1.307 + 0.32^2/2\} = 3.889$ and $5.348/(1 - .916) = 63.667$, respectively.




Fig. 3. Simulated distributions of the AMLE estimators with $n = 500$. The number of replications is $B = 100$ and the true parameter values are $(\mu_c, \tau, \mu, \sigma, \xi, \beta) = (1, 2, 0, 0.5, 0.25, 3.5)$. The dashed vertical lines are the lower and upper bounds of the uniform priors.

Table 4

Parameter estimates and standard errors obtained via AMLE and MLE in the operational risk example. Standard errors are computed via non-parametric bootstrap with $B = 100$ bootstrap replications.

	μ_c	τ	μ	σ	ξ	β
AMLE M	3.882 (0.626)	0.587 (0.352)	1.309 (0.079)	0.330 (0.053)	0.921 (0.119)	5.557 (0.819)
AMLE UK	3.410 (12.742)	0.435 (0.673)	1.303 (0.082)	0.330 (0.055)	0.912 (0.137)	5.466 (1.497)
AMLE MK	3.362 (0.846)	0.453 (0.389)	1.295 (0.106)	0.293 (0.080)	0.908 (0.197)	4.930 (1.682)
AMLE PUK	3.423 (0.692)	0.447 (0.224)	1.307 (0.090)	0.320 (0.064)	0.916 (0.160)	5.348 (1.503)
MLE	4.947 (1.402)	3.116 (1.276)	1.197 (0.103)	0.290 (0.082)	0.706 (0.191)	7.871 (2.836)




Fig. 4. The operational risk data with superimposed the dynamic lognormal-GPD density estimated via AMLE (continuous) and MLE (dashed). The upper and lower panel respectively show the body ($x \leq 50$) and the tail ($x > 50$). AMLE estimators are based on the “PUK” method.




Fig. 5. The ABC samples of the six parameters in the operational risk data.

With regard to the four AMLE methods, there is some difference mostly when estimating μ_c and τ . To understand why, Fig. 5 displays the distribution of the ABC samples corresponding to the six parameters. The distributions of the samples of $\hat{\mu}_c$ and $\hat{\tau}$ are the least symmetric. In such cases, it is not recommended to estimate the maximum with the sample mean,

Table 5
Estimated quantiles obtained via AMLE, MLE and the POT method in the operational risk example. Empirical quantiles of the observed data are reported as well.

	50%	90%	95%	99%	99.5%
AMLE M	4.493	39.404	81.632	385.415	731.770
AMLE UK	4.911	40.916	82.406	369.166	688.687
AMLE MK	4.703	37.343	75.060	323.465	617.791
AMLE PUK	4.907	40.298	82.216	390.626	750.067
MLE	4.222	36.763	66.717	230.602	385.943
GPD	-	39.941	84.153	257.865	380.208
EMP	4.990	37.176	73.569	274.323	355.255

Table 6
Estimated tail probabilities obtained via AMLE, MLE and the POT method in the operational risk example. Empirical tail probabilities of the observed data are reported as well.

	100	150	200	300	400	500	600
AMLE M	0.038	0.025	0.019	0.012	0.009	0.007	0.005
AMLE UK	0.041	0.027	0.020	0.013	0.009	0.007	0.006
AMLE MK	0.037	0.025	0.018	0.012	0.009	0.007	0.006
AMLE PUK	0.041	0.027	0.020	0.013	0.010	0.007	0.006
MLE	0.030	0.018	0.012	0.007	0.005	0.003	0.003
GPD	0.041	0.024	0.015	0.008	0.005	0.003	0.002
EMP	0.039	0.026	0.026	0.013	0.000	0.000	0.000

Table 7
Estimated Expected Shortfall obtained via AMLE, MLE and the POT method in the operational risk example. The empirical ES of the observed data is reported as well.

	50%	90%	95%	99%	99.5%
AMLE M	80.632	346.086	634.897	2535.654	4555.332
AMLE UK	85.142	368.395	680.077	2764.897	5010.873
AMLE MK	80.161	343.524	628.964	2490.805	4438.181
AMLE PUK	109.452	490.383	922.963	3961.587	7399.500
MLE	38.756	142.389	235.664	716.620	1133.496
GPD	-	137.637	216.876	528.209	747.477
EMP	35.933	123.898	194.493	368.010	393.020

hence we suggest using either “UK” or “PUK”, which are approximately identical. Notwithstanding the small sample size, “MK” also yields estimates close to “UK” and “PUK”.

Tables 5 and 6 respectively show quantiles at different levels and estimated tail probabilities $P(X \geq t)$ for various values of t . Note that, in this setup, the α -quantile can be interpreted as the Value-at-Risk at level α (see, e.g., McNeil et al., 2015, Section 2.3.2).

Given that we focus on the tail, as a benchmark we also compute the quantiles based on the GPD asymptotic approximation fitted to the excesses $x_i - u$, where u is taken to be equal to the 90-th quantile, by means of the Peaks-over-Threshold (POT) method (McNeil et al., 2015, Section 5.3.2). Since this is a theoretically well-grounded approach aimed at tail approximation, it is expected to be a reliable benchmark. Finally, we also report the sample quantiles of the observed data (Table 5) and the empirical estimate of the tail probabilities $p_t = \#\{x_i : x_i \geq t\}/n$ (Table 6). Tables 5 and 6 suggest that the AMLE approaches yield quite similar results; only MK gives slightly smaller estimated quantiles. For the smallest quantile levels (90 and 95%) and the smallest thresholds (≤ 200), the AMLE estimates are closer than MLE to both the GPD and the empirical estimates. On the other hand, when one moves farther into the tail, the MLE-based estimates are very close to the GPD-based estimates. It is interesting to notice that, even though the densities in Fig. 4 look very close to each other, when the quantile level is high the difference between AMLE and MLE becomes non-negligible.

In terms of Expected Shortfall, Table 7 shows that the dynamic mixture, especially when estimated via AMLE, yields numerical values larger than both the empirical and the GPD-based measure. However, interpreting such a comparison in terms of ES with a small sample size requires some care, since ES is very sensitive to a few very large observations.

Finally, to assess the quality of the VaR and ES estimates obtained, we use the Mean Absolute Relative Error (MARE) proposed by Naderi et al. (2020):

$$MARE = \frac{1}{n_l} \sum_{i=1}^{n_l} \left| \frac{M_i - \hat{M}_i}{M_i} \right|,$$

Table 8
Mean absolute relative error for AMLE, MLE and the POT method in the operational risk example.

	M	UK	MK	PUK	MLE	GPD
MARE VaR	0.264	0.210	0.223	0.219	0.103	0.087
MARE ES	7.316	6.262	8.248	9.298	3.001	1.670




Fig. 6. The metropolitan cities data with superimposed the dynamic lognormal-GPD density estimated via AMLE (continuous) and MLE (dashed). The upper and lower panel respectively show the body ($x \leq 300$) and the tail ($x > 300$). AMLE estimators are based on the “PUK” method.

where n_i is the number of levels and M_i and \hat{M}_i denote the empirical and estimated risk measure at the i -th level, respectively. Table 8 displays the results, which confirm that MLE is closer to the empirical values.

5.2. Metropolitan cities

In the last two decades or so, city size data have often been the focus of a lively dispute between scholars: while some argue that the whole distribution is lognormal, others claim that the body is lognormal but the tail follows a Pareto-type distribution. See, e.g., D’Acci (2019) and the references therein.

Here we study the distribution of the 2019 population estimate, divided by 10 000, of the 415 US metropolitan areas computed by the US Census Bureau²; for an earlier investigation of US metropolitan areas, see Gabaix and Ibragimov (2011). The data and the two estimated densities are shown in Fig. 6. Whereas in the previous application the two densities were very similar to each other (see Fig. 4), in this case the MLE-based density gives less weight to the body and more weight to the tail; this fact is confirmed by the estimated quantiles in Table 10.

Parameter estimates and standard errors are shown in Table 9. Once again, the Cauchy parameters are difficult to estimate, and one can notice large differences between AMLE and MLE in the estimated value of the location parameter μ_c . In terms of variability, the standard errors of AMLEs are almost always smaller than those of MLEs, and for some parameters the gain is considerable.

Analogously to the previous section, Table 10 shows selected quantiles estimated via both methods, and Table 11 reports estimated tail probabilities $P(X \geq t)$ for various values of t .

Tables 10 and 11 suggest that the AMLE-based estimated quantiles and tail probabilities are closer than the MLE-based quantities to both the GPD-based and the empirical measures. MLE-based quantiles and tail probabilities are considerably larger, especially at high levels and thresholds. The ES measures in Table 12 essentially confirm these remarks.

² www.census.gov/data/datasets/time-series/demo/pepsect/2010s-total-metro-and-micro-statistical-areas.html.

Table 9

Parameter estimates and standard errors obtained via AMLE and MLE in the metropolitan cities example. Standard errors are computed via non-parametric bootstrap with $B = 100$ bootstrap replications.

	μ_c	τ	μ	σ	ξ	β
AMLE M	26.493 (0.913)	2.667 (0.288)	2.869 (0.039)	0.455 (0.021)	0.595 (0.078)	55.770 (2.427)
AMLE UK	24.361 (1.367)	3.991 (0.391)	2.873 (0.039)	0.456 (0.022)	0.590 (0.081)	55.216 (2.354)
AMLE MK	23.975 (1.775)	2.499 (0.685)	2.887 (0.049)	0.475 (0.037)	0.583 (0.107)	50.971 (3.359)
AMLE PUK	24.714 (1.470)	3.246 (0.501)	2.895 (0.047)	0.467 (0.038)	0.547 (0.098)	54.523 (2.526)
MLE	15.002 (2.992)	4.504 (0.426)	2.829 (0.079)	0.359 (0.046)	0.680 (0.118)	56.602 (9.570)

Table 10

Estimated quantiles obtained via AMLE, MLE and the POT method in the metropolitan cities example. Empirical quantiles of the observed data are reported as well.

	50%	90%	95%	99%	99.5%
AMLE M	29.417	205.714	359.239	1108.211	1747.710
AMLE UK	27.615	190.465	326.674	1002.483	1625.888
AMLE MK	28.791	204.030	347.917	1090.096	1708.617
AMLE PUK	28.302	195.501	337.762	1069.634	1716.508
MLE	34.223	273.854	487.290	1567.529	2576.189
GPD	-	227.171	366.553	884.398	1231.582
EMP	27.32	225.09	342.88	919.46	1170.91

Table 11

Estimated tail probabilities obtained via AMLE, MLE and the POT method in the metropolitan cities example. Empirical tail probabilities of the observed data are reported as well.

	200	400	600	700	800	900	1000	1500
AMLE M	0.101	0.041	0.024	0.019	0.0153	0.0129	0.011	0.006
AMLE UK	0.095	0.039	0.022	0.018	0.015	0.012	0.011	0.006
AMLE MK	0.099	0.041	0.023	0.019	0.0153	0.0126	0.011	0.006
AMLE PUK	0.097	0.041	0.023	0.019	0.016	0.014	0.012	0.007
MLE	0.138	0.062	0.037	0.030	0.026	0.022	0.019	0.011
GPD	-	0.043	0.021	0.016	0.012	0.001	0.008	0.003
EMP	0.123	0.046	0.029	0.019	0.012	0.012	0.001	0.002

Table 12

Estimated Expected Shortfall obtained via AMLE, MLE and the POT method in the metropolitan cities example. The empirical ES of the observed data is reported as well.

	50%	90%	95%	99%	99.5%
AMLE M	194.015	648.499	1036.208	2944.561	4530.217
AMLE UK	199.711	680.113	1097.263	3222.347	5102.488
AMLE MK	190.648	631.351	996.834	2693.999	3989.616
AMLE PUK	201.276	685.638	1108.676	3272.185	5207.477
MLE	303.493	1106.343	1857.966	6094.621	10193.591
GPD	-	516.622	747.522	1605.441	2180.643
EMP	162.353	495.473	720.431	1275.270	1475.528

Table 13 displays the MARE results, which allow us to conclude that AMLE-based risk measures are preferable to MLE-based measures.

In this application it is relevant to study the degree of overlap of the two component distributions, and the possible presence of a GPD tail. Most of the investigations carried out in the literature try to find a threshold where the power-law distribution starts, as if the two distributions did not overlap. This way of proceeding is not correct, since the two distributions actually overlap. The present approach, however, does not suffer from this drawback.

Table 13
Mean absolute relative error for AMLE, MLE and the POT method in the metropolitan cities example.

	M	UK	MK	PUK	MLE	GPD
MARE VaR	0.175	0.174	0.163	0.174	0.611	0.042
MARE ES	2.790	2.892	2.937	3.331	5.381	1.021




Fig. 7. Estimated dynamic weight functions obtained via AMLE “PUK” and MLE in the metropolitan cities example.

Table 14
Estimated values $x : p(x; \mu_c, \tau) > \alpha$ obtained with all the methods for various values of α . The numbers of observations exceeding the corresponding value x_α are reported in parentheses.

	0.90	0.95	0.99	0.995
AMLE M	10.313	11.129	15.260 (304)	25.883 (219)
AMLE UK	10.401	11.577	18.167 (263)	52.672 (137)
AMLE MK	10.227	10.983	14.966 (311)	22.925 (232)
AMLE PUK	10.313	11.351	16.753 (288)	32.620 (184)
MLE	10.081	11.129	19.769 (256)	68.386 (113)

Bee (2022a) analyzes the same dataset; even though the model is different, it is interesting to compare the number of Pareto observations found in that paper to the number of observations that are likely to be GPD in the present dynamic mixture model.

The likelihood of an observation being GPD can be measured by means of the estimated weight $p(x; \mu_c, \tau)$, by setting a high probability level α and finding the smallest $x : p(x; \mu_c, \tau) > \alpha$. This value (x_α , say) can be interpreted as a threshold above which almost all the observations are GPD. Fig. 7 shows the weights estimated via AMLE “PUK” and MLE.

Table 14 displays thresholds $x_\alpha : p(x_\alpha; \mu_c, \tau) = \alpha$ estimated with all the methods for various values of α . Given α , the probability that an observation $x > x_\alpha$ is generated from the GPD is thus larger than α . The number of observations exceeding a given x_α is reported in parentheses: for example, there are 304 observations larger than 15.26. In each row, these numbers give an idea of the number of GPD observations, since in the two rightmost columns only 1% and 0.5% of the observations are not GPD. For comparison, by fitting a lognormal-Pareto model with fixed mixing weight to the same dataset, Bee (2022a) found 214 Pareto observations.

For high values of α , the difference between the estimation methods is non-negligible, but in all cases there is a rather large number of GPD observations, consistently with the findings in Bee (2022a).

6. Conclusion

In this paper we have proposed a simulation-based procedure for the estimation of the parameters of a dynamic mixture. The motivation of this choice is that numerical MLE of this model is complicated by the need of approximating numerically the normalizing constant of the distribution.

To this aim, we have implemented an approximate maximum likelihood approach in the lognormal-GPD case, using the Cramér-von Mises distance to measure the discrepancy between observed and simulated samples. On the theoretical side,

we have proved pointwise convergence of the AMLE approximation to the likelihood. In addition, we have given conditions under which the mode of the approximation converges to the mode of the likelihood. Simulation results suggest that AMLE outperforms MLE in terms of RMSE, albeit with a higher computational cost. Two empirical applications confirm that the approach is successful at fitting skewed datasets in economics and finance.

Various issues need further research. First, it is possible to carry out a thorough comparison between the dynamic mixture considered in this paper and the lognormal-Pareto model mentioned in Section 2 (Scollnik, 2007), in terms of both goodness of fit and numerical difficulty of the estimation procedures. Second, since the main issue with standard MLE is the evaluation of the normalizing constant, this approach could be dramatically improved by finding a weighting function that allows one to solve the integral in closed form, possibly at the price of some loss in flexibility. Third, computing bootstrap standard errors is time-consuming. Wang and Lin (2016) have developed the score vector and Hessian matrix for multivariate t mixtures, extending earlier work by Boldea and Magnus (2009) for multivariate Gaussian mixtures. Since their results cannot be directly applied to our setup, which is somewhat different, further research on a possible extension to dynamic mixtures is necessary. Fourth, setting up a more parsimonious model may help to reduce estimation difficulties: one possibility would be to replace the Cauchy cdf with a single-parameter cdf, such as the exponential. Similarly, exploring the performance of a model based on distributions different from the lognormal and the GPD may be of interest. Finally, when the scale parameter of the Cauchy cdf tends to zero, the dynamic weight converges to the Heaviside function: it may therefore be of interest to exploit the estimated value of this parameter to select the type of weight that is more appropriate for a given dataset.

Acknowledgements

I would like to thank an associate editor and two anonymous reviewers for their valuable comments. The R package `LNPar`, containing the data used in the paper and the implementation of the methodology, is available at the address <https://github.com/marco-bee/LNPar>.

References

- Axtell, R.L., 2001. Zipf distribution of U.S. firm sizes. *Science* 293, 1818–1820.
- Beaumont, M.A., 2010. Approximate Bayesian computation in evolution and ecology. *Annu. Rev. Ecol. Evol. Syst.* 41, 379–406.
- Beaumont, M.A., 2019. Approximate Bayesian computation. *Annu. Rev. Stat. Appl.* 6, 379–403.
- Beaumont, M.A., Zhang, W., Balding, D.J., 2002. Approximate Bayesian computation in population genetics. *Genetics* 162, 2025–2035.
- Bee, M., 2015. Estimation of the lognormal-Pareto distribution using probability weighted moments and maximum likelihood. *Commun. Stat., Simul. Comput.* 44, 2040–2060.
- Bee, M., 2022a. On discriminating between lognormal and Pareto tail: an unsupervised mixture-based approach. *Adv. Data Anal. Classif.* <https://doi.org/10.1007/s11634-022-00497-4>.
- Bee, M., 2022b. The truncated g-and-h distribution: estimation and application to loss modeling. *Comput. Stat.* 37, 1771–1794.
- Bee, M., Benedetti, R., Espa, G., 2017a. Approximate maximum likelihood estimation of the Bingham distribution. *Comput. Stat. Data Anal.* 108, 84–96.
- Bee, M., Espa, G., Giuliani, D., 2015. Approximate maximum likelihood estimation of the autologistic model. *Comput. Stat. Data Anal.* 84, 14–26.
- Bee, M., Riccaboni, M., Schiavo, S., 2017b. Where Gibrat meets Zipf: scale and scope of French firms. *Physica A* 481, 265–275.
- Bee, M., Trapin, L., 2018. A characteristic function-based approach to approximate maximum likelihood estimation. *Commun. Stat., Theory Methods* 47, 3138–3160.
- Bernton, E., Jacob, P.E., Gerber, M., Robert, C.P., 2019. Approximate Bayesian computation with the Wasserstein distance. *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 81, 235–269.
- Boldea, O., Magnus, J.R., 2009. Maximum likelihood estimation of the multivariate normal mixture model. *J. Am. Stat. Assoc.* 104, 1539–1549.
- Casella, G., Robert, C., 2004. *Monte Carlo Statistical Methods*, second ed. Springer.
- D'Acci, L., 2019. *The Mathematics of Urban Morphology*. Birkhäuser.
- Di Giovanni, J., Levchenko, A.A., Rancière, R., 2011. Power laws in firm size and openness to trade: measurement and implications. *J. Int. Econ.* 85, 42–52. <https://doi.org/10.1016/j.jinteco.2011.05.003>.
- Drovandi, C., Frazier, D., 2022. A comparison of likelihood-free methods with and without summary statistics. *Stat. Comput.* 32, Paper No. 42.
- Embrechts, P., Klüppelberg, C., Mikosch, T., 1997. *Modelling Extremal Events for Insurance and Finance*. Springer.
- Frigessi, A., Haug, O., Rue, H., 2002. A dynamic mixture model for unsupervised tail estimation without threshold selection. *Extremes* 3, 219–235.
- Gabaix, X., Ibragimov, R., 2011. Rank-1/2: A simple way to improve the OLS estimation of tail exponents. *J. Bus. Econ. Stat.* 29, 24–39.
- Heyer, H., 1982. *Theory of Statistical Experiments*. Springer.
- Hubert, M., Vandervieren, E., 2008. An adjusted boxplot for skewed distributions. *Comput. Stat. Data Anal.* 52, 5186–5201.
- Kondo, I., Lewis, L., Stella, A., 2021. Heavy tailed, but not Zipf: firm and establishment size in the U.S. U.S. Census working paper number CES-21-15.
- McNeil, A., Frey, R., Embrechts, P., 2015. *Quantitative Risk Management: Concepts, Techniques, Tools*, second ed. Princeton University Press.
- Naderi, M., Jamalizadeh, A., Wang, W.L., Lin, T.I., 2020. Evaluating risk measures using the normal mean-variance Birnbaum-Saunders distribution. In: *Computational and Methodological Statistics and Biostatistics: Contemporary Essays in Advancement*. Springer, pp. 187–209.
- Park, M., Jitkrittum, W., Sejdinovic, D., 2016. K2-ABC: approximate Bayesian computation with kernel embeddings. In: Gretton, A., Robert, C. (Eds.), *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pp. 398–407.
- Rubio, F.J., Johansen, A.M., 2013. A simple approach to maximum intractable likelihood estimation. *Electron. J. Stat.* 7, 1632–1654.
- Scollnik, D., 2007. On composite lognormal-Pareto models. *Scand. Actuar. J.* 1, 20–33.
- Sunnäker, M., Busetto, A., Numminen, E., Corander, J., Foll, M., Dessimoz, C., 2013. Approximate Bayesian computation. *PLoS Comput. Biol.* 9, e1002803.
- Tafakori, L., Bee, M., Soltani, A., 2022. Some analytical results on bivariate stable distributions with an application in operational risk. *Quant. Finance* 22, 1355–1369.
- Tang, A., 2015. Does Gibrat's law hold for Swedish energy firms? *Empir. Econ.* 49, 659–674.
- Tavaré, S., 2018. On the history of ABC. In: Sisson, S.A., Fan, Y., Beaumont, M. (Eds.), *Handbook of Approximate Bayesian Computation*. Chapman and Hall.
- Tukey, J., 1977. *Exploratory Data Analysis*. Addison-Wesley.
- Wang, W.L., Lin, T.I., 2016. Maximum likelihood inference for the multivariate t mixture model. *J. Multivar. Anal.* 149, 54–64.