



Minimum L^q -distance estimators for non-normalized parametric models

Steffen BETSCH^{1*}, Bruno EBNER¹ , and Bernhard KLAR¹

¹*Institute of Stochastics, Karlsruhe Institute of Technology (KIT), D-76131 Karlsruhe, Germany*

Key words and phrases: Burr Type XII distribution; empirical processes; exponential-polynomial models; measurable selections; minimum distance estimators; Rayleigh distribution; Stein discrepancies
MSC 2010: Primary 62F10; secondary 62F12, 62E10.

Abstract: We propose and investigate a new estimation method for the parameters of models consisting of smooth density functions on the positive half axis. The procedure is based on a recently introduced characterization result for the respective probability distributions, and is to be classified as a minimum distance estimator, incorporating as a distance function the L^q -norm. Throughout, we deal rigorously with issues of existence and measurability of these implicitly defined estimators. Moreover, we provide consistency results in a common asymptotic setting, and compare our new method with classical estimators for the exponential, the Rayleigh and the Burr Type XII distribution in Monte Carlo simulation studies. We also assess the performance of different estimators for non-normalized models in the context of an exponential-polynomial family. *The Canadian Journal of Statistics* 00: 000–000; 2020 © 2020 Statistical Society of Canada

Résumé: Les auteurs proposent et étudient une nouvelle méthode d'estimation pour les paramètres de modèles constitués de fonctions de densité lisses sur le demi-axe positif. Basée sur un résultat récent à propos de la représentation des distributions de probabilité en question, leur procédure peut être classée comme un estimateur de distance minimale incorporant la norme L^q comme fonction de distance. Les auteurs gèrent de façon rigoureuse les questions d'existence et de mesurabilité de ces estimateurs définis implicitement. Ils fournissent également des résultats de convergence dans un contexte asymptotique commun, puis présentent des simulations de Monte Carlo où leur nouvelle méthode est comparée à des estimateurs classiques pour les distributions exponentielles, Rayleigh et Burr de type XII. Les auteurs évaluent également la performance de différents estimateurs pour les modèles non normalisés dans le contexte d'une famille exponentielle polynomiale. *La revue canadienne de statistique* 00: 000–000; 2020 © 2020 Société statistique du Canada

1. INTRODUCTION

One of the most classical problems in statistics is the estimation of the parameter vector of a parametrized family of probability distributions. It presents itself in a significant share of applications because parametric models often contribute a reasonable compromise between flexibility in the shape of the statistical model and meaningfulness of the conclusions that can be drawn from the model. As a consequence, all kinds of professions are confronted with the issue of parameter estimation, be it meteorologists, engineers or biologists. Throughout the last decades,

* Author to whom correspondence may be addressed.
E-mail: Steffen.Betsch@kit.edu

a vast amount of highly focused estimation procedures for all kinds of situations have been provided, but the procedure that is arguably used most often remains the maximum likelihood estimator. Apart from its (asymptotic) optimality properties, its popularity is presumably in direct relation with its universality: for the professions mentioned above, and many more, whose prime interest is not the study of sophisticated statistical procedures, it is essential to have at hand a method that is both, easily communicated and applicable to a wide range of model assumptions. A second class of methods incorporates the idea of using as an estimator the value that minimizes some goodness-of-fit measure. To implement this type of estimators, the empirical distribution, quantile or characteristic function is compared to its theoretical counterpart from the underlying parametric model in a suitable distance, and the term is minimized over the parameter space, see Wolfowitz (1957), or Parr (1981) for an early bibliography. These procedures provide some freedom in adapting the estimation method to the intended inferences from the model and they regularly possess good robustness properties (see Parr & Schucany, 1980, as well as Millar, 1981). An example which was discussed recently, and which goes by the name of minimum CRPM estimation, see Gneiting et al. (2005), is tailored to the practice of issuing forecasts: as argued by Gneiting, Balabdaoui & Raftery (2007), a good probabilistic forecast minimizes a (strictly) proper scoring rule such as the “CRPM” (Gneiting & Raftery, 2007), and after constructing a suitable model it appears somewhat more natural to use as an estimator the one that minimizes the scoring rule instead of a classical estimation method like maximum likelihood (for a comparison see Gneiting et al., 2018). As it happens, these rather universal procedures listed above easily run into computational hardships. Just consider that even for “basic” models, density functions can take complicated forms, and distribution or characteristic functions, or even normalization constants, may be nowhere near to an explicit formula. This is where we want to tie on. In a recent work, Betsch & Ebner (2019a) established distributional characterizations that, from a practical point of view, are comparable to the characterization of a probability distribution through its distribution function. Their results, which are given in terms of the derivative of a density function and the density itself, provide explicit formulae that simplify the dependence of the terms on the parameters (even for rather complicated models), and extend characterizations via the zero-bias or equilibrium transformation (Goldstein & Reinert, 1997; Peköz & Röllin, 2011, respectively) that arise in the context of Stein’s method, cf. Chen, Goldstein & Shao (2011). The aim of this work is to investigate these characterizations, which were already used to construct goodness-of-fit tests (see Betsch & Ebner, 2019b, 2020), more closely in the context of parameter estimation. An advantage of the resulting estimators lies in the way the density function of the underlying model appears in the characterization, and thus also in the estimation method. When considering for some (positive) density function p the quotient $\frac{p'}{p}$, the term no longer depends on the integration constant which ensures that the function integrates to one, but only on the functional form of the density. As indicated before, our estimators depend on the underlying model precisely via this quotient, so they are applicable in cases where the normalization constant is unknown. Models of this type occur (though often in discrete settings) in such applied areas as image modelling (using Markov random fields, see Li, 2009) and machine learning, or in any other area where models are complex enough to render the calculation of the normalization constant impractical. For more specific discussions of such applications, we refer to the introduction of the work by Uehara et al. (2019a). The problem was already addressed by Hyvärinen (2005), who set out to find an estimation method which only takes into account the functional form of a density. The approach introduced there goes by the name of “score matching,” and the estimation method involves terms of the form $\frac{p''}{p} - \frac{1}{2} \left(\frac{p'}{p}\right)^2$ and hence does not depend on the normalization constant either. In the univariate case we discuss here, our method provides a good supplement as it contains no second derivatives and may thus be applicable to cases where other methods fail. Also note that several other approaches by Pihlaja,

Gutmann & Hyvärinen (2010), Matsuda & Hyvärinen (2019) and Uehara, Matsuda & Kim (2020) are available. Later on we also discuss noise-contrastive estimation, a concept introduced by Gutmann & Hyvärinen (2010). All these references indicate that statistical inference for non-normalized models is a topic of very recent investigation that also interests researcher in machine learning, a fact which we further allude to at the end of the following section.

In Section 2 we introduce this new class of parameter estimators that are comparable, in their range of applicability in the given setting, to the maximum likelihood and minimum Cramér–von Mises distance estimators (as discussed by Parr & Schucany, 1980, or Parr & De Wet, 1981). We rigorously deal with the existence and measurability of our estimators in Section 3. In Section 4 we provide results on consistency. Thereafter, we give as (normalized) examples the exponential (Section 5), the Rayleigh (Section 7) and the Burr Type XII distribution (Section 8). For each of the three parametric models we compare our new method to classical methods like the maximum likelihood and minimum Cramér–von Mises distance estimator in competitive Monte Carlo simulation studies. The Burr distribution (cf. Burr, 1942; Rodriguez, 1977; Tadikamalla, 1980; Section 6.2 of Kleiber & Kotz, 2003; or Kumar, 2017) as a model is relevant in econometrics, initiated by Singh & Maddala (1976) (see also Schmittlein, 1983), and other areas like engineering, hydrology and quality assurance, see Shah & Gokhale (1993) for corresponding references. However, the parameter estimation is non-trivial and can even cause computational issues. Thus, providing a new estimation method could prove useful in applications. In Section 9 we discuss an exponential-polynomial model for which the normalization constant is intractable, and we compare the new estimators with the score matching and noise-contrastive estimation approaches.

2. THE NEW ESTIMATORS

To be specific, recall that the problem of parameter estimation for continuous, univariate probability distributions presents itself as follows. Consider for $\Theta \subset \mathbb{R}^d$ a parametric family of probability density functions

$$\mathfrak{P}_\Theta = \{p_\vartheta \mid \vartheta \in \Theta\},$$

and let X_1, \dots, X_n be a sample consisting of independent real-valued random variables with a distribution from \mathfrak{P}_Θ , that is, there exists some $\vartheta_0 \in \Theta$ such that X_i has density function p_{ϑ_0} ($X_i \sim p_{\vartheta_0}$, for short) for $i = 1, \dots, n$. Denote with P_ϑ the distribution function corresponding to p_ϑ . The task is to construct an estimator of the unknown ϑ_0 based on X_1, \dots, X_n .

For the construction of our new estimation method, we first recall in a non-technical fashion a famous distributional characterization that can be traced back to Charles Stein, see Chapter VI of Stein (1986). In the more elaborated version of Ley & Swan (2013) it establishes that, given a suitable probability density function p , the distribution of a real-valued random variable X is given through the density function at hand if, and only if,

$$\mathbb{E} \left[f'(X) + \frac{p'(X)}{p(X)} f(X) \right] = 0$$

for a large enough class of suitably chosen test functions f . Motivated by the well-known zero-bias distribution, Betsch & Ebner (2019a) used the above characterization in a recent publication to derive explicit identities which retain the essence of the characterizing property. Indeed, they were able to derive from the Stein characterization that, for a suitable density function p on the positive axis with few technical assumptions (which we adopt below), the

distribution of a positive random variable X (satisfying a weak integrability property) is given through p if, and only if, the distribution function F_X corresponding to X satisfies

$$F_X(t) = \mathbb{E} \left[-\frac{p'(X)}{p(X)} \min\{X, t\} \right], \quad t > 0. \quad (1)$$

As we intent to use this result as a foundation for our estimation method in parametric models for non-negative quantities, assume that the support of each density function in \mathfrak{P}_Θ is $(0, \infty)$. In particular, suppose that each p_ϑ is positive and continuously differentiable on $(0, \infty)$. Also assume that

$$\int_0^\infty |x| |p'_\vartheta(x)| dx < \infty \quad \text{and} \quad \sup_{x>0} \frac{p'_\vartheta(x) \min\{P_\vartheta(x), 1 - P_\vartheta(x)\}}{p_\vartheta^2(x)} < \infty.$$

Moreover, suppose that $\lim_{x \searrow 0} \frac{P_\vartheta(x)}{p_\vartheta(x)} = 0$. These presumptions were made by Betsch & Ebner (2019a) to derive the characterization given above, and they are straight forward to check for most common density functions. Particularly the last condition is exhaustively discussed in Proposition 3.7 of Döbler (2015). Let X be a positive random variable with

$$\mathbb{E} \left[\frac{p'_\vartheta(X)}{p_\vartheta(X)} X \right] < \infty, \quad \vartheta \in \Theta, \quad (2)$$

and define the function

$$\eta(t, \vartheta) = \mathbb{E} \left[-\frac{p'_\vartheta(X)}{p_\vartheta(X)} \min\{X, t\} \right] - F_X(t)$$

for $(t, \vartheta) \in (0, \infty) \times \Theta$. Then the characterization of Betsch & Ebner (2019a), as built up in Equation (1) and as given in their Corollary 3, asserts that X has density function p_ϑ if, and only if, $\eta(t, \vartheta) = 0$ for every $t > 0$. Therefore, if we assume initially that $X \sim p_{\vartheta_0}$ (note that (2) is satisfied by requirement on p_ϑ), then

$$\|\eta(\cdot, \vartheta)\|_{L^q} = 0 \quad \text{if, and only if,} \quad \vartheta = \vartheta_0.$$

Here, $L^q = L^q((0, \infty), \mathcal{B}(0, \infty), w(t) dt)$, $1 \leq q < \infty$, denote the usual L^q -spaces over $(0, \infty)$, w is a positive and integrable weight function, and for $f \in L^q$, $g \in L^{q'}$ ($1/q + 1/q' = 1$)

$$\|f\|_{L^q} = \left(\int_0^\infty |f(t)|^q w(t) dt \right)^{1/q}, \quad \langle f, g \rangle_{L^q} = \int_0^\infty f(t) g(t) w(t) dt$$

are the usual norm and duality in L^q . Thus, with an empirical version

$$\eta_n(t, \vartheta) = -\frac{1}{n} \sum_{j=1}^n \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} \min\{X_j, t\} - \frac{1}{n} \sum_{j=1}^n \mathbb{1}\{X_j \leq t\} \quad (3)$$

of η , based on a sample of independent and identically distributed (i.i.d.) random variables X_1, \dots, X_n with $X_1 \sim p_{\vartheta_0}$, a reasonable estimator for the unknown ϑ_0 is

$$\hat{\vartheta}_{n,q} = \arg \min \{ \|\eta_n(\cdot, \vartheta)\|_{L^q} \mid \vartheta \in \Theta \} \quad \left(= \arg \min \{ \|\eta_n(\cdot, \vartheta)\|_{L^q}^q \mid \vartheta \in \Theta \} \right), \quad (4)$$

that is, we choose $\hat{\vartheta}_{n,q}$ such that $\|\eta_n(\cdot, \hat{\vartheta}_{n,q})\|_{L^q} \leq \|\eta_n(\cdot, \vartheta)\|_{L^q}$ for each $\vartheta \in \Theta$. Heuristically, $\|\eta_n(\cdot, \vartheta)\|_{L^q}$ approximates $\|\eta(\cdot, \vartheta)\|_{L^q}$, so $\hat{\vartheta}_{n,q}$ should provide an estimate for the minimum of $\vartheta \mapsto \|\eta(\cdot, \vartheta)\|_{L^q}$ which coincides with ϑ_0 , the (unique) zero of this function. At this point of course, there arise questions of existence and measurability of such an estimator, and we will handle these questions in full detail in Section 3. Intuitively, one might argue to replace F_X and the empirical distribution function in the definition of η and η_n , respectively, with the theoretical distribution function P_ϑ . However, there is a bit of a technical point involved, and the characterizations by Betsch & Ebner (2019a) do not include results that give a rigorous handle for this slightly (yet decisively) different situation. There are, however, similar characterizations for univariate distribution with other supports than the positive half axis. We allude to that setting in Section 10. Note that the availability of the term $\frac{p'_\vartheta}{p_\vartheta}$ for the model in consideration is rather essential. If this term is not amenable explicitly, it might still be calculable using numerical differentiation (and so η_n and the estimator could be calculated numerically), but it would make it hard to theoretically justify the validity of the conditions on p_ϑ . In our experience, however, the term $\frac{p'_\vartheta}{p_\vartheta}$ is readily available whenever p_ϑ can be differentiated explicitly, and this seems a manageable assumption.

As we have outlined above, our new estimators are eventually based on Stein characterizations which rely on some suitable class of test functions (for an overview in the univariate case, and a record of the vast amount of literature on these identities, see Ley, Reinert & Swan, 2017b). The goal of Betsch & Ebner (2019a) was to derive from these Stein identities new characterizations that no longer involve the classes of test functions. While this approach apparently leads to feasible applications in statistics, other methods are based directly on the Stein characterizations, using Stein discrepancies which gradually become popular in machine learning. The idea in the context of parameter estimation, in heuristic terms, boils down to choosing as a parameter estimator the value which (approximately) minimizes

$$\sup_f \left| \mathbb{E} \left[f'(X) + \frac{p'_\vartheta(X)}{p_\vartheta(X)} f(X) \right] \right|,$$

where the supremum is over all test functions in consideration. By the Stein characterization detailed above, the expectation is 0 for every test function precisely when $\vartheta = \vartheta_0$, as we assume that $X \sim p_{\vartheta_0}$. However, it is not clear how to calculate the supremum in practice taking that the class of test functions is very large. The theory developed around Stein discrepancies has produced different formal methods to evaluate such terms. Other than the fact that they are based on the Stein characterization, the identities derived by Betsch & Ebner (2019a) are not related to the framework of Stein discrepancies, and so it is surprising that merely measuring the difference between the functions in (1) in an L^q -norm, which is what we do to construct our estimators, leads back to so-called feature Stein discrepancies. Indeed, upon defining the “feature” function $\Phi(x, t) = \min\{x, t\}$, $x, t > 0$, and considering the Langevin-Stein operators

$$(\mathcal{J}_\vartheta f)(x, t) = p_\vartheta(x)^{-1} \partial_x (p_\vartheta(x) \cdot f(x, t))$$

as applied to suitable functions $f : (0, \infty)^2 \rightarrow \mathbb{R}$, we obtain

$$\|\eta(\cdot, \vartheta)\|_{L^q}^2 = \left\| \mathbb{E} \left[-\frac{p'_\vartheta(X)}{p_\vartheta(X)} \min\{X, \cdot\} \right] - F_X(\cdot) \right\|_{L^q}^2 = \left\| \int_0^\infty (\mathcal{J}_\vartheta \Phi)(x, \cdot) dP_{\vartheta_0}(x) \right\|_{L^q}^2,$$

which is the right-hand side of equation (1) in the paper by Huggins & Mackey (2018). So by retracing their calculation,

$$\|\eta(\cdot, \vartheta)\|_{L^q}^2 = \Phi\text{SD}_{\mathcal{G},q}^2(P_{\vartheta_0}, P_{\vartheta}) = \sup_{g \in \mathcal{G}_{\Phi,q}} \left| \mathbb{E} \left[\frac{\partial_x(g(X) \cdot p_{\vartheta}(X))}{p_{\vartheta}(X)} \right] \right|^2,$$

where $\mathcal{G}_{\Phi,q}$ is the class of test functions as defined by Huggins & Mackey (2018) (the precise form of which is inessential at this point). This means that we can embed our setting into the framework of these feature Stein discrepancies, as the construction of our estimator cumulating in (4) corresponds to minimizing the quantity at the beginning of this paragraph which sought to motivate these discrepancies. Now, of course, the starting point of our estimation method being the characterizations by Betsch & Ebner (2019a), we already had our estimator at hand explicitly and could choose the feature function accordingly. Still, the fact that both the characterization of Betsch & Ebner (2019a) and the (feature) Stein discrepancy approach (for the above feature function), when translated into an estimation method, lead to the same procedure is remarkable and deems it worthwhile to study the method further, as we were assured that it can be rather hard to find explicit examples for which the Stein discrepancy approach is feasible in practice.

To complete this insightful tour into the realm of Stein discrepancies, we mention some contributions of various solutions to statistical problems based on those discrepancies. In particular, Chwialkowski, Strathmann & Gretton (2016), Liu, Lee & Jordan (2016) and Yang et al. (2018) construct tests of fit, Gorham & Mackey (2015) measure sample quality and Barp et al. (2019) develop estimation methods for non-normalized statistical models.

3. EXISTENCE AND MEASURABILITY

We discuss the measurability properties of η_n and derive an existence result for a measurable version of (approximate) estimators of the type in (4). The result that is central to us in this section can be found in Chapter III of Castaing & Valadier (1977) (see the references therein and Chapter 8 by Cohn, 2013, for further background). Before we summarize these results, recall that a Suslin space is a Hausdorff topological space which is the image of a separable, completely metrizable topological space under a continuous map (for an overview, consult Chapter II of Schwartz, 1973). See also Remark 9 in the Appendix for more information.

Theorem 1. *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space and $(\mathfrak{S}, \mathcal{O}_{\mathfrak{S}})$ a Suslin topological space with Borel- σ -field $\mathcal{B}(\mathfrak{S})$. Assume that Γ maps Ω into the non-empty subsets of \mathfrak{S} , and that*

$$\text{graph}(\Gamma) = \{(\omega, x) \in \Omega \times \mathfrak{S} \mid x \in \Gamma(\omega)\} \in \mathcal{F} \otimes \mathcal{B}(\mathfrak{S}).$$

Then, there exists an $(\mathcal{F}, \mathcal{B}(\mathfrak{S}))$ -measurable map $\hat{\vartheta} : \Omega \rightarrow \mathfrak{S}$ such that $\hat{\vartheta}(\omega) \in \Gamma(\omega)$ for every $\omega \in \Omega$. Additionally, if $\psi : \Omega \times \mathfrak{S} \rightarrow \overline{\mathbb{R}}$ is $(\mathcal{F} \otimes \mathcal{B}(\mathfrak{S}), \overline{\mathcal{B}})$ -measurable, then

$$m(\omega) = \inf_{x \in \Gamma(\omega)} \psi(\omega, x) \quad \text{and} \quad M(\omega) = \sup_{x \in \Gamma(\omega)} \psi(\omega, x)$$

are $(\mathcal{F}, \overline{\mathcal{B}})$ -measurable. Here, $(\overline{\mathbb{R}}, \overline{\mathcal{B}})$ denotes the extended real line with its usual σ -field, and we write \otimes for the product of σ -fields.

To apply Theorem 1, we first have to investigate the measurability properties of η_n . In the setting of Section 2, assume the following regularity condition.

(R1) The map $\Theta \ni \vartheta \mapsto \frac{p'_{\vartheta}(x)}{p_{\vartheta}(x)}$ is continuous for every $x > 0$.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space, which is assumed to underlie all random quantities of the previous and subsequent sections. Notice that the function η_n defined in (3) depends on the random variables X_1, \dots, X_n defined on Ω , hence η_n (as a random quantity) can be understood as a map $\Omega \times (0, \infty) \times \Theta \rightarrow \mathbb{R}$. Exploiting the structure of η_n , we obtain the following lemma. The proof is simple, and the basic thoughts can be found in the Appendix.

Lemma 2. *The map $\eta_n : \Omega \times (0, \infty) \times \Theta \rightarrow \mathbb{R}$ from (3) is $(\mathcal{F} \otimes \mathcal{B}(0, \infty) \otimes \mathcal{B}(\Theta), \mathcal{B}^1)$ -measurable. Moreover, as an element in L^q , $\eta_n : \Omega \times \Theta \rightarrow L^q$ is $(\mathcal{F} \otimes \mathcal{B}(\Theta), \mathcal{B}(L^q))$ -measurable. In particular,*

$$(\omega, \vartheta) \mapsto \|\eta_n(\omega, \cdot, \vartheta)\|_{L^q}$$

is an $(\mathcal{F} \otimes \mathcal{B}(\Theta), \mathcal{B}[0, \infty))$ -measurable mapping.

Similar measurability results hold for $\eta : (0, \infty) \times \Theta \rightarrow \mathbb{R}$. For the remainder of this work assume that

(R2) the parameter space $\emptyset \neq \Theta \subset \mathbb{R}^d$ is a Borel set in \mathbb{R}^d .

As such, Θ is a Suslin topological space (see Proposition 8.2.10 of Cohn, 2013) with the subspace topology induced by \mathbb{R}^d . It is also a metric space with the standard metric in \mathbb{R}^d restricted to Θ . For $n \in \mathbb{N}$, let ε_n be positive random variables such that $\varepsilon_n \rightarrow 0$ \mathbb{P} -almost surely (a.s.), as $n \rightarrow \infty$. Define $\psi_{n,q}(\omega, \vartheta) = \|\eta_n(\omega, \cdot, \vartheta)\|_{L^q}$ which, by Lemma 2, is a product-measurable function from $\Omega \times \Theta$ into $[0, \infty)$. Theorem 1 implies that $m_{n,q}(\omega) = \inf_{\vartheta \in \Theta} \psi_{n,q}(\omega, \vartheta)$ is $(\mathcal{F}, \mathcal{B}^1)$ -measurable. Hence the set-valued function

$$\Gamma_{n,q}(\omega) = \left\{ \vartheta \in \Theta \mid \psi_{n,q}(\omega, \vartheta) \leq m_{n,q}(\omega) + \varepsilon_n(\omega) \right\} \tag{5}$$

has a measurable graph. By construction, $\Gamma_{n,q}$ takes as values only non-empty subsets of Θ . In fact, $\Gamma_{n,q}(\omega)$ is also closed in Θ for every $\omega \in \Omega$, see Remark 11 in the Appendix. Theorem 1 yields the existence of an $(\mathcal{F}, \mathcal{B}(\Theta))$ -measurable map $\hat{\vartheta}_{n,q} : \Omega \rightarrow \Theta$ with $\hat{\vartheta}_{n,q}(\omega) \in \Gamma_{n,q}(\omega)$, which is, by definition of $\Gamma_{n,q}$,

$$\|\eta_n(\omega, \cdot, \hat{\vartheta}_{n,q}(\omega))\|_{L^q} \leq \inf_{\vartheta \in \Theta} \|\eta_n(\omega, \cdot, \vartheta)\|_{L^q} + \varepsilon_n(\omega) \tag{6}$$

for each $\omega \in \Omega$ or, equivalently,

$$\psi_{n,q}(\omega, \hat{\vartheta}_{n,q}(\omega)) \leq \inf_{\vartheta \in \Theta} \psi_{n,q}(\omega, \vartheta) + \varepsilon_n(\omega).$$

Whenever we refer to an estimator that satisfies (4), we mean precisely such an (approximate) measurable version. This settles the existence problem and for our asymptotic studies we have measurability of $\hat{\vartheta}_{n,q}$ at hand.

4. CONSISTENCY

In this section, we investigate the asymptotic behaviour of our estimators. Unfortunately, we cannot apply the general results for minimum distance estimators given by Millar (1984), since a major assumption in that work is that the term in the norm is differentiable (with respect to ϑ) with derivative not depending on ω , that is, in a sense, the parameter and the ‘‘uncertainty’’

have to be separated, which is clearly not the case in our setting. Thus, we need to deal with the empirical process involved.

Assume the setting from Section 2. For brevity, we keep the notation from the previous section, $\psi_{n,q}(\vartheta) (= \psi_{n,q}(\omega, \vartheta)) = \|\eta_n(\omega, \cdot, \vartheta)\|_{L^q}$ and set $\psi_q(\vartheta) = \|\eta(\cdot, \vartheta)\|_{L^q}$. Recall from the construction that $\hat{\vartheta}_{n,q}$ (approximately) minimizes $\psi_{n,q}$ (see (6)), and ϑ_0 is the unique minimum of ψ_q . The heuristic of the consistency statement proven in this section is as follows. If $\psi_{n,q}$ converges to ψ_q in a suitable function space, then the random minimal points $\hat{\vartheta}_{n,q}$ converge to ϑ_0 . In order to establish convergence of $\psi_{n,q}$, we need the functions to be sufficiently smooth in ϑ . In most applications the mapping $\vartheta \mapsto \frac{p'_{\vartheta}(x)}{p_{\vartheta}(x)}$ will be continuously differentiable for every $x > 0$, which can often be used to derive the following regularity condition.

(R3) For each non-empty compact subset K of Θ there exists some $0 < \alpha = \alpha_K < \infty$ and a measurable function $H = H_K : (0, \infty) \rightarrow [0, \infty)$ with $\mathbb{E}[H(X)X] < \infty$ such that

$$\left| \frac{p'_{\vartheta(2)}(x)}{p_{\vartheta(2)}(x)} - \frac{p'_{\vartheta(1)}(x)}{p_{\vartheta(1)}(x)} \right| \leq H(x) |\vartheta(2) - \vartheta(1)|^\alpha, \quad x > 0, \quad \vartheta(1), \vartheta(2) \in K.$$

Now, let $K \neq \emptyset$ be an arbitrary compact subset of Θ . Then on Ω and for $\vartheta(1), \vartheta(2) \in K$, we have

$$\left| \psi_{n,q}(\vartheta(2)) - \psi_{n,q}(\vartheta(1)) \right| \leq |\vartheta(2) - \vartheta(1)|^\alpha \cdot \left(\int_0^\infty w(t) dt \right)^{1/q} \cdot \frac{1}{n} \sum_{j=1}^n H(X_j) X_j$$

with H and α as in (R3). In particular, $K \ni \vartheta \mapsto \psi_{n,q}(\omega, \vartheta)$ is continuous for every $\omega \in \Omega$, and, by Lemma 2, it constitutes a product measurable map. This already implies that $\vartheta \mapsto \psi_{n,q}(\vartheta)$ is a random element of $C(K)^+$ (see Lemma 3.1 of Kallenberg, 2002), the space of continuous functions from K to $[0, \infty)$ which is a complete, separable metric space (endowed with the usual metric that induces the uniform topology). From (R3) it also follows that $K \ni \vartheta \mapsto \psi_q(\vartheta)$ is an element of $C(K)^+$. We can now state the convergence results for $\psi_{n,q}$ that are essential for our consistency proof.

Lemma 3. *In the setting of Section 2, assume that (R1)–(R3) are satisfied. Let $K \neq \emptyset$ be a compact subset of Θ . Then $\psi_{n,q} \rightarrow \psi_q$ in $C(K)^+$ \mathbb{P} -a.s., as $n \rightarrow \infty$. Moreover,*

$$\inf_{\vartheta \in F} \|\eta_n(\cdot, \vartheta)\|_{L^q} = \inf_{\vartheta \in F} \psi_{n,q}(\vartheta) \rightarrow \inf_{\vartheta \in F} \psi_q(\vartheta) = \inf_{\vartheta \in F} \|\eta(\cdot, \vartheta)\|_{L^q}$$

\mathbb{P} -a.s., as $n \rightarrow \infty$, for every non-empty closed subset F of K .

The proof of this lemma is rather technical and deferred to the Appendix. Note that the term $\inf_{\vartheta \in F} \psi_{n,q}(\vartheta)$ is a random variable by Theorem 1 (cf. the measurability of $m_{n,q}$ in the previous section). The following theorem uses the above lemma to establish consistency. In the second statement, we assume that the parameter space Θ is compact, thus rendering Lemma 3 applicable on the whole of Θ , which will turn out essential to prove strong consistency. For most practical purposes this is sufficient, when parameters relevant for modelling in applications can be taken to stem from some (huge) compact set. Note that with this compactness assumption we actually do not need the ε_n -term in (6) since $\psi_{n,q}$ is lower semi-continuous by (R1) and Fatou’s lemma, and thus attains its minimum in Θ . The first statement of the following theorem shows that if the sequence $\hat{\vartheta}_{n,q}$ is already known to be tight, no compactness assumption is needed, but we can only expect weak consistency in general, thus denoting by “ $\xrightarrow{\mathbb{P}}$ ” convergence in probability. After the proof of the theorem, we provide an insight in which cases this is possible (Remark 5).

Theorem 4 (Consistency). *Take the setting from Section 2, let $\psi_{n,q}, \psi_q$ be as above, and consider $\widehat{\vartheta}_{n,q}$ from (6). Further assume that (R1)–(R3) are satisfied.*

- (i) *If $\{\widehat{\vartheta}_{n,q}\}_{n \in \mathbb{N}}$ is tight in Θ , then $\widehat{\vartheta}_{n,q} \xrightarrow{\mathbb{P}} \vartheta_0$, as $n \rightarrow \infty$.*
- (ii) *If Θ is compact, then $\widehat{\vartheta}_{n,q} \rightarrow \vartheta_0$ \mathbb{P} -a.s., as $n \rightarrow \infty$.*

Proof. In the proof of (i) we follow Theorem 3.2.2 from van der Vaart & Wellner (2000), but we adapt the reasoning to our setting, using the measurability properties we established in Section 3, and Lemma 3. For completeness, as well as to prepare the proof of the second result, we give a full proof. We start with a preliminary observation, establishing that the minimum at ϑ_0 is (locally) well separated. If K is a compact subset of Θ , and O an open subset of \mathbb{R}^d which contains ϑ_0 , then

$$0 = \psi_q(\vartheta_0) < \inf_{\vartheta \in K \setminus O} \psi_q(\vartheta). \tag{7}$$

Indeed, if this is not the case, we find a sequence $\vartheta^{(k)} \in K \setminus O$ such that $\psi_q(\vartheta^{(k)}) \rightarrow 0$ as $k \rightarrow \infty$. Since $K \setminus O$ is compact, there exists a subsequence $\{\vartheta^{(k_j)}\}_{j \in \mathbb{N}}$ and some $\vartheta^* \in K \setminus O$ such that $\vartheta^{(k_j)} \rightarrow \vartheta^*$ as $j \rightarrow \infty$. By continuity of ψ_q , it holds that $\psi_q(\vartheta^*) = \lim_{j \rightarrow \infty} \psi_q(\vartheta^{(k_j)}) = 0$, but $K \setminus O \ni \vartheta^* \neq \vartheta_0 \in O$ which is a contradiction to the fact that ϑ_0 is the unique zero of ψ_q .

Now, let $\varepsilon, \delta > 0$. Choose a compact subset $K = K_\delta \subset \Theta$ with $\sup_{n \in \mathbb{N}} \mathbb{P}(\widehat{\vartheta}_{n,q} \notin K) < \delta$, and define $F = K \setminus B_\varepsilon(\vartheta_0)$, where $B_\varepsilon(\vartheta_0)$ denotes the open ball in \mathbb{R}^d of radius ε around ϑ_0 . Applying Lemma 3 and (7) to K and F , together with (6) and the Portmanteau theorem (cf. Theorem 2.1 of Billingsley, 1968), we get

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbb{P}\left(|\widehat{\vartheta}_{n,q} - \vartheta_0| \geq \varepsilon\right) &\leq \limsup_{n \rightarrow \infty} \mathbb{P}\left(\widehat{\vartheta}_{n,q} \in F\right) + \limsup_{n \rightarrow \infty} \mathbb{P}\left(\widehat{\vartheta}_{n,q} \notin K\right) \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{P}\left(\inf_{\vartheta \in F} \psi_{n,q}(\vartheta) \leq \psi_{n,q}(\vartheta_0) + \varepsilon_n\right) + \delta \\ &\leq \mathbb{P}\left(\inf_{\vartheta \in F} \psi_q(\vartheta) \leq \psi_q(\vartheta_0)\right) + \delta \\ &= \delta. \end{aligned}$$

Note that if $F = \emptyset$, the inequality holds trivially. Since both ε and δ were arbitrary, the claim follows. For this first part of the proof, we only needed the convergences from Lemma 3 to be valid in probability. For the following proof of (ii), we rely on the stronger result. The arguments we use are scattered over Section 3 of the work by Sahler (1970). For reasons alluded to in Remark 9, and since that work contains some typos, we provide the adapted arguments. Let $\varepsilon > 0$ and define $\beta_\varepsilon = \inf_{\vartheta \in \Theta \setminus B_\varepsilon(\vartheta_0)} \psi_q(\vartheta)$. By (7), we have $\beta_\varepsilon > 0$. Using the well-known equivalent criterion for almost sure convergence, Lemma 3 gives

$$\lim_{n \rightarrow \infty} \mathbb{P}\left(\bigcup_{k \geq n} \left\{ \sup_{\vartheta \in \Theta \setminus B_\varepsilon(\vartheta_0)} |\psi_{k,q}(\vartheta) - \psi_q(\vartheta)| \geq \frac{\beta_\varepsilon}{2} \right\}\right) = 0.$$

By definition of β_ε this implies

$$\lim_{n \rightarrow \infty} \mathbb{P}\left(\bigcup_{k \geq n} \left\{ \inf_{\vartheta \in \Theta \setminus B_\varepsilon(\vartheta_0)} \psi_{k,q}(\vartheta) \leq \frac{\beta_\varepsilon}{2} \right\}\right) = 0.$$

Moreover, $\psi_{n,q}(\vartheta_0) + \varepsilon_n \rightarrow \psi_q(\vartheta_0) = 0$ \mathbb{P} -a.s., as $n \rightarrow \infty$, and thus

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\bigcup_{k \geq n} \left\{ \left| \psi_{k,q}(\vartheta_0) + \varepsilon_k \right| \geq \frac{\beta_\varepsilon}{2} \right\} \right) = 0.$$

Putting everything together,

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \mathbb{P} \left(\bigcup_{k \geq n} \left\{ \left| \hat{\vartheta}_{k,q} - \vartheta_0 \right| \geq \varepsilon \right\} \right) \\ & \leq \limsup_{n \rightarrow \infty} \mathbb{P} \left(\bigcup_{k \geq n} \left\{ \inf_{\vartheta \in \Theta \setminus B_\varepsilon(\vartheta_0)} \psi_{k,q}(\vartheta) \leq \psi_{k,q}(\vartheta_0) + \varepsilon_k \right\} \right) \\ & \leq \limsup_{n \rightarrow \infty} \left(\mathbb{P} \left(\bigcup_{k \geq n} \left\{ \inf_{\vartheta \in \Theta \setminus B_\varepsilon(\vartheta_0)} \psi_{k,q}(\vartheta) \leq \frac{\beta_\varepsilon}{2} \right\} \right) + \mathbb{P} \left(\bigcup_{k \geq n} \left\{ \psi_{k,q}(\vartheta_0) + \varepsilon_k \geq \frac{\beta_\varepsilon}{2} \right\} \right) \right) \\ & = 0, \end{aligned}$$

that is, $\hat{\vartheta}_{n,q} \rightarrow \vartheta_0$ \mathbb{P} -a.s., as $n \rightarrow \infty$. ■

Remark 5 (A priori tightness of the estimators). *We provide a tool for proving tightness of the estimators before having established consistency, which we can use in Theorem 4 to get consistency even for unbounded parameter spaces. The statement essentially yields that if $\psi_{n,q}$ is strictly convex, $\{\hat{\vartheta}_{n,q}\}_{n \in \mathbb{N}}$ is tight. More precisely, suppose that conditions (R1)–(R3) hold. Let Θ be convex with $\vartheta_0 \in \Theta^\circ$, the interior of Θ . Further, let $\psi_{n,q}$ be strictly convex (almost surely). Then the sequence of estimators $\hat{\vartheta}_{n,q}$ is tight in Θ . The proof is straight-forward and some hints are given in exercise problem 4 in Section 3.2 of van der Vaart & Wellner (2000) (for more details, find the proof in the Appendix).*

5. EXAMPLE: THE EXPONENTIAL DISTRIBUTION

Let $\Theta = (0, \infty)$ and $p_\vartheta(x) = \vartheta \exp(-\vartheta x)$, $x > 0$. This trivially is an admissible class of density functions. Moreover, let $\vartheta_0 \in \Theta$, $X \sim p_{\vartheta_0}$, and take a sample X_1, \dots, X_n of i.i.d. copies of X . An easy calculation gives

$$\begin{aligned} (\psi_q(\vartheta))^q &= \int_0^\infty \left| \mathbb{E}[\vartheta \min\{X, t\}] - (1 - \exp(-\vartheta_0 t)) \right|^q w(t) dt \\ &= \left| \frac{\vartheta}{\vartheta_0} - 1 \right|^q \int_0^\infty \left| 1 - \exp(-\vartheta_0 t) \right|^q w(t) dt, \end{aligned}$$

which nicely illustrates that ϑ_0 is indeed the unique zero of this functions. For the particular choice of weight $w(t) = \exp(-at)$, $t > 0$, with some tuning parameter $a > 0$, and in the case $q = 2$, straight-forward calculations give

$$(\psi_{n,2}(\vartheta))^2 = \vartheta^2 \Psi_n^{(1)} + \vartheta \Psi_n^{(2)} + \Psi_n^{(3)},$$

where

$$\begin{aligned} \Psi_n^{(1)} &= \frac{2}{a^3} + \frac{2}{a^2 n^2} \sum_{j=1}^n e^{-aX_{(j)}} \left(X_{(j)}(-n + j - 1) - \frac{1}{a}(2n - 2j + 1) \right) \\ &\quad - \frac{2}{a^2 n^2} \sum_{1 \leq j < k \leq n} X_{(j)} e^{-aX_{(k)}}, \end{aligned}$$

$$\Psi_n^{(2)} = \frac{2}{an^2} \sum_{j=1}^n e^{-aX_{(j)}} \left(X_{(j)}(-n + j - 1) - \frac{1}{a}(n - 2j + 1) \right) - \frac{2}{an^2} \sum_{1 \leq j < k \leq n} X_{(j)} e^{-aX_{(k)}},$$

$$\Psi_n^{(3)} = \frac{1}{an^2} \sum_{j=1}^n e^{-aX_{(j)}}(2j - 1),$$

and $X_{(1)} < \dots < X_{(n)}$ is the ordered sample. Using that $e^{-aX_{(k)}} \leq e^{-aX_{(j)}}$ \mathbb{P} -a.s. for $j < k$, we obtain

$$\Psi_n^{(1)} \geq \frac{2}{a^3 n^2} \sum_{j=1}^n (2n - 2j + 1) \left(1 - (1 + aX_{(j)})e^{-aX_{(j)}} \right),$$

and since $1 + aX_{(j)} < e^{aX_{(j)}}$ \mathbb{P} -a.s., we have $\Psi_n^{(1)} > 0$ almost surely. Therefore, $(\psi_{n,2})^2$ is strictly convex (almost surely), and has a unique minimum. By Remark 5 and Theorem 4 (i), the estimator

$$\begin{aligned} \hat{\vartheta}_{n,2}^{(a)} &= \arg \min \{ \psi_{n,2}(\vartheta) \mid \vartheta > 0 \} = \arg \min \{ (\psi_{n,2}(\vartheta))^2 \mid \vartheta > 0 \} \\ &= \arg \min \{ \vartheta^2 \Psi_n^{(1)} + \vartheta \Psi_n^{(2)} + \Psi_n^{(3)} \mid \vartheta > 0 \} \end{aligned}$$

is consistent for ϑ_0 (over the whole of Θ). Note that we have not made the dependence of $\Psi_n^{(1)}$, $\Psi_n^{(2)}$, and $\Psi_n^{(3)}$ on “ a ” explicit to prevent overloading the notation. With a similar argument as above, we may show that $\Psi_n^{(2)} < 0$ almost surely, thus we can calculate $\hat{\vartheta}_{n,2}^{(a)}$ explicitly as

$$\hat{\vartheta}_{n,2}^{(a)} = -\frac{\Psi_n^{(2)}}{2\Psi_n^{(1)}}.$$

To provide insight on the performance of this estimator, we compare it with the maximum likelihood estimator and the minimizer of the mean squared error (MSE) (for $n \geq 3$) which are given as

$$\hat{\vartheta}_n^{ML} = \left(\frac{1}{n} \sum_{j=1}^n X_j \right)^{-1} \quad \text{and} \quad \hat{\vartheta}_n^{MSE} = \left(\frac{1}{n-2} \sum_{j=1}^n X_j \right)^{-1},$$

respectively, as well as with the minimum Cramér–von Mises distance estimator discussed in the introduction, namely

$$\begin{aligned} \hat{\vartheta}_n^{CvM} &= \arg \min \left\{ \int_0^\infty \left(\hat{F}_n(t) - P_\vartheta(t) \right)^2 dP_\vartheta(t) \mid \vartheta > 0 \right\} \\ &= \arg \min \left\{ \frac{1}{n} \sum_{j=1}^n \left[\exp(-2\vartheta X_{(j)}) + \exp(-\vartheta X_{(j)}) \cdot \left(\frac{2j-1}{n} - 2 \right) \right] \mid \vartheta > 0 \right\}, \end{aligned}$$

where $P_\vartheta(x) = 1 - \exp(-\vartheta x)$, $x > 0$, denotes the distribution function of the exponential distribution, and where \hat{F}_n is the empirical distribution function of X_1, \dots, X_n . For this comparison we simulate (for fixed values of n and ϑ_0) $D = 100,000$ samples of size n from an exponential

TABLE 1: Approximated biases calculated with 100,000 exponentially distributed Monte Carlo samples.

ϑ_0	n	$\hat{\vartheta}_n^{ML}$	$\hat{\vartheta}_n^{MSE}$	$\hat{\vartheta}_n^{CvM}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$
0.5	10	0.0557	-0.0554	0.0510	0.0428	0.0376	0.0333	0.0302	0.0291
	25	0.0212	-0.0205	0.0187	0.0161	0.0144	0.0129	0.0118	0.0114
	50	0.0098	-0.0106	0.0089	0.0075	0.0067	0.0062	0.0057	0.0055
	100	0.005	-0.0051	0.0045	0.0039	0.0035	0.0032	0.0030	0.0029
	200	0.0023	-0.0027	0.0020	0.0018	0.0016	0.0015	0.0014	0.0013
2	10	0.2193	-0.2245	0.2004	0.2011	0.1871	0.1680	0.1476	0.1371
	25	0.0830	-0.0836	0.0746	0.0754	0.0701	0.0634	0.0569	0.0538
	50	0.0398	-0.0418	0.0345	0.0358	0.0332	0.0298	0.0263	0.0246
	100	0.0191	-0.0213	0.0165	0.0171	0.0158	0.0142	0.0126	0.0118
	200	0.0095	-0.0106	0.0074	0.0084	0.0077	0.0067	0.0057	0.0052
5	10	0.5437	-0.5651	0.4863	0.5238	0.5059	0.4753	0.4303	0.3997
	25	0.2102	-0.2066	0.1832	0.2015	0.1940	0.1818	0.1649	0.1540
	50	0.1048	-0.0994	0.0923	0.1004	0.0967	0.0908	0.0829	0.0779
	100	0.0520	-0.0491	0.0440	0.0496	0.0477	0.0446	0.0404	0.0378
	200	0.0264	-0.0238	0.0224	0.0253	0.0243	0.0229	0.0209	0.0196
10	10	1.1230	-1.1016	1.0316	1.1028	1.0837	1.0484	0.9885	0.9401
	25	0.4177	-0.4157	0.3719	0.4089	0.4008	0.3863	0.3628	0.3448
	50	0.2041	-0.2040	0.1826	0.1996	0.1955	0.1883	0.1768	0.1681
	100	0.0991	-0.1029	0.0873	0.0967	0.0945	0.0908	0.0848	0.0804
	200	0.0556	-0.0450	0.0483	0.0544	0.0533	0.0513	0.0483	0.0460

distribution with parameter ϑ_0 , calculate the values of the estimator for each sample yielding values $\hat{\vartheta}_1, \dots, \hat{\vartheta}_D$, and approximate the bias and MSE via

$$\frac{1}{D} \sum_{k=1}^D (\hat{\vartheta}_k - \vartheta_0) \quad \text{and} \quad \frac{1}{D} \sum_{k=1}^D (\hat{\vartheta}_k - \vartheta_0)^2$$

for each of the above estimators. We perform all simulations with Python 3.7.2 (as provided by the Python Software Foundation, <https://www.python.org>, accessed 28 August 2019). For the minimization required to calculate the minimum Cramér–von Mises distance estimator, we choose as initial value the maximum likelihood estimator and use a sequential least squares programming method (“SLSQP”) (cf. Kraft, 1988) implemented in the “optimize.minimize” function of the Python module “scipy,” see Virtanen. P. et al., (2020). Tables 1 and 2 contain the results for the bias and MSE values.

As for the biases, the maximum likelihood estimator and the minimum MSE estimator perform almost identically in terms of the absolute bias, and the minimum Cramér–von Mises distance estimator has a slight edge. Our new estimator outperforms all other methods (virtually) uniformly. More precisely, it seems as if for larger tuning parameters “ a ” the bias decreases. We

TABLE 2: Approximated MSE calculated with 100,000 exponentially distributed Monte Carlo samples.

ϑ_0	n	$\hat{\vartheta}_n^{ML}$	$\hat{\vartheta}_n^{MSE}$	$\hat{\vartheta}_n^{CvM}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$
0.5	10	0.0416	0.0277	0.0593	0.0409	0.0428	0.0496	0.0662	0.0837
	25	0.0123	0.0105	0.0162	0.0127	0.0138	0.0167	0.0229	0.0294
	50	0.0055	0.0051	0.0072	0.0058	0.0064	0.0078	0.0109	0.0140
	100	0.0026	0.0025	0.0034	0.0028	0.0031	0.0038	0.0053	0.0069
	200	0.0013	0.0013	0.0017	0.0014	0.0015	0.0019	0.0026	0.0034
2	10	0.6645	0.4449	0.9504	0.6569	0.6525	0.6537	0.6845	0.7346
	25	0.1949	0.1661	0.2573	0.1942	0.1952	0.2006	0.2184	0.2403
	50	0.0887	0.0821	0.1165	0.0889	0.0898	0.0932	0.1029	0.1141
	100	0.0418	0.0402	0.0549	0.0420	0.0426	0.0444	0.0493	0.0549
	200	0.0205	0.0201	0.0269	0.0207	0.0210	0.0220	0.0244	0.0272
5	10	4.0739	2.7374	5.6848	4.0529	4.0350	4.0092	3.9977	4.0335
	25	1.2302	1.0465	1.621	1.2272	1.2259	1.2284	1.2493	1.2842
	50	0.5522	0.5087	0.7246	0.5518	0.5524	0.5561	0.5706	0.5908
	100	0.2635	0.2529	0.3445	0.2636	0.2641	0.2665	0.2745	0.2850
	200	0.1295	0.1268	0.1692	0.1296	0.1299	0.1313	0.1355	0.1410
10	10	16.8106	11.1652	23.5189	16.7647	16.7219	16.6460	16.5344	16.4779
	25	4.8850	4.1598	6.4565	4.8785	4.8739	4.8702	4.8831	4.9188
	50	2.2069	2.0371	2.8967	2.2053	2.2048	2.2069	2.2213	2.2464
	100	1.0473	1.007	1.3747	1.0471	1.0474	1.0497	1.0594	1.0740
	200	0.5126	0.5014	0.6658	0.5127	0.513	0.5144	0.5198	0.5274

will show, however, that this observation is not correct in that generality. The results for the MSE reveal that the minimum MSE estimator is the best method with respect to this measure of quality, which is no surprise as it is constructed to minimize the MSE. For sample size $n = 10$ the superiority is particularly obvious, but for larger samples, the maximum likelihood estimator is only slightly worse. Our new estimator shows almost identical results (for $a = 0.25$) as the maximum likelihood estimator, undermining that the method is sound and powerful. In contrast to the observation with the bias values, the MSE appears to increase with “ a .” This nicely illustrates the variance-bias trade-off commonly observed in the context of estimation problems.

6. THE CASE $a \rightarrow \infty$

As discussed previously, the simulation results for the exponential distribution somewhat indicate that as the tuning parameter “ a ” grows, the bias decreases while the MSE increases. Interestingly, we can lay observations for $a \rightarrow \infty$ on a rigorous theoretical basis. To be precise, observe the following general result.

Theorem 6. *Consider the setting from Section 2 with weight function $w(t) = e^{-at}$. For the quantity $\psi_{n,q}(\vartheta, a) = \psi_{n,q}(\vartheta) = \|\eta_n(\cdot, \vartheta)\|_{L^q}$ from the end of Section 3, we make the dependence*

on the tuning parameter “ a ” explicit. Then

$$\lim_{a \rightarrow \infty} a^{q+1} (\psi_{n,q}(\vartheta, a))^q = \Gamma(q + 1) \left| \frac{1}{n} \sum_{j=1}^n \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} \right|^q,$$

on a set of measure one, where $\Gamma(\cdot)$ denotes the Gamma function.

The proof consists of an almost trivial application of an Abelian theorem for the Laplace transform, see p. 182 of Widder (1959), or the work by Baringhaus, Gürtler & Henze (2000). Since $a, q > 0$, the functions $\psi_{n,q}(\vartheta)$ and $a^{q+1} (\psi_{n,q}(\vartheta))^q$ attain their minimum in the same point. Thus, in the limit $a \rightarrow \infty$, our procedure essentially yields as an estimators the minimizer of the quantity

$$\left| \frac{1}{n} \sum_{j=1}^n \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} \right|^q.$$

In the situation of the exponential distribution as discussed in Section 5, the result reduces to $\lim_{a \rightarrow \infty} a^3 (\psi_{n,2}(\vartheta, a))^2 = 2\vartheta^2$, so in the limit $a \rightarrow \infty$, the procedure will choose $\hat{\vartheta} = 0 \notin \Theta$ as the estimator, which leads to a bias of $-\vartheta_0$ and an MSE of ϑ_0^2 . The observation from the simulations is, therefore, not universal. An example for which the limit in Theorem 6 is less trivial is the Rayleigh distribution.

7. EXAMPLE: THE RAYLEIGH DISTRIBUTION

Let $\Theta = (0, \infty)$ and take the density function of the Rayleigh distribution with parameter $\vartheta \in \Theta$,

$$p_\vartheta(x) = \frac{x}{\vartheta^2} \exp\left(-\frac{x^2}{2\vartheta^2}\right), \quad x > 0.$$

It is easy to check that the Rayleigh density satisfies all regularity conditions stated throughout the work, and that we have $\frac{p'_\vartheta(x)}{p_\vartheta(x)} = \frac{1}{x} - \frac{x}{\vartheta^2}$. The limit in Theorem 6 thus takes the form

$$\Gamma(q + 1) \left| \frac{1}{n} \sum_{j=1}^n \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} \right|^q = \Gamma(q + 1) \left| \frac{1}{n} \sum_{j=1}^n \left(\frac{1}{X_j} - \frac{X_j}{\vartheta^2} \right) \right|^q,$$

where X_1, \dots, X_n are i.i.d. random variables which follow the Rayleigh law, $X_1 \sim p_{\vartheta_0}$, for some unknown scale parameter $\vartheta_0 \in \Theta$. In the case $q = 2$, it is easy to calculate that the minimum of the above function over $\vartheta > 0$ is given through

$$\hat{\vartheta}_n^{AM} = \sqrt{\frac{\frac{1}{n} \sum_{j=1}^n X_j}{\frac{1}{n} \sum_{j=1}^n \frac{1}{X_j}}}.$$

Strikingly, this asymptotically derived moment-type estimator is itself consistent for ϑ_0 , as

$$\hat{\vartheta}_n^{AM} = \sqrt{\frac{\frac{1}{n} \sum_{j=1}^n X_j}{\frac{1}{n} \sum_{j=1}^n \frac{1}{X_j}}} \rightarrow \sqrt{\frac{\mathbb{E}[X_1]}{\mathbb{E}\left[\frac{1}{X_1}\right]}} = \sqrt{\frac{\sqrt{\frac{\pi}{2}} \vartheta_0}{\sqrt{\frac{\pi}{2}} \cdot \frac{1}{\vartheta_0}}} = \vartheta_0$$

\mathbb{P} -a.s., as $n \rightarrow \infty$, where we used the law of large numbers, as well as the fact that X_1, \dots, X_n all follow the Rayleigh distribution with parameter ϑ_0 . We compare this estimator with other methods. Among them is our new estimator

$$\hat{\vartheta}_{n,2}^{(a)} = \arg \min \{ (\psi_{n,2}(\vartheta))^2 \mid \vartheta > 0 \} = \arg \min \{ \vartheta^{-4} \tilde{\Psi}_n^{(1)} + \vartheta^{-2} \tilde{\Psi}_n^{(2)} + \tilde{\Psi}_n^{(3)} \mid \vartheta > 0 \},$$

where

$$\begin{aligned} \tilde{\Psi}_n^{(1)} &= \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left[X_{(j)} X_{(k)} \cdot \frac{2}{a^3} (1 - e^{-aX_{(j)}}) - \frac{X_{(j)}^2 X_{(k)}}{a^2} (e^{-aX_{(j)}} + e^{-aX_{(k)}}) \right] \\ &\quad + \frac{1}{n^2} \sum_{j=1}^n \left[\frac{2X_{(j)}^2}{a^3} (1 - e^{-aX_{(j)}}) - \frac{2X_{(j)}^3}{a^2} e^{-aX_{(j)}} \right], \end{aligned}$$

$$\begin{aligned} \tilde{\Psi}_n^{(2)} &= \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left[\frac{X_{(j)}^2 e^{-aX_{(k)}}}{a} \left(\frac{1}{aX_{(k)}} - 1 \right) + \frac{X_{(j)} e^{-aX_{(j)}}}{a} \left(\frac{X_{(j)}}{aX_{(k)}} - X_{(k)} \right) \right. \\ &\quad \left. - \frac{2}{a^3} (1 - e^{-aX_{(j)}}) \left(\frac{X_{(k)}}{X_{(j)}} + \frac{X_{(j)}}{X_{(k)}} \right) \right] \\ &\quad + \frac{1}{n^2} \sum_{j=1}^n \left[\frac{2e^{-aX_{(j)}}}{a} \cdot X_{(j)} \left(\frac{2j}{a} - X_{(j)} \right) - \frac{4}{a^3} (1 - e^{-aX_{(j)}}) \right], \end{aligned}$$

and

$$\begin{aligned} \tilde{\Psi}_n^{(3)} &= \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left[\frac{X_{(j)}}{aX_{(k)}} e^{-aX_{(j)}} + \frac{2}{a^3 X_{(j)} X_{(k)}} (1 - e^{-aX_{(j)}}) \right] \\ &\quad + \frac{1}{n^2} \sum_{j=1}^n \left[\frac{2}{a^3 X_{(j)}^2} (1 - e^{-aX_{(j)}}) + \frac{e^{-aX_{(j)}}}{a} \left(4j - 1 - \frac{2}{aX_{(j)}} (2j - 1) \right) \right], \end{aligned}$$

and $X_{(1)} < \dots < X_{(n)}$ denotes the ordered sample. It is easily seen that if both $\tilde{\Psi}_n^{(1)} > 0$ and $\tilde{\Psi}_n^{(2)} < 0$ \mathbb{P} -a.s., then the minimum can be calculated explicitly as

$$\hat{\vartheta}_{n,2}^{(a)} = \sqrt{-\frac{2\tilde{\Psi}_n^{(1)}}{\tilde{\Psi}_n^{(2)}}},$$

and indeed, using that $e^{-aX_{(k)}} < e^{-aX_{(j)}}$ and $1 - e^{-aX_{(j)}} - aX_{(j)}e^{-aX_{(j)}} > 0$ \mathbb{P} -a.s., we have

$$\begin{aligned} \tilde{\Psi}_n^{(1)} &> \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \frac{2X_{(j)} X_{(k)}}{a^3} (1 - e^{-aX_{(j)}} - aX_{(j)} e^{-aX_{(j)}}) \\ &\quad + \frac{1}{n^2} \sum_{j=1}^n \frac{2X_{(j)}^2}{a^3} (1 - e^{-aX_{(j)}} - aX_{(j)} e^{-aX_{(j)}}) \\ &> 0 \quad \mathbb{P} - \text{a.s.}, \end{aligned}$$

TABLE 3: Approximated biases calculated with 100,000 Rayleigh-distributed Monte Carlo samples.

ϑ_0	n	$\hat{\vartheta}_n^{ML}$	$\hat{\vartheta}_n^{Mom}$	$\hat{\vartheta}_n^{AM}$	$\hat{\vartheta}_n^{CvM}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$
0.5	10	-0.0061	0.0001	0.0269	0.0029	-0.0050	-0.0040	-0.0023	0.0006	0.0028
	25	-0.0025	0.0001	0.0140	0.0012	-0.0020	-0.0016	-0.0009	0.0002	0.0011
	50	-0.0011	0.0001	0.0085	0.0007	-0.0009	-0.0007	-0.0003	0.0002	0.0007
	100	-0.0006	0	0.0047	0.0003	-0.0005	-0.0004	-0.0002	0.0001	0.0003
	200	-0.0004	0	0.0027	0.0001	-0.0003	-0.0002	-0.0002	0	0.0001
2	10	-0.0246	0.0001	0.1074	0.0107	-0.0093	0.0019	0.0189	0.0434	0.0598
	25	-0.0097	0	0.0549	0.0039	-0.0036	0.0007	0.0073	0.0169	0.0237
	50	-0.0057	-0.0007	0.0320	0.0013	-0.0025	-0.0003	0.0030	0.0079	0.0114
	100	-0.0026	-0.0001	0.0190	0.0008	-0.0011	0	0.0017	0.0041	0.0058
	200	-0.0015	-0.0003	0.0104	0	-0.0007	-0.0002	0.0005	0.0016	0.0025
5	10	-0.0624	-0.0009	0.2642	0.0255	0.0156	0.0640	0.1293	0.1926	0.2199
	25	-0.0245	-0.0002	0.1388	0.0097	0.0063	0.0251	0.0519	0.0817	0.0973
	50	-0.0132	-0.0002	0.0848	0.0049	0.0030	0.0129	0.0270	0.0432	0.0523
	100	-0.0059	0	0.0477	0.0021	0.0016	0.0062	0.0128	0.0206	0.0253
	200	-0.0028	0.0002	0.0279	0.0010	0.0010	0.0033	0.0066	0.0106	0.0129
10	10	-0.1248	-0.0004	0.5383	0.0537	0.1302	0.2617	0.3919	0.4777	0.5055
	25	-0.0565	-0.0076	0.2699	0.0123	0.0430	0.0965	0.1564	0.2074	0.2293
	50	-0.0261	-0.0021	0.1582	0.0083	0.0225	0.0480	0.0783	0.1073	0.1214
	100	-0.0109	0.0013	0.0979	0.0057	0.0138	0.0272	0.0430	0.0586	0.0671
	200	-0.0077	-0.0010	0.0545	0.0011	0.0057	0.0128	0.0207	0.0289	0.0334

and with similar thoughts, $\tilde{\Psi}_n^{(2)} < 0$ \mathbb{P} -a.s. Additionally, we consider the maximum likelihood estimator and a moment estimator, which are given as

$$\hat{\vartheta}_n^{ML} = \sqrt{\frac{1}{2n} \sum_{j=1}^n X_j^2} \quad \text{and} \quad \hat{\vartheta}_n^{Mom} = \sqrt{\frac{2}{\pi}} \cdot \frac{1}{n} \sum_{j=1}^n X_j,$$

respectively. Note in particular that the moment estimator is unbiased and we can expect it to outperform the other estimators in this regard. Finally, we include the minimum Cramér–von Mises distance estimator given through

$$\hat{\vartheta}_n^{CvM} = \arg \min \left\{ \frac{1}{n} \sum_{j=1}^n \left[\left(\frac{2j-1}{n} - 2 \right) \exp \left(-\frac{X_{(j)}^2}{2\vartheta^2} \right) + \exp \left(-\frac{-X_{(j)}^2}{\vartheta^2} \right) \right] \mid \vartheta > 0 \right\},$$

where we solve the minimization numerically via a sequential least squares programming method as in the case of the exponential distribution in Section 5, using as initial value the maximum

TABLE 4: Approximated MSE calculated with 100,000 Rayleigh-distributed Monte Carlo samples.

ϑ_0	n	$\hat{\vartheta}_n^{ML}$	$\hat{\vartheta}_n^{Mom}$	$\hat{\vartheta}_n^{AM}$	$\hat{\vartheta}_n^{CvM}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$
0.5	10	0.0061	0.0067	0.0135	0.0082	0.0062	0.0062	0.0064	0.0068	0.0072
	25	0.0025	0.0027	0.0061	0.0033	0.0025	0.0025	0.0025	0.0027	0.0028
	50	0.0013	0.0014	0.0033	0.0017	0.0013	0.0013	0.0013	0.0013	0.0014
	100	0.0006	0.0007	0.0019	0.0008	0.0006	0.0006	0.0006	0.0007	0.0007
	200	0.0003	0.0003	0.001	0.0004	0.0003	0.0003	0.0003	0.0003	0.0004
2	10	0.0992	0.1089	0.2185	0.1316	0.1030	0.1090	0.1248	0.1570	0.1800
	25	0.0401	0.0438	0.0977	0.0527	0.0412	0.0433	0.04880	0.0601	0.0690
	50	0.0199	0.0219	0.0543	0.0264	0.0205	0.0215	0.0242	0.0297	0.0341
	100	0.0100	0.0109	0.0300	0.0130	0.0102	0.0106	0.0119	0.0146	0.0167
	200	0.0050	0.0055	0.0167	0.0065	0.0051	0.0053	0.0060	0.0073	0.0083
5	10	0.6205	0.6827	1.3695	0.8271	0.7057	0.8359	1.0635	1.2775	1.3473
	25	0.2521	0.2760	0.6122	0.3314	0.2803	0.3258	0.4097	0.5088	0.5566
	50	0.1250	0.1371	0.3398	0.1648	0.1385	0.1606	0.2015	0.2529	0.2811
	100	0.0627	0.0684	0.1876	0.0819	0.0688	0.0793	0.0989	0.1242	0.1392
	200	0.0311	0.0341	0.1039	0.0409	0.0343	0.0395	0.0491	0.0618	0.0695
10	10	2.4749	2.7278	5.4528	3.3202	3.3485	4.2534	5.0993	5.4722	5.5230
	25	0.9966	1.0933	2.4305	1.3171	1.2922	1.6219	2.0109	2.3024	2.3989
	50	0.5000	0.5445	1.3528	0.6504	0.6342	0.7926	0.9955	1.1792	1.2600
	100	0.2499	0.2735	0.7504	0.3290	0.3179	0.3959	0.4966	0.5962	0.6473
	200	0.1248	0.1364	0.4200	0.1637	0.1579	0.1961	0.2470	0.2999	0.3288

likelihood estimator. The execution of the comparison is as in the example on the exponential distribution, and the results are displayed in Tables 3 and 4.

Apparently, the moment estimator $\hat{\vartheta}_n^{Mom}$ outperforms the other estimators with respect to the bias values, while the maximum likelihood estimator $\hat{\vartheta}_n^{ML}$ gets the smallest MSE. The estimator we obtained via the limit results from the previous section seems sound in itself but is completely negligible compared to the other methods. In terms of bias, the minimum Cramér–von Mises distance estimator is preferable to the maximum likelihood method, and both are outdone by our new estimator, which even keeps up with the unbiased moment estimator for the smaller values of the parameter ϑ_0 . Notice that the maximum likelihood and moment estimator tend to underestimate the parameter, while the other procedures tend to a slight overestimation. As for the MSE, the moment estimator and our new method perform similarly and follow the maximum likelihood estimator closely. The minimum Cramér–von Mises distance estimator is a bit behind. To summarize, the maximum likelihood and moment estimator for the Rayleigh parameter are both simple and very convincing, but the newly proposed method keeps up (for suitably chosen tuning parameter) and appears to find a good compromise between bias and MSE. The only graver weakness shows for the large parameter value $\vartheta_0 = 10$ and small sample sizes $n = 10, 25$.

8. EXAMPLE: THE BURR TYPE XII DISTRIBUTION

Consider the density function of the Burr distribution,

$$p_{\vartheta}(x) = c k x^{c-1} (1 + x^c)^{-k-1}, \quad x > 0,$$

where $\vartheta = (c, k) \in (0, \infty)^2 = \Theta$. It is not exactly trivial, but still straight-forward, to prove that this is an admissible distribution in terms of the setting in Section 2 (see also Betsch & Ebner, 2019a) and the regularity conditions (R1)–(R3). With $q = 2$ and weight $w(t) = e^{-at}$, $a > 0$, the function $\psi_{n,2}(\vartheta) = \|\eta_n(\cdot, \vartheta)\|_{L^2}$ from Section 3 (see also Section 2) can be calculated explicitly as

$$\begin{aligned} (\psi_{n,2}(\vartheta))^2 &= \frac{2}{n^2} \sum_{1 \leq j < \ell \leq n} \left\{ A_{(\ell)}(c, k) \left[\frac{2A_{(j)}(c, k)}{a^3} (1 - e^{-aX_{(j)}}) \right. \right. \\ &\quad \left. \left. + \frac{B_{(j)}(c, k)}{a^2} (e^{-aX_{(j)}} + e^{-aX_{(\ell)}}) \right. \right. \\ &\quad \left. \left. + \frac{c-2}{a^2} e^{-aX_{(j)}} - \frac{X_{(j)}}{a} e^{-aX_{(j)}} \right] + \frac{B_{(j)}(c, k)}{a} e^{-aX_{(\ell)}} \right\} \\ &+ \frac{1}{n^2} \sum_{j=1}^n \left\{ (A_{(j)}(c, k))^2 \left(-\frac{2X_{(j)}}{a^2} e^{-aX_{(j)}} - \frac{2}{a^3} e^{-aX_{(j)}} + \frac{2}{a^3} \right) \right. \\ &\quad \left. + \frac{2(j-1)c}{a^2} A_{(j)}(c, k) e^{-aX_{(j)}} + \frac{2B_{(j)}(c, k)}{a} e^{-aX_{(j)}} \right\} \\ &+ \frac{2c}{an^2} \sum_{j=1}^n j e^{-aX_{(j)}} - \frac{1}{an^2} \sum_{j=1}^n e^{-aX_{(j)}}, \end{aligned}$$

where

$$A_{(j)}(c, k) = c(k+1) \frac{X_{(j)}^{c-1}}{1 + X_{(j)}^c} - \frac{c-1}{X_{(j)}}, \quad B_{(j)}(c, k) = -c(k+1) \frac{X_{(j)}^c}{1 + X_{(j)}^c},$$

and where $X_{(1)} < \dots < X_{(n)}$ denotes the ordered sample. Our estimator $\hat{\vartheta}_{n,2}^{(a)} = (\hat{c}_n^{(a)}, \hat{k}_n^{(a)})$, as defined in (4), can be calculated as the minimizer of the above function over Θ . We use the “L-BFGS-B”-method (L-BFGS-B algorithm, see Byrd et al., 1995; Zhu et al., 1997) implemented in the “optimize.minimize” function of “scipy” to solve the minimization numerically, using (1, 1) as initial values. (Note that in preliminary simulations we have tried several other optimization routines, like a truncated Newton algorithm or the “SLSQP” from previous sections, but the “L-BFGS-B”-method appeared to be the most reliable for our purpose.) As competitors to our estimator we consider the maximum likelihood estimator with implementation as suggested by Shah & Gokhale (1993) (for a different algorithm, see Wingo, 1983). More precisely we use the Newton–Raphson method (with initial value $c = 1$) to find the root

$$\frac{n}{c} + \sum_{j=1}^n \log(X_j) - \left[\left(\frac{1}{n} \sum_{j=1}^n \log(1 + X_j^c) \right)^{-1} + 1 \right] \cdot \sum_{j=1}^n \frac{X_j^c \log(X_j)}{1 + X_j^c} \stackrel{!}{=} 0$$

giving an estimate \hat{c}_n^{ML} for c which we then introduce into

$$\hat{k}_n^{ML} = \left(\frac{1}{n} \sum_{j=1}^n \log (1 + X_j^{\hat{c}_n^{ML}}) \right)^{-1}.$$

Both relations are easily derived from the likelihood equations. Additionally, we consider the minimum Cramér–von Mises distance estimator, which can be calculated from

$$\begin{aligned} \hat{\vartheta}_n^{CvM} &= (\hat{c}_n^{CvM}, \hat{k}_n^{CvM}) \\ &= \arg \min \left\{ \frac{1}{n} \sum_{j=1}^n (1 + X_{(j)}^c)^{-k} \left[\frac{2j-1}{n} - 2 + (1 + X_{(j)}^c)^{-k} \right] \mid c, k > 0 \right\} \end{aligned}$$

(the minimization is solved numerically, similar to our new estimator). Note that there have been further contributions to the estimation of the Burr parameters (see Schmittlein, 1983; Shah & Gokhale, 1993; Wingo, 1993; Wang & Cheng, 2010).

Like for the exponential- and Rayleigh distribution, we approximate bias and MSE of these estimators and show the results in Tables 5 and 6. For each value of ϑ_0 and n , the first line corresponds to the bias/MSE of the estimator for the c -parameter, and the second line corresponds to the k -parameter. As before, it becomes evident that our new procedure outperforms the maximum likelihood and minimum Cramér–von Mises distance estimator in terms of the bias. Unlike for the exponential distribution, the dependence on the tuning parameter “ a ” is less clear: for a great deal of parameter values and sample sizes, the estimator $\hat{\vartheta}_{n,2}^{(3)}$ yields the best result, but in some cases (mostly for the k -parameter) the estimator $\hat{\vartheta}_{n,2}^{(0.25)}$, with tuning parameter from the other end of the spectrum, performs best. Also observe the oddity that in some cases the estimator fares noticeably worse for $a = 0.5$ than for both smaller and larger tuning parameters. Thus, if one seeks to minimize some measure of quality of the estimators, an optimal, data-dependent choice of the tuning parameter would be useful (more on this in Section 10). In the light of our simulations, we suggest the use of $\hat{\vartheta}_{n,2}^{(3)}$ in practice as long as no adaptive tuning is available. Both in the bias and in the MSE simulation, the maximum likelihood estimator ran into computational issues for sample size $n = 10$. The minimum Cramér–von Mises distance estimator is more stable in this regard, but still a lot less so than our new estimators which show notably slighter outliers only for large values of the Burr parameters. Once samples get larger ($n = 50+$), the asymptotic optimality properties of the maximum likelihood estimator appear to kick in, as its performance stabilizes. Still for suitably chosen tuning parameter, our estimators are very close in virtually all instances. The small sample behaviour of the maximum likelihood estimator poses a huge drawback for applications and the problem is well-known.

9. EXAMPLE: EXPONENTIAL-POLYNOMIAL MODELS

We now proceed to consider an example of a non-normalized parametric model, one of the major motivations to this work. In particular, let

$$p_{\vartheta}(x) = C(\vartheta)^{-1} \exp (\vartheta_1 x + \dots + \vartheta_d x^d), \quad x > 0, \tag{8}$$

$\vartheta = (\vartheta_1, \dots, \vartheta_d) \in \mathbb{R}^{d-1} \times (-\infty, 0) = \Theta$, where

$$C(\vartheta) = \int_0^{\infty} \exp (\vartheta_1 x + \dots + \vartheta_d x^d) dx.$$

TABLE 5: Approximated biases calculated with 100,000 Burr-distributed Monte Carlo samples.

$\vartheta_0 = \binom{c_0}{k_0}$	n	$\hat{\vartheta}_n^{ML}$	$\hat{\vartheta}_n^{CvM}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$
$\binom{0.8}{2}$	10	–	0.1420	0.0094	–0.1608	0.0375	0.0636	0.0585
		–	1.3014	0.1750	–0.2579	0.0243	0.1635	0.1382
	25	0.0406	0.0451	–0.0377	–0.1960	0.0240	0.0223	0.0152
		0.1057	0.1311	0.0184	–0.3010	0.0503	0.0490	0.0131
	50	0.0197	0.0207	–0.0411	–0.1673	0.0121	0.0102	0.0034
		0.0491	0.0565	–0.0089	–0.2435	0.0256	0.0214	–0.0122
	100	0.0097	0.0102	–0.0307	–0.1102	0.0062	0.0056	0.0010
		0.0234	0.0266	–0.0128	–0.1505	0.0125	0.0125	–0.0095
	200	0.0046	0.0050	–0.0120	–0.0510	0.0030	0.0029	0.0012
		0.0114	0.0131	–0.0048	–0.0660	0.0064	0.0071	–0.0001
$\binom{2}{5}$	10	0.2956	0.3458	–0.2755	–1.0773	0.2152	0.1987	0.1841
		2.8551	37.1075	1.6188	–2.1985	2.3500	2.2870	2.1681
	25	0.1027	0.1082	–0.1434	–1.2772	0.0725	0.0655	0.0618
		0.6208	0.8619	0.2341	–2.6011	0.5033	0.4754	0.4647
	50	0.0476	0.0492	–0.0347	–1.4268	0.0326	0.0298	0.0283
		0.2669	0.3415	0.1278	–2.8090	0.2126	0.2039	0.2021
	100	0.0233	0.0233	0.0079	–1.5877	0.0159	0.0145	0.0138
		0.1285	0.1565	0.0946	–3.0394	0.1025	0.0989	0.0983
	200	0.0120	0.0113	0.0089	–1.7622	0.0082	0.0076	0.0073
		0.0627	0.0732	0.0526	–3.3064	0.0500	0.0485	0.0483
$\binom{5}{0.8}$	10	–	2.1411	1.0267	1.0622	1.0927	1.1037	1.1374
		–	0.0451	0.0322	0.0313	0.0327	0.0380	0.0432
	25	0.3731	0.4635	0.3177	0.3233	0.3167	0.3101	0.3146
		0.0143	0.0113	0.0096	0.0095	0.0106	0.0130	0.0151
	50	0.1748	0.2071	0.1519	0.1530	0.1488	0.1453	0.1463
		0.0063	0.0046	0.0038	0.0039	0.0045	0.0057	0.0067
	100	0.0835	0.0960	0.0731	0.0729	0.0708	0.0690	0.0693
		0.0031	0.0023	0.0019	0.0019	0.0022	0.0028	0.0033
	200	0.0421	0.0481	0.0375	0.0370	0.0360	0.0348	0.0347
		0.0016	0.0012	0.0010	0.0010	0.0012	0.0015	0.0017

These density functions correspond to a so-called exponential-polynomial model, which constitutes a special type of exponential family. It is trivial to see that these density functions obey the regularity assumptions (R1)–(R3), and also not hard to verify that the regularity conditions stated by Betsch & Ebner (2019a) (as summarized in Section 2) are satisfied. Thus, we can first of all note, as a corollary to Theorem 3 of Betsch & Ebner (2019a), the following characterization

TABLE 6: Approximated MSE values calculated with 100,000 Burr-distributed Monte Carlo samples.

$\vartheta_0 = \binom{c_0}{k_0}$	n	$\hat{\vartheta}_n^{ML}$	$\hat{\vartheta}_n^{CvM}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$	
$\binom{0.8}{2}$	10	–	0.1790	0.1398	0.2562	0.1209	0.1092	0.1199	
		–	22686.9274	1.9083	0.9712	1.0083	2.2673	2.5852	
	25	0.0228	0.0293	0.0724	0.2017	0.0310	0.0344	0.0406	
		0.2354	0.3578	0.2867	0.3025	0.3438	0.5130	0.7007	
	50	0.0090	0.0121	0.0518	0.1577	0.0137	0.0159	0.0200	
		0.0957	0.1242	0.1437	0.1853	0.1557	0.2298	0.3430	
	100	0.0042	0.0056	0.0337	0.1009	0.0066	0.0076	0.0100	
		0.0483	0.0545	0.0781	0.1136	0.0737	0.1066	0.1710	
	200	0.0020	0.0027	0.0144	0.0472	0.0032	0.0036	0.0045	
		0.2070	0.0255	0.0368	0.0582	0.0358	0.0507	0.0769	
	$\binom{2}{5}$	10	0.4819	0.9394	1.0985	2.6895	0.4657	0.4352	0.4291
			260.2839	1671094.1072	88.3851	22.0669	164.7637	180.7299	168.6383
25		0.1139	0.1717	0.4778	2.7278	0.1143	0.1135	0.1163	
		3.5036	10.6678	4.6982	9.5310	3.3494	3.3022	3.3652	
50		0.0477	0.0699	0.1719	2.9353	0.0497	0.0504	0.0526	
		1.0590	1.8952	1.5019	10.0855	1.0549	1.0890	1.1556	
100		0.0221	0.0321	0.0410	3.2160	0.0237	0.0243	0.0254	
		0.4312	0.6940	0.5032	10.9842	0.4435	0.4665	0.5016	
200		0.0107	0.0153	0.0119	3.5414	0.0116	0.0119	0.0125	
		0.1954	0.2996	0.2009	12.0603	0.2039	0.2156	0.2330	
$\binom{5}{0.8}$		10	–	63.9726	14.0301	14.8418	16.3923	16.7479	17.2525
			–	0.2941	0.1349	0.1370	0.1405	0.1479	0.1565
	25	1.6759	2.7939	1.6869	1.6671	1.6310	1.6163	1.6767	
		0.0403	0.0496	0.0443	0.0450	0.0461	0.0483	0.0508	
	50	0.6177	0.8648	0.6609	0.6460	0.6299	0.6307	0.6597	
		0.0189	0.0227	0.0212	0.0215	0.0220	0.0229	0.0239	
	100	0.2746	0.3641	0.2958	0.2891	0.2827	0.2844	0.2992	
		0.0091	0.0108	0.0103	0.0105	0.0107	0.0111	0.0116	
	200	0.1293	0.1671	0.1395	0.1366	0.1338	0.1350	0.1430	
		0.0045	0.0054	0.0052	0.0052	0.0053	0.0055	0.0058	

Corollary 7. A positive random variable X with $\mathbb{E}X^d < \infty$ follows the exponential-polynomial model in (8) if, and only if, the distribution function F_X of X satisfies

$$F_X(t) = \mathbb{E} \left[- \left(\sum_{k=1}^d k \vartheta_k X^{k-1} \right) \min\{X, t\} \right], \quad t > 0.$$

This is the characterization which underlies our estimation method as constructed in Section 2.

Notice that $C(\vartheta)$ cannot be written in a closed form, so maximum likelihood estimators are not readily available for the model in (8). Using the method of holonomic gradient descent, introduced by Nakayama et al. (2011), Hayakawa & Takemura (2016) identify a differential equation which allows to numerically calculate $C(\vartheta)$ and its derivatives, and thus to get an approximation of the ML estimator. In our simulations, however, we focus on methods that do not try to approximate $C(\vartheta)$ numerically, but get rid of the normalization constant altogether. Namely, we consider our new method and compare it to the well-known score matching approach of Hyvärinen (2007), in generalization of his method introduced in Hyvärinen (2005), as well as to the noise-contrastive estimation technique of Gutmann & Hyvärinen (2012). In the case of non-negative, univariate observations, the score matching approach boils down to finding the minimum of

$$\begin{aligned} \tilde{J}_{NN}(\vartheta) &= \frac{1}{n} \sum_{j=1}^n \left[2X_j \cdot \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} + \frac{p''_\vartheta(X_j)}{p_\vartheta(X_j)} X_j^2 - \frac{1}{2} \cdot \frac{(p'_\vartheta(X_j))^2}{(p_\vartheta(X_j))^2} X_j^2 \right] \\ &= \frac{1}{n} \sum_{j=1}^n \left[\left(\sum_{k=1}^d k(k+1) \vartheta_k X_j^k \right) + \frac{1}{2} \left(\sum_{k=1}^d k \vartheta_k X_j^k \right)^2 \right] \end{aligned}$$

(see Section 3 of Hyvärinen, 2007), where X_1, \dots, X_n are i.i.d. random variables such that $X_1 \sim p_{\vartheta_0}$, for some unknown $\vartheta_0 \in \Theta$. Clearly, the quantity does not rely on $C(\vartheta)$. As for the estimator constructed in this article, fixing $q = 2$ and the weight $w(t) = e^{-at}$, where $a > 0$ is a tuning parameter, we may calculate $\psi_{n,2}(\vartheta) = \|\eta_n(\cdot, \vartheta)\|_{L^2}$ (see Sections 2 and 3) explicitly as

$$\begin{aligned} \psi_{n,2}(\vartheta) &= \frac{1}{n^2} \sum_{j=1}^n \left\{ \frac{e^{-aX_j}}{a} \left[2 \sum_{\ell=1}^d \ell \vartheta_\ell X_j^\ell + 1 - \left(\sum_{\ell=1}^d \ell \vartheta_\ell X_j^{\ell-1} \right)^2 \left(\frac{2X_j}{a} + \frac{2}{a^2} \right) \right] \right. \\ &\quad \left. + \frac{2}{a^3} \left(\sum_{\ell=1}^d \ell \vartheta_\ell X_j^{\ell-1} \right)^2 \right\} \\ &\quad - \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left\{ \left(\sum_{\ell=1}^d \ell \vartheta_\ell X_{(j)}^\ell + 1 \right) \left[-\frac{e^{-aX_{(k)}}}{a} + \frac{e^{-aX_{(k)}}}{a^2} \sum_{\ell=1}^d \ell \vartheta_\ell X_{(k)}^{\ell-1} \right] \right. \\ &\quad \left. + \left(\sum_{\ell=1}^d \ell \vartheta_\ell X_{(k)}^{\ell-1} \right) \left[\frac{e^{-aX_{(j)}}}{a^2} \left(\sum_{\ell=1}^d \ell \vartheta_\ell X_{(j)}^\ell \right) - \frac{X_{(j)} e^{-aX_{(j)}}}{a} - \frac{e^{-aX_{(j)}}}{a^2} \right] \right. \\ &\quad \left. - \left(\sum_{\ell=1}^d \ell \vartheta_\ell X_{(j)}^{\ell-1} \right) \frac{2}{a^3} (1 - e^{-aX_{(j)}}) \right\}, \end{aligned}$$

where $X_{(1)} < \dots < X_{(n)}$ are the ordered values X_1, \dots, X_n . This formula is notably more complicated than the one resulting from the score matching approach, but in the two-parameter setting we now turn to, both estimators can be calculated explicitly. More precisely, to keep the presentation clear, we intent to focus on a two parameter case, but in order not to end up with a Gaussian-type model, we consider $d = 3$ and fix $\vartheta_2 = 0$, thus effectively considering the model

$$p_\vartheta(x) = C(\vartheta_1, \vartheta_3)^{-1} \exp(\vartheta_1 x + \vartheta_3 x^3), \quad x > 0, \quad \vartheta_1 \in \mathbb{R}, \vartheta_3 \in (-\infty, 0).$$

In this case, each time by solving a quadratic equation in ϑ_1 and ϑ_3 (which is obtained by simplifying the above quantities \bar{J}_{NN} and $\psi_{n,2}$ further), we obtain the estimators explicitly. The score matching estimators for $\vartheta = (\vartheta_1, \vartheta_3)$ are given as

$$\hat{\vartheta}_n^{SM} = \left(-\frac{2m_1}{m_2} - \frac{3m_4}{m_2} \cdot \frac{4m_2m_3 - 2m_1m_4}{3(m_4)^2 - 3m_2m_6}, \frac{4m_2m_3 - 2m_1m_4}{3(m_4)^2 - 3m_2m_6} \right),$$

where $m_k = \sum_{j=1}^n X_j^k$, and our new estimators are

$$\hat{\vartheta}_{n,2}^{(a)} = \left(\frac{\bar{\Psi}_n^{(3)}\bar{\Psi}_n^{(5)} - 2\bar{\Psi}_n^{(2)}\bar{\Psi}_n^{(4)}}{4\bar{\Psi}_n^{(1)}\bar{\Psi}_n^{(2)} - (\bar{\Psi}_n^{(3)})^2}, \frac{\bar{\Psi}_n^{(3)}\bar{\Psi}_n^{(4)} - 2\bar{\Psi}_n^{(1)}\bar{\Psi}_n^{(5)}}{4\bar{\Psi}_n^{(1)}\bar{\Psi}_n^{(2)} - (\bar{\Psi}_n^{(3)})^2} \right),$$

where

$$\begin{aligned} \bar{\Psi}_n^{(1)} &= \frac{2}{a^3} + \frac{1}{n^2} \sum_{j=1}^n e^{-aX_{(j)}} \left(-\frac{2X_{(j)}}{a^2} - \frac{2}{a^3} (2n - 2j + 1) \right) \\ &\quad + \frac{2}{n^2} \sum_{1 \leq j < k \leq n} -\frac{X_{(j)}}{a^2} (e^{-aX_{(k)}} + e^{-aX_{(j)}}), \end{aligned}$$

$$\begin{aligned} \bar{\Psi}_n^{(2)} &= \frac{1}{n^2} \sum_{j=1}^n \left\{ -\frac{18X_{(j)}^5}{a^2} e^{-aX_{(j)}} + \frac{18X_{(j)}^4}{a^3} (1 - e^{-aX_{(j)}}) \right\} \\ &\quad + \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left\{ -\frac{9X_{(j)}^3 X_{(k)}^2}{a^2} (e^{-aX_{(k)}} + e^{-aX_{(j)}}) + \frac{18X_{(j)}^2 X_{(k)}^2}{a^3} (1 - e^{-aX_{(j)}}) \right\}, \end{aligned}$$

$$\begin{aligned} \bar{\Psi}_n^{(3)} &= \frac{1}{n^2} \sum_{j=1}^n \left\{ -\frac{12X_{(j)}^3}{a^2} e^{-aX_{(j)}} + \frac{12X_{(j)}^2}{a^3} (1 - e^{-aX_{(j)}}) \right\} \\ &\quad + \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left\{ (e^{-aX_{(k)}} + e^{-aX_{(j)}}) \left(-\frac{3X_{(j)}X_{(k)}^2}{a^2} - \frac{3X_{(j)}^3}{a^2} \right) \right. \\ &\quad \left. + \frac{6}{a^3} (1 - e^{-aX_{(j)}}) (X_{(j)}^2 + X_{(k)}^2) \right\}, \end{aligned}$$

$$\bar{\Psi}_n^{(4)} = \frac{1}{n^2} \sum_{j=1}^n e^{-aX_{(j)}} \left(\frac{2X_{(j)}}{a} + \frac{2(n - 2j + 1)}{a^2} \right) + \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \frac{X_{(j)}}{a} (e^{-aX_{(k)}} + e^{-aX_{(j)}}),$$

and

$$\begin{aligned} \bar{\Psi}_n^{(5)} &= \frac{2}{n^2} \sum_{1 \leq j < k \leq n} \left\{ \frac{3X_{(j)}^3}{a} e^{-aX_{(k)}} + \frac{3X_{(j)}X_{(k)}^2}{a} e^{-aX_{(j)}} + \frac{3X_{(k)}^2}{a^2} (e^{-aX_{(j)}} - e^{-aX_{(k)}}) \right\} \\ &\quad + \frac{1}{n^2} \sum_{j=1}^n \frac{6X_{(j)}^3}{a} e^{-aX_{(j)}}. \end{aligned}$$

Moreover, we consider the noise-contrastive estimators in the refined version of Gutmann & Hyvärinen (2012) that generalizes the initial results of Gutmann & Hyvärinen (2010). The idea is motivated by a binary classification problem and proceeds to consider the unknown normalization constant as an additional parameter to be estimated. The objective function is constructed in such a way that it ensures that the obtained estimator for the normalization constant truly provides (in numerical approximation) a normalized density without any further constraints on the optimization. Following Gutmann & Hyvärinen (2012), we implement this technique as follows. Given the sample X_1, \dots, X_n , choose the noise sample size $T_n = \nu \cdot n$ (for some fixed $\nu \in \mathbb{N}$, in our case $\nu = 10$) and sample from the noise distribution (in our case, the exponential distribution with rate parameter $\lambda_n = n / \sum_{j=1}^n X_j$) to obtain values Y_1, \dots, Y_{T_n} . Then, minimize the objective function

$$J(\vartheta_1, \vartheta_3, c) = \frac{1}{n} \sum_{j=1}^n \log \left(1 + \nu \lambda_n \exp \left(-(\lambda_n + \vartheta_1)X_j - \vartheta_3 X_j^3 - c \right) \right) \\ - \frac{\nu}{T_n} \sum_{k=1}^{T_n} \log \left(1 - \frac{1}{1 + \nu \lambda_n \exp \left(-(\lambda_n + \vartheta_1)Y_k - \vartheta_3 Y_k^3 - c \right)} \right)$$

to obtain an estimator $\hat{\vartheta}_n^{NC}$ for the unknown parameters $(\vartheta_1^{(0)}, \vartheta_3^{(0)})$ as well as for the logarithm of the inverse of the normalization constant. In our simulations we used the “L-BFGS-B”-method, which we have also applied in previous examples, for this optimization (with initial values $(0, -0.1, 0)$ and with the second parameter constrained to the negative numbers). As in the previous simulations, we approximate bias and MSE of the competing estimators. The results are presented in Tables 7 and 8. In the tables, for each underlying parameter $\vartheta_0 = (\vartheta_1^{(0)}, \vartheta_3^{(0)})$ and each sample size, the first line corresponds to the bias/MSE of the ϑ_1 -parameter, while the second line corresponds to the ϑ_3 -parameter.

It is immediate that our new estimator and the noise-contrastive estimator outperform the score matching method distinctively over all tuning parameters, sample sizes, and parameter values for both the bias and the MSE, with the only exception being the parameter vector $(1, -0.05)$, for which the estimator $\hat{\vartheta}_{n,2}^{(5)}$ fares worse than the score matching approach in MSE terms. We propose as a very good compromise choice of the tuning parameter the use of $\hat{\vartheta}_{n,2}^{(1)}$ as an estimator. This particular estimator outperforms the score matching method by factors of (at least) 4 in terms of MSE and also fares notably better in terms of the bias. It also outperforms the noise-contrastive estimation method uniformly, except for four instances in the MSE values (in three of which our method still performs better when another tuning parameter is chosen). The simulation in this non-normalized models conforms with the observation from previous examples that the new method fares remarkably well bias-wise. We also note that all of the estimators admit a large MSE for very small sample sizes, a behaviour to be expected. From our simulations we conclude that the new estimation method is to be preferred clearly over the other approaches in this univariate setting of the exponential-polynomial models, but of course larger scale simulations involving different types of multi-parameter versions of the model would be needed to further strengthen this position (also, generalizations of the score matching technique, like Yu, Drton & Shojaie (2019), could be taken into account). One massive advantage of the score matching and noise-contrastive estimation approaches, however, is that they readily generalize to the multivariate situation, a generalization we were not (yet) able to establish for our approach (see the last paragraph of Section 10).

Remark 8. *We observed in our simulations that the noise-contrastive estimators can run into computational problems when the exponentials in the objective function raise an overflow warning. A step by step analysis of the code suggests that for large noise sample sizes T_n*

TABLE 7: Approximated biases calculated with 100,000 Monte Carlo samples.

ϑ_0	n	$\hat{\vartheta}_n^{SM}$	$\hat{\vartheta}_n^{NC}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$	$\hat{\vartheta}_{n,2}^{(5)}$	
$\begin{pmatrix} 1 \\ -0.05 \end{pmatrix}$	10	1.4804	(0.4925)	0.4559	0.3977	0.3223	0.2779	0.3353	0.7036	
		-0.0619	(-0.0268)	-0.0248	-0.0223	-0.0192	-0.0176	-0.0206	-0.0387	
	25	0.5724	(0.1589)	0.1531	0.1320	0.1067	0.0903	0.0960	0.1651	
		-0.0221	(-0.0084)	-0.0080	-0.0072	-0.0062	-0.0057	-0.0062	-0.0100	
	50	0.2982	0.0740	0.0725	0.0620	0.0495	0.0406	0.0418	0.0641	
		-0.0111	(-0.0039)	-0.0038	-0.0034	-0.0029	-0.0026	-0.0028	-0.0042	
	100	0.1521	(0.0344)	0.0339	0.0290	0.0233	0.0199	0.0215	0.0335	
		-0.0056	(-0.0019)	-0.0018	-0.0016	-0.0014	-0.0013	-0.0014	-0.0022	
	200	0.0790	(0.0172)	0.0170	0.0145	0.0117	0.0101	0.0109	0.0162	
		-0.0029	(-0.0010)	-0.0009	-0.0008	-0.0007	-0.0007	-0.0007	-0.0011	
	$\begin{pmatrix} 0 \\ -0.5 \end{pmatrix}$	10	4.3787	0.9147	0.9046	0.8328	0.7070	0.5174	0.3913	0.2530
			-1.9839	-0.7007	-0.6992	-0.6667	-0.6091	-0.5203	-0.4607	-0.4004
25		1.6908	0.3153	0.3014	0.2722	0.2237	0.1570	0.1166	0.0738	
		-0.6081	-0.2038	-0.1984	-0.1868	-0.1673	-0.1407	-0.1254	-0.1127	
50		0.8920	0.1540	0.1456	0.1310	0.1077	0.0771	0.0592	0.0402	
		-0.2963	-0.0951	-0.0917	-0.0861	-0.0771	-0.0657	-0.0595	-0.0549	
100		0.4785	0.0729	0.0691	0.0617	0.0499	0.0344	0.0250	0.0148	
		-0.1512	-0.0443	-0.0430	-0.0402	-0.0358	-0.0301	-0.0269	-0.0243	
200		0.2571	0.0375	0.0351	0.0314	0.0255	0.0177	0.0129	0.0075	
		-0.0789	-0.0223	-0.0213	-0.0199	-0.0177	-0.0149	-0.0133	-0.0119	
$\begin{pmatrix} -0.5 \\ -3 \end{pmatrix}$		10	8.5068	1.7608	1.7578	1.6817	1.5396	1.2927	1.0904	0.7944
			-14.5606	-5.1057	-5.1765	-5.0482	-4.8060	-4.3777	-4.0194	-3.4851
	25	3.2683	0.6169	0.5935	0.5616	0.5036	0.4084	0.3361	0.2394	
		-4.3203	-1.481	-1.4566	-1.4099	-1.3245	-1.1834	-1.0761	-0.9369	
	50	1.7089	0.2881	0.2739	0.2577	0.2288	0.1832	0.1499	0.1066	
		-2.0575	-0.6602	-0.6428	-0.6200	-0.5793	-0.5151	-0.4688	-0.4119	
	100	0.9238	0.1453	0.1376	0.1293	0.1148	0.0921	0.0758	0.0545	
		-1.0552	-0.3193	-0.3114	-0.2999	-0.2798	-0.2486	-0.2266	-0.1996	
	200	0.4917	0.0730	0.0674	0.0634	0.0564	0.0457	0.0380	0.0282	
		-0.5441	-0.1588	-0.1518	-0.1463	-0.1367	-0.1222	-0.1122	-0.1002	

(that is, for large ν) one tends to obtain some large values in the sample Y_1, \dots, Y_{T_n} which are cubed in the exponential terms and thus become very (if not too) large. The behaviour seems to appear more often for small parameter values $\vartheta_3^{(0)}$, but it seems to affect only single evaluations of the objective function during the optimization routine. We believe that most values for the noise-contrastive estimation approach in the table are intact and they also replicated when we reran the whole simulation, with a bit of an exception in the case of the parameter vector $(1, -0.05)$, where the values show a rather noticeable dependence on the initial value chosen

TABLE 8: Approximated MSE values calculated with 100,000 Monte Carlo samples.

ϑ_0	n	$\hat{\vartheta}_n^{SM}$	$\hat{\vartheta}_n^{NC}$	$\hat{\vartheta}_{n,2}^{(0.25)}$	$\hat{\vartheta}_{n,2}^{(0.5)}$	$\hat{\vartheta}_{n,2}^{(1)}$	$\hat{\vartheta}_{n,2}^{(2)}$	$\hat{\vartheta}_{n,2}^{(3)}$	$\hat{\vartheta}_{n,2}^{(5)}$	
$\binom{1}{-0.05}$	10	5.9091	(1.5668)	1.4652	1.3332	1.2118	1.5497	3.2052	25.4377	
		0.0120	(0.0041)	0.0038	0.0035	0.0031	0.0037	0.0068	0.0434	
	25	1.1954	(0.2810)	0.2982	0.2726	0.2568	0.3570	0.7326	4.1544	
		0.0019	(0.0006)	0.0006	0.0006	0.0006	0.0007	0.0015	0.0082	
	50	0.4758	0.1143	0.1228	0.1133	0.1090	0.1558	0.3253	1.6347	
		0.0007	(0.0002)	0.0002	0.0002	0.0002	0.0003	0.0006	0.0033	
	100	0.2113	(0.0514)	0.0550	0.0511	0.0500	0.0735	0.1541	0.7475	
		0.0003	(0.0001)	0.0001	0.0001	0.0001	0.0001	0.0003	0.0015	
	200	0.1012	(0.0249)	0.0262	0.0245	0.0241	0.0358	0.0752	0.3573	
		0.0001	(0)	0	0	0	0.0001	0.0001	0.0007	
	$\binom{0}{-0.5}$	10	42.1253	4.8720	5.7900	5.5517	5.1880	4.8596	5.0047	6.6628
			15.2937	3.2687	3.4841	3.3680	3.1697	2.8990	2.7841	3.0120
25		7.3114	1.1328	1.2154	1.1777	1.1311	1.1409	1.2750	1.9169	
		1.1392	0.2851	0.2938	0.2836	0.2690	0.2614	0.2814	0.4100	
50		2.6352	0.4930	0.4993	0.4874	0.4765	0.4997	0.5766	0.8948	
		0.3115	0.0934	0.0929	0.0905	0.0878	0.0907	0.1040	0.1660	
100		1.1005	0.2248	0.2221	0.2175	0.2144	0.2292	0.2688	0.4246	
		0.1097	0.0362	0.0358	0.0350	0.0345	0.0370	0.0439	0.0736	
200		0.5090	0.1099	0.1066	0.1046	0.1038	0.1120	0.1320	0.2093	
		0.0456	0.0164	0.0159	0.0156	0.0156	0.0170	0.0205	0.0351	
$\binom{-0.5}{-3}$		10	158.3410	17.0727	21.1996	20.7230	19.8873	18.6499	17.9361	17.9169
			911.2872	171.7483	206.1279	202.4045	195.4680	183.5396	174.0526	161.9207
	25	26.1605	3.9728	4.2928	4.2173	4.0966	3.9673	3.9750	4.3622	
		56.2301	14.0880	14.6718	14.3849	13.8935	13.2242	12.9520	13.5457	
	50	9.2065	1.6764	1.7049	1.6816	1.6482	1.6295	1.6687	1.9025	
		14.5421	4.2357	4.2429	4.1757	4.0717	3.9785	4.0302	4.5485	
	100	3.7881	0.7814	0.7736	0.7647	0.7531	0.7527	0.7795	0.9053	
		4.9664	1.6415	1.6254	1.6052	1.5777	1.5715	1.6264	1.9134	
	200	1.7190	0.3773	0.3638	0.3601	0.3559	0.3585	0.3740	0.4388	
		1.9962	0.7272	0.7056	0.6988	0.6914	0.6987	0.7334	0.8830	

for the optimization (though this does not happen for the other parameter values). Therefore, one possible way to reduce the occurrence of overflows, which lies in choosing small initial values for the ϑ_3 -parameter in the optimization routine, has an impact on the performance of the estimator. Another way out could be to adopt noise distributions with extremely short tails. It could prove useful to see if our observations replicate in other simulation studies. Note that no computational issues arise for the score matching and our new approach, where the estimators can be calculated explicitly.

10. NOTES AND COMMENTS

Note that there remain some problems for further research on our newly proposed estimators, the discussion or extension of which would be too extensive for this contribution. First, for all estimators we considered explicitly, we incorporate a tuning parameter “ a ” on which the performance depends strongly. It would be beneficial to have an adaptive choice of this parameter (see Allison & Santana, 2015, and the refinements by Tenreiro, 2019, who discuss such a method in the context of goodness-of-fit testing problems), probably adaptable to which criterion (minimal bias etc.) the estimator should satisfy. In the context of deriving results for $a \rightarrow \infty$, we obtained another consistent estimator for the Rayleigh parameter, and it would be interesting to see if such results can be derived for other distributions. Also, we have not used in practice the flexibility gained by providing all results for the general L^q -spaces, but restricted our attention to the case $q = 2$, mostly because of the explicit formulae obtainable in that case. If no closed formula for $\psi_{n,q}$ is feasible, either because of the use of some $q \neq 2$ or because some advanced weight function w is chosen, the integral in $\psi_{n,q}$ has to be solved numerically which could lead to a computationally highly demanding procedure overall. As for the choice of a specific weight function w , to the best of our knowledge there exist no theoretical results which favour specific choices over others. Considering the vast amount of weighted L^2 -statistics put to use in goodness-of-fit testing problems, it seems we cannot hope for general results in that direction. As such, the choice of the weight function provides some flexibility, but without clear guidance to satisfy specific objectives other than $\psi_{n,q}$ being calculable explicitly.

We have proven in a quite usual setting the consistency of our estimators. Surely, a limit theorem of the type

$$s(n)(\hat{\vartheta}_{n,q} - \vartheta_0) \xrightarrow{d} \mathcal{P},$$

where $s(n) \rightarrow \infty$, as $n \rightarrow \infty$, and where \mathcal{P} is some limit distribution (e.g., the normal distribution) is desirable. Such a result would pave the way for constructing confidence regions for the true parameter based on our method. The main hurdle in direct approaches of proving such a limit result, like some Taylor expansion or methods from empirical process theory, is that the terms involved in such calculations become too complicated and make the endeavour appear impractical to us. One hope is that, since Barp et al. (2019) provide limit results for special classes of Stein discrepancy-based estimators, the interpretation of our estimation method in terms of the feature Stein discrepancy might at some point lead to advances.

Moreover, a larger-scale simulation study, involving more underlying parameters, sample sizes, and distributions could provide further insight into the estimation method. Improvements from a numerical point of view would, of course, benefit the approach. From a theoretical perspective, an important step in this last direction is to study whether the minimization method that is used in cases where the estimators cannot be calculated explicitly will always find a global minimum, or if not, in which situations it is likely to get stuck in some local minimum.

Note that Betsch & Ebner (2019a) also give characterization results for density functions on bounded intervals or on the whole real line. These can be used to construct similar estimation methods in the corresponding cases. To sketch the idea in the case of parametric models on the whole real line, assume that the support of each density function p_ϑ in \mathfrak{P}_Θ is the whole real line (and that some mild regularity conditions hold). Let \tilde{X} be a real-valued random variable with

$$\mathbb{E} \left[\left| \frac{p'_\vartheta(\tilde{X})}{p_\vartheta(\tilde{X})} \right| (|\tilde{X}| + 1) \right] < \infty, \quad \vartheta \in \Theta,$$

and consider

$$\tilde{\eta}(t, \vartheta) = \mathbb{E} \left[\frac{p'_{\vartheta}(\tilde{X})}{p_{\vartheta}(\tilde{X})} (t - \tilde{X}) \mathbb{1}\{\tilde{X} \leq t\} \right] - F_{\tilde{X}}(t)$$

for $(t, \vartheta) \in \mathbb{R} \times \Theta$. Then, similar to our elaborations in Section 2, Theorem 4.1 of Betsch & Ebner (2019a) shows that $\tilde{X} \sim p_{\vartheta_0}$ if, and only if, $\tilde{\eta}(t, \vartheta_0) = 0$ for every $t \in \mathbb{R}$. Therefore, if, initially, $\tilde{X} \sim p_{\vartheta_0}$, then $\|\tilde{\eta}(\cdot, \vartheta)\|_{L^q} = 0$ if, and only if, $\vartheta = \vartheta_0$. Here, $L^q = L^q(\mathbb{R}, \mathcal{B}^1, \tilde{w}(t) dt)$, $1 \leq q < \infty$, with a positive weight function \tilde{w} satisfying

$$\int_{\mathbb{R}} (|t|^q + 1) \tilde{w}(t) dt < \infty.$$

Thus, with

$$\tilde{\eta}_n(t, \vartheta) = \frac{1}{n} \sum_{j=1}^n \frac{p'_{\vartheta}(\tilde{X}_j)}{p_{\vartheta}(\tilde{X}_j)} (t - \tilde{X}_j) \mathbb{1}\{\tilde{X}_j \leq t\} - \frac{1}{n} \sum_{j=1}^n \mathbb{1}\{\tilde{X}_j \leq t\},$$

a reasonable estimator for ϑ_0 is

$$\tilde{\vartheta}_{n,q} = \arg \min \{ \|\tilde{\eta}_n(\cdot, \vartheta)\|_{L^q} \mid \vartheta \in \Theta \}.$$

Apparently, once we switch to density function supported by the whole real line, the characterization result due to Betsch & Ebner (2019a), and thus our estimator, have slightly different forms, but using the results from Section 3, we could still prove existence and measurability for this type of estimator, and give a formal definition as in (6). Moreover, a classical proof via the law of large numbers for random elements in separable Banach spaces and the Arzelà-Ascoli theorem (considering the modulus of continuity, as employed by Billingsley, 1968) yields the convergence results from Lemma 3 for $\tilde{\psi}_{n,q} = \|\tilde{\eta}_n(\cdot, \vartheta)\|_{L^q}$, but with all convergences only in probability. That result can then be used to derive consistency as in Theorem 4, again with all convergences only in probability. However, choosing a fixed (i.e., parameter-independent) weight function on \mathbb{R} with a mere scale-tuning, as we employ it throughout (using the weight $t \mapsto e^{-at}$), appears not to be sufficient to account for the possible location-dependence of the model. Thus, in simulations (for instance with the Cauchy distribution) the problem, to us, seems empirically more involved and is therefore not addressed in the work at hand.

Still, we deem it possible to apply our new type of estimator to models which are supported by any connected subset of \mathbb{R} as indicated in the previous lines. Of course, the next question which forces itself on us is whether a similar method can be devised for multivariate models. Here the frontiers are somewhat blurry: the Stein density approach identity which appears at the beginning of Section 2 is not yet fully understood in the multivariate case (as stated in Remark 1.1 by Ley, Reinert & Swan, 2017a), and the characterizations derived by Betsch & Ebner (2019a) rely on further calculations, the generalization of which is not immediate. Thus, we have to state at this point that, to us, it is an open question how a generalization to the multivariate setting could look like (with no clear indication of it being possible at all).

APPENDIX

Remark 9 (Comments on Theorem 1). *There is another result which gives measurable selections without the completeness assumption on the probability space (as provided by Brown & Purves, 1973), but it requires σ -compactness of the parameter space, thus essentially reducing the study to euclidean parameters (a Banach space is σ -compact if, and only if, it is of finite dimension, which follows easily from Baire’s category theorem). Of course this is enough for our*

purposes, but currently the interest in statistical inference for infinite dimensional models grows remarkably. Hence if a statistician was to investigate measurability of an estimator for some infinite dimensional quantity, she would have to resort to a result in the generality of Theorem 1. Another reason for us to build on Theorem 1 is that other measurability results known to us do not quite fit the construction of our estimators. For instance, Sahler (1970) considers minimum discrepancy estimators, where discrepancies are (certain) functions on the Cartesian product of a suitable set of probability measures with itself. It is (formally) not possible to identify such a set of probability measures in our setting, as we ought to introduce the empirical distribution of a sample into the discrepancy function, while only considering parametric distributions with a continuously differentiable density. Even though we believe this to be a purely formal issue which might be resolved to render results from Sahler (1970) applicable, additional caution is needed that Theorem 1 does not require. Likewise, the setting considered by Pfanzagl (1969) does not cover our estimators. Note that since completing (the σ -field of) an underlying probability space does not interfere with measurability properties of random maps, nor does it meddle with push-forward measures, the corresponding assumption in Theorem 1 is no restriction. If \mathfrak{S} is a complete, separable metric space and the map Γ from Theorem 1 takes compact subsets of \mathfrak{S} as values, the condition imposed on the graph is equivalent to Γ being measurable with respect to the Borel- σ -field generated by the Hausdorff topology (see Theorems III.2 and III.30 by Castaing & Valadier, 1977). Likewise, if \mathfrak{S} is a locally compact, complete, separable metric space and Γ maps into the closed subsets of \mathfrak{S} , the condition is equivalent to Γ being measurable with respect to the Borel- σ -field generated by the Fell topology (this can be proven using results from Beer, 1993, and Castaing & Valadier, 1977).

Proof of Lemma 2. First recall the following lemma on product-measurability, the proof of which is an easy exercise.

Lemma 10. *Let (S, \mathcal{A}, μ) be a measure space, $I \subset \mathbb{R}$ an open interval, and let $(\mathcal{T}, \mathcal{O}_{\mathcal{T}})$ be a topological vector space. Furthermore, let $h : S \times I \rightarrow \mathcal{T}$ be a map such that*

- $s \mapsto h(s, x)$ is $(\mathcal{A}, \mathcal{B}(\mathcal{T}))$ -measurable for every $x \in I$, and
- $x \mapsto h(s, x)$ is right-continuous for every $s \in S$.

Then h is $(\mathcal{A} \otimes \mathcal{B}(I), \mathcal{B}(\mathcal{T}))$ -measurable.

Notice that for fixed $(t, \vartheta) \in (0, \infty) \times \Theta$ the map $\omega \mapsto \eta_n(\omega, t, \vartheta)$ is $(\mathcal{F}, \mathcal{B}^1)$ -measurable, and for fixed $(\omega, t) \in \Omega \times (0, \infty)$ the map $\vartheta \mapsto \eta_n(\omega, t, \vartheta)$ is continuous. By a statement analogous to Lemma 10 (see Lemma III.14 by Castaing & Valadier, 1977), the map $(\omega, \vartheta) \mapsto \eta_n(\omega, t, \vartheta)$ is $(\mathcal{F} \otimes \mathcal{B}(\Theta), \mathcal{B}^1)$ -measurable for fixed $t > 0$. Since $t \mapsto \eta_n(\omega, t, \vartheta)$ is continuous for fixed $(\omega, \vartheta) \in \Omega \times \Theta$, Lemma 10 implies that η_n is $(\mathcal{F} \otimes \mathcal{B}(0, \infty) \otimes \mathcal{B}(\Theta), \mathcal{B}^1)$ -measurable. Consequently, the maps

$$(\omega, \vartheta) \mapsto \langle \eta_n(\omega, \cdot, \vartheta), g \rangle_{L^q}$$

are measurable for every $g \in L^q$ by Fubini's theorem, and since L^q is a separable Banach space, the mapping $(\omega, \vartheta) \mapsto \eta_n(\omega, \cdot, \vartheta)$ is $(\mathcal{F} \otimes \mathcal{B}(\Theta), \mathcal{B}(L^q))$ -measurable (cf. Corollary 1.1.2 of Hytönen et al., 2016). ■

Remark 11 ($\Gamma_{n,q}$ from (5) is closed). *Note that (R1) and Fatou's lemma imply the lower semi-continuity of $\vartheta \mapsto \psi_{n,q}(\omega, \vartheta)$. Thus if $\vartheta^{(k)} \in \Gamma_{n,q}(\omega)$, $k \in \mathbb{N}$, converges (with respect to the metric in Θ) to $\vartheta^* \in \Theta$ as $k \rightarrow \infty$, then*

$$\psi_{n,q}(\omega, \vartheta^*) = \|\eta_n(\omega, \cdot, \vartheta^*)\|_{L^q} \leq \liminf_{k \rightarrow \infty} \|\eta_n(\omega, \cdot, \vartheta^{(k)})\|_{L^q} \leq m_{n,q}(\omega) + \varepsilon_n(\omega),$$

that is, $\vartheta^* \in \Gamma_{n,q}(\omega)$, so $\Gamma_{n,q}(\omega)$ is closed in Θ for every $\omega \in \Omega$. Hence we can note that if Θ is closed, and therefore locally compact (cf. p. 42 of Kuratowski, 1968) and complete, $\Gamma_{n,q}$ is a random element in the space of all closed subsets of Θ endowed with the Fell topology (see also Remark 9).

Proof of Lemma 3. First note that for any non-empty closed subset F of K ,

$$\left| \inf_{\vartheta \in F} \psi_{n,q}(\vartheta) - \inf_{\vartheta \in F} \psi_q(\vartheta) \right| \leq \sup_{\vartheta \in K} |\psi_{n,q}(\vartheta) - \psi_q(\vartheta)|,$$

so the second claim of Lemma 3 follows from the first. For the first claim, let $K \neq \emptyset$ be a compact subset of Θ . Note that

$$\begin{aligned} \sup_{\vartheta \in K} |\psi_{n,q}(\vartheta) - \psi_q(\vartheta)| &\leq \sup_{\vartheta \in K} \|\eta_n(\cdot, \vartheta) - \eta(\cdot, \vartheta)\|_{L^q} \\ &\leq C \cdot \sup_{\substack{\vartheta \in K \\ t > 0}} \left| \frac{1}{n} \sum_{j=1}^n \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} \min\{X_j, t\} - \mathbb{E} \left[\frac{p'_\vartheta(X)}{p_\vartheta(X)} \min\{X, t\} \right] \right| \\ &\quad + C \cdot \sup_{t > 0} \left| \frac{1}{n} \sum_{j=1}^n \mathbb{1}\{X_j \leq t\} - F_X(t) \right|, \end{aligned} \tag{A.1}$$

where $C = \left(\int_0^\infty w(t) dt \right)^{1/q}$. The second term on the right-hand side of (A.1) converges to 0 almost surely by the classical Glivenko–Cantelli theorem. For a function $f : (0, \infty) \rightarrow \mathbb{R}$ we write $\mathbb{P}_n f = \frac{1}{n} \sum_{j=1}^n f(X_j)$ and $\mathbb{P}^X f = \mathbb{E}[f(X)]$. Then the first term on the right-hand side of (A.1) can be written as

$$\begin{aligned} \sup_{\substack{\vartheta \in K \\ t > 0}} \left| \frac{1}{n} \sum_{j=1}^n \frac{p'_\vartheta(X_j)}{p_\vartheta(X_j)} \min\{X_j, t\} - \mathbb{E} \left[\frac{p'_\vartheta(X)}{p_\vartheta(X)} \min\{X, t\} \right] \right| &= \sup_{\vartheta \in K} \left| \mathbb{P}_n f_{t,\vartheta} - \mathbb{P}^X f_{t,\vartheta} \right| \\ &= \sup_{f \in \mathcal{H}_\Theta} \left| \mathbb{P}_n f - \mathbb{P}^X f \right|, \end{aligned} \tag{A.2}$$

where $f_{t,\vartheta}(x) = \frac{p'_\vartheta(x)}{p_\vartheta(x)} \min\{x, t\}$, $x > 0$, is a measurable function for every $\vartheta \in K$ and $t > 0$, and $\mathcal{H}_\Theta = \{f_{t,\vartheta} \mid \vartheta \in K, t > 0\}$ is the collection of all such functions. Note that the supremum in (A.2) is finite (\mathbb{P} -a.s.) by (R1), (2), and (R3), and that the terms in (A.2) constitute measurable maps from (Ω, \mathcal{F}) to $(\mathbb{R}, \mathcal{B}^1)$ by Theorem 1.

As is commonly done, we denote, for given functions $l, u : (0, \infty) \rightarrow \mathbb{R}$, by $[l, u]$ the set of all functions f such that $l \leq f \leq u$ pointwise. An ε -bracket with respect to $L^1(\mathbb{P}^X) = L^1((0, \infty), \mathcal{B}(0, \infty), \mathbb{P}^X)$ is one such set $[l, u]$ with $\|u - l\|_{L^1(\mathbb{P}^X)} < \varepsilon$. The bracketing number $\mathcal{N}_{[\cdot]}(\varepsilon, \mathcal{H}_\Theta, L^1(\mathbb{P}^X))$ of \mathcal{H}_Θ is the minimum number of ε -brackets needed to cover \mathcal{H}_Θ . If the bracketing number of \mathcal{H}_Θ is finite for every $\varepsilon > 0$, then \mathcal{H}_Θ is a Glivenko-Cantelli class, that is, $\sup_{f \in \mathcal{H}_\Theta} |\mathbb{P}_n f - \mathbb{P}^X f| \rightarrow 0$ almost surely (see Theorem 2.4.1 by van der Vaart & Wellner, 2000), which, combined with (A.1) and (A.2), implies the claim. Note that the result by van der Vaart & Wellner (2000) is formulated to give convergence outer almost surely, but as we work on a complete probability space, the transition to an outer probability measure is not necessary (since we can provide enough measurability on a complete probability space and the notions of almost sure convergence and outer almost sure convergence agree).

Thus, to prove Lemma 3, it remains to show that the bracketing numbers of \mathcal{H}_Θ are finite. The following argument combines ideas from the classical Glivenko-Cantelli theorem and from

Example 19.7 of van der Vaart (1998). Let $\varepsilon > 0$ be arbitrary, and set $\delta = \varepsilon^{1/\alpha} (4 \mathbb{E}[H(X)X])^{-1/\alpha}$, where H and α are as in (R3). Since K is compact there exist $\vartheta_1, \dots, \vartheta_m \in K$, $m = m_\varepsilon \in \mathbb{N}$, such that $\bigcup_{i=1}^m B_\delta(\vartheta_i) \supset K$. Additionally, since for each $i = 1, \dots, m$ the function

$$[0, \infty) \ni t \mapsto E_i(t) = \mathbb{E} \left[\left| \frac{p'_{\vartheta_i}(X)}{p_{\vartheta_i}(X)} \right| \min\{X, t\} \right]$$

is continuous, monotonically increasing, and satisfies $E_i(0) = \lim_{t \searrow 0} E_i(t) = 0$, as well as $E_i(\infty) = \lim_{t \nearrow \infty} E_i(t) = \mathbb{E} \left[\frac{p'_{\vartheta_i}(X)}{p_{\vartheta_i}(X)} X \right] < \infty$, there exist some $0 = t_0 < t_1 < \dots < t_\ell = \infty$, $\ell = \ell_\varepsilon \in \mathbb{N}$, such that

$$E_i(t_j) - E_i(t_{j-1}) < \varepsilon/4$$

for $j = 1, \dots, \ell$ and $i = 1, \dots, m$. Upon setting $f_{0,\vartheta}(x) = 0$, $f_{\infty,\vartheta}(x) = \frac{p'_\vartheta(x)}{p_\vartheta(x)} x$, for $x > 0$ and $\vartheta \in K$, we define the brackets

$$\mathcal{H}_{i,j} = \left[f_{t_{j-1},\vartheta_i} - |f_{t_j,\vartheta_i} - f_{t_{j-1},\vartheta_i}| - \delta^\alpha \cdot H^*, f_{t_{j-1},\vartheta_i} + |f_{t_j,\vartheta_i} - f_{t_{j-1},\vartheta_i}| + \delta^\alpha \cdot H^* \right],$$

for $j = 1, \dots, \ell$ and $i = 1, \dots, m$, where $H^*(x) = H(x) \cdot x$, $x > 0$. These brackets cover \mathcal{H}_Θ . Indeed, if $\vartheta \in K$ and $t > 0$ are arbitrary, there exist $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, \ell\}$ such that $\vartheta \in B_\delta(\vartheta_i)$ and $t_{j-1} \leq t < t_j$, so $f_{t,\vartheta} \in \mathcal{H}_{i,j}$ since for every $x > 0$

$$\begin{aligned} |f_{t,\vartheta}(x) - f_{t_{j-1},\vartheta_i}(x)| &\leq |f_{t,\vartheta}(x) - f_{t,\vartheta_i}(x)| + |f_{t,\vartheta_i}(x) - f_{t_{j-1},\vartheta_i}(x)| \\ &= \left| \frac{p'_\vartheta(x)}{p_\vartheta(x)} - \frac{p'_{\vartheta_i}(x)}{p_{\vartheta_i}(x)} \right| \min\{x, t\} \\ &\quad + \left| \frac{p'_{\vartheta_i}(x)}{p_{\vartheta_i}(x)} \right| \left(\min\{x, t\} - \min\{x, t_{j-1}\} \right) \\ &\leq H(x) \cdot x \cdot |\vartheta - \vartheta_i|^\alpha + \left| \frac{p'_{\vartheta_i}(x)}{p_{\vartheta_i}(x)} \right| \left(\min\{x, t_j\} - \min\{x, t_{j-1}\} \right) \\ &\leq \delta^\alpha \cdot H^*(x) + |f_{t_j,\vartheta_i}(x) - f_{t_{j-1},\vartheta_i}(x)|. \end{aligned}$$

Moreover, the brackets $\mathcal{H}_{i,j}$ are ε -brackets with respect to $L^1(\mathbb{P}^X)$, as

$$\begin{aligned} \left\| 2 \left(|f_{t_j,\vartheta_i} - f_{t_{j-1},\vartheta_i}| + \delta^\alpha \cdot H^* \right) \right\|_{L^1(\mathbb{P}^X)} &= 2 \mathbb{E} \left[|f_{t_j,\vartheta_i}(X) - f_{t_{j-1},\vartheta_i}(X)| + \delta^\alpha \cdot H(X)X \right] \\ &= 2(E_i(t_j) - E_i(t_{j-1})) + \frac{\varepsilon}{2} \\ &< \varepsilon. \end{aligned}$$

Hence $\mathcal{N}_{[\cdot]}(\varepsilon, \mathcal{H}_\Theta, L^1(\mathbb{P}^X)) \leq m_\varepsilon \cdot \ell_\varepsilon < \infty$. ■

Proof of Remark 5. From Lemma 3 we know that $\psi_{n,q}(\vartheta) \rightarrow \psi_q(\vartheta)$ \mathbb{P} -a.s., as $n \rightarrow \infty$, for each $\vartheta \in \Theta$. Since $\vartheta_0 \in \Theta^\circ$, there exists a $\delta > 0$ such that $B_{2\delta}(\vartheta_0) \subset \Theta$. Then the closed ball $\overline{B} = \overline{B_\delta(\vartheta_0)}$ also lies in Θ . Denote by $R = \partial B_\delta(\vartheta_0)$ the boundary of that ball. It follows from Lemma 3 that

$$\inf_{\vartheta \in \overline{B}} \psi_{n,q}(\vartheta) \longrightarrow \inf_{\vartheta \in \overline{B}} \psi_q(\vartheta) = 0 \quad \text{and} \quad \inf_{\vartheta \in R} \psi_{n,q}(\vartheta) \longrightarrow \inf_{\vartheta \in R} \psi_q(\vartheta) > 0,$$

both \mathbb{P} -a.s., as $n \rightarrow \infty$, where the positiveness of the last term follows from (7). Now, let $\varepsilon > 0$ and choose $n_0 = n_0(\varepsilon) \in \mathbb{N}$ such that

$$\mathbb{P}\left(\inf_{\vartheta \in \bar{B}} \psi_{n,q}(\vartheta) + \varepsilon_n < \inf_{\vartheta \in R} \psi_{n,q}(\vartheta)\right) \geq 1 - \frac{\varepsilon}{2}, \quad n \geq n_0.$$

Next, note that if $\inf_{\vartheta \in \bar{B}} \psi_{n,q}(\vartheta) + \varepsilon_n < \inf_{\vartheta \in R} \psi_{n,q}(\vartheta)$ then $\psi_{n,q}$ has a local minimum in $B_\delta(\vartheta_0)$ (since $\varepsilon_n > 0$) which, by strict convexity, is the unique global minimum. Additionally, we have $\inf_{\vartheta \in \bar{B}} \psi_{n,q}(\vartheta) + \varepsilon_n < \inf_{\vartheta \in \Theta \setminus B_\delta(\vartheta_0)} \psi_{n,q}(\vartheta)$. On the other hand, if we have the relation $\hat{\vartheta}_{n,q} \in \Theta \setminus B_\delta(\vartheta_0)$, then

$$\inf_{\vartheta \in \Theta \setminus B_\delta(\vartheta_0)} \psi_{n,q}(\vartheta) \leq \psi_{n,q}(\hat{\vartheta}_{n,q}) \leq \inf_{\vartheta \in \Theta} \psi_{n,q}(\vartheta) + \varepsilon_n = \inf_{\vartheta \in \bar{B}} \psi_{n,q}(\vartheta) + \varepsilon_n.$$

Consequently, for all $n \geq n_0$,

$$\begin{aligned} 1 - \frac{\varepsilon}{2} &\leq \mathbb{P}\left(\inf_{\vartheta \in \bar{B}} \psi_{n,q}(\vartheta) + \varepsilon_n < \inf_{\vartheta \in R} \psi_{n,q}(\vartheta)\right) \\ &\leq \mathbb{P}\left(\inf_{\vartheta \in \bar{B}} \psi_{n,q}(\vartheta) + \varepsilon_n < \inf_{\vartheta \in \Theta \setminus B_\delta(\vartheta_0)} \psi_{n,q}(\vartheta)\right) \\ &\leq \mathbb{P}\left(\hat{\vartheta}_{n,q} \in \bar{B}\right). \end{aligned}$$

Since $\{\mathbb{P}^{\hat{\vartheta}_{n,q}} \mid n \leq n_0\}$ is a finite set of measures on \mathbb{R}^d , there exists a compact set $K \subset \mathbb{R}^d$ such that $\mathbb{P}(\hat{\vartheta}_{n,q} \in K) \geq 1 - \frac{\varepsilon}{2}$ for all $n \leq n_0$. The set $K \cap \bar{B} \subset \Theta$ is a compact subset of \mathbb{R}^d and thus also of Θ , for a compact metric space is a compact subset of every metric space it embeds into continuously (see p. 21, Theorem 3, of Kuratowski, 1968). By choice of the sets,

$$\mathbb{P}\left(\hat{\vartheta}_{n,q} \in K \cap \bar{B}\right) \geq 1 - \varepsilon,$$

which is the claim. ■

ACKNOWLEDGEMENTS

The authors express their gratitude to the editor, associate editor as well as two anonymous referees for their careful reading and the helpful suggestions. The authors also thank Francois-Xavier Briol for his reaction to a preprint version of this paper and his insights into the literature on Stein discrepancies. Open access funding enabled and organized by Projekt DEAL.

BIBLIOGRAPHY

- Allison, J. S. & Santana, L. (2015). On a data-dependent choice of the tuning parameter appearing in certain goodness-of-fit tests. *Journal of Statistical Computation and Simulation*, 85, 3276–3288.
- Baringhaus, L., Gürtler, N., & Henze, N. (2000). Weighted integral test statistics and components of smooth tests of fit. *Australian & New Zealand Journal of Statistics*, 42, 179–192.
- Barp, A., Briol, F. X., Duncan, A. B., Girolami, M. A., & Mackey, L. W. (2019). Minimum Stein discrepancy estimators. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., & Garnett, R. (Eds.) *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 12964–12976.
- Beer, G. (1993). *Topologies on Closed and Closed Convex Sets, Mathematics and Its Application*. Kluwer Academic Publishers, Dordrecht.
- Betsch, S. & Ebner, B. (2019a). Fixed point characterizations of continuous univariate probability distributions and their applications. *Annals of the Institute of Statistical Mathematics*, <https://doi.org/10.1007/s10463-019-00735-1>, (to appear in print).

- Betsch, S. & Ebner, B. (2019b). A new characterization of the Gamma distribution and associated goodness-of-fit tests. *Metrika*, 82, 779–806.
- Betsch, S. & Ebner, B. (2020). Testing normality via a distributional fixed point property in the Stein characterization. *TEST*, 29, 105–138.
- Billingsley, P. (1968). Convergence of probability measures. *Wiley Series in Probability and Mathematical Statistics*, John Wiley & Sons, New York.
- Brown, L. D. & Purves, R. (1973). Measurable selections of extrema. *The Annals of Statistics*, 1, 902–912.
- Burr, I. W. (1942). Cumulative frequency functions. *The Annals of Mathematical Statistics*, 13, 215–232.
- Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal of Scientific Computing*, 16, 1190–1208.
- Castaing, C. & Valadier, M. (1977). *Convex Analysis and Measurable Multifunctions. Lecture Notes in Mathematics 580*. Springer-Verlag, Berlin, Heidelberg, New York.
- Chen, L. H. Y., Goldstein, L., & Shao, Q. M. (2011). *Normal approximation by Stein's method*. Springer-Verlag, Berlin Heidelberg.
- Chwialkowski, K., Strathmann, H., & Gretton, A. (2016). A kernel test of goodness of fit. *Proceedings of the 33rd International Conference on Machine Learning –Volume 48, ICML'16*; 2606–2615.
- Cohn, D. L. (2013). *Measure Theory*, 2nd ed., Birkhäuser, New York.
- Döbler, C. (2015). Stein's method of exchangeable pairs for the Beta distribution and generalizations. *Electronic Journal of Probability*, 20, 1–34.
- Gebetsberger, M., Messner, J. W., Mayr, G. J., & Zeileis, A. (2018). Estimation methods for nonhomogeneous regression models: Minimum continuous ranked probability score versus maximum likelihood. *Monthly Weather Review*, 146, 4323–4338.
- Gneiting, T. & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 359–378.
- Gneiting, T., Raftery, A. E., Westveld, A. H., & Goldman, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133, 1098–1118.
- Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69, 243–268.
- Goldstein, L. & Reinert, G. (1997). Stein's method and the zero bias transformation with application to simple random sampling. *The Annals of Applied Probability*, 7, 935–952.
- Gorham, J. & Mackey, L. (2015). Measuring sample quality with Stein's method. In Cortes, C., Lawrence, N.D., Lee, D. D., Sugiyama, M., & Garnett, R. (Eds.) *Advances in Neural Information Processing Systems 28*, Curran Associates, Inc., 226–234.
- Gutmann, M. U. & Hyvärinen, A. (2010). Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Teh, Y. W. & Titterton, M. (Eds.) *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS), Journal of Machine Learning Research –Proceedings Track, JMLR WCP*, Vol. 9, 297–304.
- Gutmann, M. U. & Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13, 307–361.
- Hayakawa, J. & Takemura, A. (2016). Estimation of exponential-polynomial distribution by holonomic gradient descent. *Communications in Statistics –Theory and Methods*, 45, 6860–6882.
- Huggins, J. H. & Mackey, L. (2018). Random feature stein discrepancies. *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, Curran Associates Inc., Red Hook, NY, 1903–1913.
- Hytönen, T., van Neerven, J., Veraar, M., & Weis, L. (2016). *Analysis in Banach Spaces –Volume I: Martingales and Littlewood-Paley Theory, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge/A Series of Modern Surveys in Mathematics*, Vol. 63. Springer International Publishing AG, Cham.
- Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6, 695–709.
- Hyvärinen, A. (2007). Some extensions of score matching. *Computational Statistics & Data Analysis*, 51, 2499–2512.
- Kallenberg, O. (2002). Foundations of Modern Probability, 2nd ed., *Probability and Its Applications*, Springer-Verlag, New York.

- Kleiber, C. & Kotz, S. (2003). *Statistical size distributions in economics and actuarial sciences*. Wiley Series in Probability and Statistics, John Wiley and Sons, Inc., Hoboken, NJ.
- Kraft, D. (1988). *A Software Package for Sequential Quadratic Programming*. Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht, Wiss. Berichtswesen d. DFVLR, Band 88, Ausgabe 28.
- Kumar, D. (2017). The Burr type XII distribution with some statistical properties. *Journal of Data Science*, 15, 509–533.
- Kuratowski, K. (1968). *Topology Volume II*. Academic Press/Polish Scientific Publishers, New York/Warsaw.
- Ley, C. & Swan, Y. (2013). Stein's density approach and information inequalities. *Electronic Communications in Probability*, 18, 1–14.
- Ley, C., Reinert, G., & Swan, Y. (2017a). Distances between nested densities and a measure of the impact of the prior in Bayesian statistics. *The Annals of Applied probability*, 27, 216–241.
- Ley, C., Reinert, G., & Swan, Y. (2017b). Stein's method for comparison of univariate distributions. *Probability Surveys*, 14, 1–52.
- Li, S. Z. (2009). *Markov Random Field Modeling in Image Analysis*, 3rd ed., Springer-Verlag, London.
- Liu, Q., Lee, J. D., & Jordan, M. (2016). A kernelized Stein discrepancy for goodness-of-fit tests. *Proceedings of the 33rd International Conference on Machine Learning –Volume 48, ICML'16*, 276–284.
- Matsuda, T. & Hyvärinen, A. (2019). Estimation of non-normalized mixture models. In Chaudhuri, K. & Sugiyama, M. (Eds.) *Proceedings of Machine Learning Research, PMLR*, Vol. 89, 2555–2563.
- Millar, P. W. (1981). Robust estimation via minimum distance methods. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 55, 73–89.
- Millar, P. W. (1984). A general approach to the optimality of minimum distance estimators. *Transactions of The American Mathematical Society*, 286, 377–418.
- Nakayama, H., Nishiyama, K., Noro, M., Ohara, K., Sei, T., Takayama, N., & Takemura, A. (2011). Holonomic gradient descent and its application to the Fisher–Bingham integral. *Advances in Applied Mathematics*, 47, 639–658.
- Parr, W. C. (1981). Minimum distance estimation: A bibliography. *Communications in Statistics –Theory and Methods*, 10, 1205–1224.
- Parr, W. C. & De Wet, T. (1981). On minimum Cramer–von Mises–norm parameter estimation. *Communications in Statistics –Theory and Methods*, 10, 1149–1166.
- Parr, W. C. & Schucany, W. R. (1980). Minimum distance and robust estimation. *Journal of the American Statistical Association*, 75, 616–624.
- Peköz, E. A. & Röllin, A. (2011). New rates for exponential approximation and the theorems of Rényi and Yaglom. *The Annals of Probability*, 39, 587–608.
- Pfanzagl, J. (1969). On the measurability and consistency of minimum contrast estimates. *Metrika*, 14, 249–272.
- Pihlaja, M., Gutmann, M. U., & Hyvärinen, A. (2010). A family of computationally efficient and simple estimators for unnormalized statistical models. *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI'10*, AUAI Press, Arlington., Catalina Island, CA, 442–449.
- Rodriguez, R. N. (1977). A guide to the Burr type XII distributions. *Biometrika*, 64, 129–134.
- Sahler, W. (1970). Estimation by minimum-discrepancy methods. *Metrika*, 16, 85–106.
- Schmittlein, D. C. (1983). Some sampling properties of a model for income distribution. *Journal of Business & Economic Statistics*, 1, 147–153.
- Schwartz, L. (1973). *Radon Measures on Arbitrary Topological Spaces and Cylindrical Measures*. Oxford University Press, London.
- Shah, A. & Gokhale, D. V. (1993). On maximum product of spacings (MPS) estimation for Burr XII distributions. *Communications in Statistics –Simulation and Computation*, 22, 615–641.
- Singh, S. K. & Maddala, G. S. (1976). A function for size distribution of incomes. *Econometrica*, 44, 963–970.
- Stein, C. (1986). Approximate computation of expectations. *Lecture Notes –Monograph Series 7*, Institute of Mathematical Statistics.
- Tadikamalla, P. R. (1980). A look at the Burr and related distributions. *International Statistical Review/Revue Internationale de Statistique*, 48, 337–344.

- Tenreiro, C. (2019). On the automatic selection of the tuning parameter appearing in certain families of goodness-of-fit tests. *Journal of Statistical Computation and Simulation*, 89, 1780–1797.
- Uehara, M., Kanamori, T., Takenouchi, T., & Matsuda, T. (2020). A unified statistically efficient estimation framework for unnormalized models. In Chiappa, S. & Calandra, R. (Eds.) *Proceedings of Machine Learning Research, PMLR*, Vol. 108, 809–819.
- Uehara, M., Matsuda, T., & Kim, J. K. (2020). Imputation estimators for unnormalized models with missing data. In Chiappa, S. & Calandra, R. (Eds.) *Proceedings of Machine Learning Research, PMLR*, Vol. 108, 831–841.
- van der Vaart, A. W. (1998). *Asymptotic Statistics. Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- van der Vaart, A. W. & Wellner, J. A. (2000). Weak convergence and empirical processes – With applications to statistics. *Springer Series in Statistics*, Springer-Verlag, New York.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C., Polat, I., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., & SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17, 261–272.
- Wang, F. K. & Cheng, Y. F. (2010). Robust regression for estimating the Burr XII parameters with outliers. *Journal of Applied Statistics*, 37, 807–819.
- Widder, D. V. (1959). *The Laplace Transform*, 5th printing, Princeton University Press, Princeton.
- Wingo, D. R. (1983). Maximum likelihood methods for fitting the Burr type XII distribution to life test data. *Biometrical Journal*, 25, 77–84.
- Wingo, D. R. (1993). Maximum likelihood methods for fitting the Burr type XII distribution to multiply (progressively) censored life test data. *Metrika*, 40, 203–210.
- Wolfowitz, J. (1957). The minimum distance method. *The Annals of Mathematical Statistics*, 28, 75–88.
- Yang, J., Liu, Q., Rao, V., & Neville, J. (2018). Goodness-of-fit testing for discrete distributions via Stein discrepancy. In Dy, J. & Krause, A. (Eds.) *Proceedings of the 35th International Conference on Machine Learning – Volume 80, PMLR, Proceedings of Machine Learning Research*, 5561–5570.
- Yu, S., Drton, M., & Shojaie, A. (2019). Generalized score matching for non-negative data. *Journal of Machine Learning Research*, 20, 1–70.
- Zhu, C., Byrd, R. H., Lu, P., & Nocedal, J. (1997). Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software*, 23, 550–560.

Received 01 October 2019

Accepted 16 April 2020