Ann Inst Stat Math (2015) 67:573–594
DOI 10.1007/s10463-014-0465-x

The sinh-arcsinhed logistic family of distributions: properties and inference

Arthur Pewsey · Toshihiro Abe

Received: 11 April 2013 / Revised: 17 January 2014 / Published online: 4 May 2014

Abstract The sinh-arcsinh transform is used to obtain a flexible four-parameter model that provides a natural framework with which to perform inference robust to wide-ranging departures from the logistic distribution. Its basic properties are established and its distribution and quantile functions, and properties related to them, shown to be highly tractable. Two important subfamilies are also explored. Maximum likelihood estimation is discussed, and reparametrisations designed to reduce the asymptotic correlations between the maximum likelihood estimates provided. A likelihood-ratio test for logisticness, which outperforms standard empirical distribution function based tests, follows naturally. The application of the proposed model and inferential methods is illustrated in an analysis of carbon fibre strength data. Multivariate extensions of the model are explored.

Keywords Copula · L_U distribution · Quantile measures · Sinh-arcsinh transform · Skewness · Tailweight · Test for logisticness

---

Electronic supplementary material The online version of this article (doi:10.1007/s10463-014-0465-x) contains supplementary material, which is available to authorized users.

A. Pewsey (✉)
Mathematics Department, Escuela Politécnica, University of Extremadura, 10003 Cáceres, Spain
e-mail: apewsey@unex.es

T. Abe
Department of Management Science, Faculty of Engineering, Tokyo University of Science,
1-3 Kagurazaka, Shinjuku-ku, Tokyo 162-0825, Japan
e-mail: abe@ms.kagu.tus.ac.jp

1 Introduction

Let Z denote a standardised absolutely continuous random variable that is symmetric about the origin, with distribution function $F_Z(z)$, quantile function $Q_Z(u) = F_Z^{-1}(u), 0 < u < 1$, and density $f_Z(z)$. Jones and Pewsey (2009) define the canonical sinh-arcsinhed counterpart of Z, $X_{\varepsilon,\delta}$, through the application of the sinh-arcsinh transformation

$$Z = S_{\varepsilon,\delta}(X_{\varepsilon,\delta}) = \sinh\{\delta \sinh^{-1}(X_{\varepsilon,\delta}) - \varepsilon\}, \tag{1}$$

and show that $\varepsilon \in \mathbb{R}$ is a skewness parameter in the sense of van Zwet's(1964) skewness ordering, and $\delta > 0$ is a parameter controlling tailweight. For $\delta < 1$ ($\delta > 1$), the tails of $X_{\varepsilon,\delta}$ are heavier (lighter) than those of the base random variable Z.

Jones and Pewsey (2009) mainly provide results for the sinh-arcsinhed normal distribution obtained using a standard normal Z. However, motivated by the idea of deriving powerful likelihood-ratio tests for logisticness, in their Section 8.3 they refer to the possibility of substituting a standard logistic Z instead. Historically, the logistic distribution has been employed as a mathematically tractable alternative to the normal, its distribution and quantile functions having simple closed forms. A book-length treatment of the logistic distribution, its generalisations, extensions and applications is provided by Balakrishnan (1992). Chapter 23 of Johnson et al. (1994) is also an essential extensive source. Of the related distributions mentioned in the latter, those proposed by Tadikamalla and Johnson (1982) and Johnson and Tadikamalla (1992), obtained by transforming a logistic random variable, above all their L_U system based on the inverse-sinh transformation, are particularly relevant to the approach adopted here. More recently, Nadarajah (2009) studied a skew-logistic distribution generated using the perturbation-based construction of Azzalini (1985); see also Gupta and Kundu (2010).

In this paper, we extend the general results available for distributions obtained using the sinh-arcsinh transformation and consider the four-parameter sinh-arcsinhed logistic (SAS-logistic) family of distributions in its own right. The latter provides an attractive alternative to the SAS-normal family of Jones and Pewsey (2009) and SAS-t family of Rosco et al. (2011), which have the normal distribution as a central case, and a limiting case, respectively. The new family is appealing in terms of the tractability of its distribution and quantile functions and those properties related to them, especially its ease of simulation: an important consideration for computer intensive approaches to inference, which we illustrate too. It also provides a natural framework for performing inference robust to wide-ranging departures from the logistic distribution. One particularly important inferential spin-off is a likelihood-ratio test that outperforms the empirical distribution function based tests proposed by Stephens (1979); particularly the Anderson-Darling test known to generally provide a powerful test for logisticness.

Section 2 focuses on the sinh-arcsinh transformation and its general properties. In Sect. 3, we define the SAS-logistic model and consider properties such as its modality, limiting distributions, tailweight behaviour, moments, quantile-based measures of location, scale, skewness and kurtosis, and two important subfamilies. Likelihood-based inference is discussed in Sect. 4. There, we consider maximum likelihood estimation, reparametrisations that reduce the dependencies between the maximum likelihood estimates, and likelihood-ratio tests for logisticness and symmetry. An analysis of carbon fibre strength data, presented in Sect. 5, illustrates the application of the model. Potential multivariate extensions of the SAS-logistic distribution are discussed in Sect. 6. Conclusions are drawn in Sect. 7. Results for use in likelihood-based inference are collected together in an Appendix. A proof of the modality of the SAS-logistic distribution, together with results for the moments of two of its subfamilies, is given in the Supplementary Material available from the journal's website.

2 General properties of the sinh-arcsinh transformation

The following three representations of the sinh-arcsinh transformation in (1) prove useful in various mathematical contexts:

$$\begin{aligned} S_{\varepsilon,\delta}(x) &= \tfrac{1}{2}\left[e^{-\varepsilon}\exp\{\delta\sinh^{-1}(x)\} - e^{\varepsilon}\exp\{-\delta\sinh^{-1}(x)\}\right] \\ &= \tfrac{1}{2}\left[e^{-\varepsilon}\{(x^2+1)^{1/2}+x\}^{\delta} - e^{\varepsilon}\{(x^2+1)^{1/2}+x\}^{-\delta}\right] \\ &= \tfrac{1}{2}\left[e^{-\varepsilon}\{(x^2+1)^{1/2}+x\}^{\delta} - e^{\varepsilon}\{(x^2+1)^{1/2}-x\}^{\delta}\right]. \end{aligned} \tag{2}$$

Inverting (1),

$$X_{\varepsilon,\delta} = S^{-1}_{\varepsilon,\delta}(Z) = \sinh[\delta^{-1}\{\sinh^{-1}(Z)+\varepsilon\}] = S_{-\varepsilon/\delta,1/\delta}(Z). \tag{3}$$

The distribution function of $X_{\varepsilon,\delta}$, $F_{\varepsilon,\delta}(x)$, is related to that of Z through

$$\begin{aligned} F_{\varepsilon,\delta}(x) &= P(X_{\varepsilon,\delta} \le x) = P(S_{-\varepsilon/\delta,1/\delta}(Z) \le x) = P(Z \le S^{-1}_{-\varepsilon/\delta,1/\delta}(x)) \\ &= P(Z \le S_{\varepsilon,\delta}(x)) = F_Z(S_{\varepsilon,\delta}(x)), \end{aligned} \tag{4}$$

whilst its quantile function is given by

$$Q_{\varepsilon,\delta}(u) = F^{-1}_{\varepsilon,\delta}(u) = S^{-1}_{\varepsilon,\delta}(F_Z^{-1}(u)) = S_{-\varepsilon/\delta,1/\delta}(Q_Z(u)), \quad 0 < u < 1. \tag{5}$$

As Z is symmetric about 0, the median of a sinh-arcsinh transformed random variable is $Q_{\varepsilon,\delta}(1/2) = S_{-\varepsilon/\delta,1/\delta}(Q_Z(1/2)) = S_{-\varepsilon/\delta,1/\delta}(0) = \sinh(\varepsilon/\delta)$. Differentiating (4) with respect to x, the density of $X_{\varepsilon,\delta}$ is available as

$$\begin{aligned} f_{\varepsilon,\delta}(x) &= F'_{\varepsilon,\delta}(x) = F'_Z(S_{\varepsilon,\delta}(x))S'_{\varepsilon,\delta}(x) = f_Z(S_{\varepsilon,\delta}(x))S'_{\varepsilon,\delta}(x) \\ &= f_Z(S_{\varepsilon,\delta}(x))\frac{\delta C_{\varepsilon,\delta}(x)}{(1+x^2)^{1/2}}, \end{aligned} \tag{6}$$

where $C_{\varepsilon,\delta}(x) = \cosh\{\delta\sinh^{-1}(x) - \varepsilon\} = \{1 + S^2_{\varepsilon,\delta}(x)\}^{1/2}$. Using the third identity in (2), it follows that $f_{-\varepsilon,\delta}(x) = f_{\varepsilon,\delta}(-x)$. Thus, densities with negative values of the skewness parameter ε are the reflections about the origin of their counterparts with positive ε. In applied work, one will generally be interested in the location-scale extension of $X_{\varepsilon,\delta}$, $X_{\xi,\eta,\varepsilon,\delta} = \xi + \eta X_{\varepsilon,\delta}$, with location and scale parameters, $\xi \in \mathbb{R}$ and $\eta > 0$, respectively. $X_{\xi,\eta,\varepsilon,\delta}$ has density $\eta^{-1} f_{\varepsilon,\delta}((x-\xi)/\eta)$.

3 The SAS-logistic family

3.1 Definition and basic properties

Substituting a standard logistic random variable, L, for the generic symmetric random variable Z in (3), a canonical sinh-arcsinhed logistic random variable is defined through the relation $X_{\varepsilon,\delta} = S_{-\varepsilon/\delta,1/\delta}(L)$. Using (4)–(6), $X_{\varepsilon,\delta}$ has distribution function

$$F_{\varepsilon,\delta}(x) = \frac{1}{1+e^{-S_{\varepsilon,\delta}(x)}} = \tfrac{1}{2}\{1+\tanh(\tfrac{1}{2}S_{\varepsilon,\delta}(x))\}, \quad -\infty < x < \infty, \tag{7}$$

quantile function

$$Q_{\varepsilon,\delta}(u) = S_{-\varepsilon/\delta,1/\delta}(\log\{u/(1-u)\}), \tag{8}$$

and density

$$f_{\varepsilon,\delta}(x) = \frac{\delta C_{\varepsilon,\delta}(x)}{(1+x^2)^{1/2}} \frac{e^{-S_{\varepsilon,\delta}(x)}}{(1+e^{-S_{\varepsilon,\delta}(x)})^2} = \frac{\delta C_{\varepsilon,\delta}(x)}{4(1+x^2)^{1/2}} \operatorname{sech}^2(\tfrac{1}{2}S_{\varepsilon,\delta}(x)). \tag{9}$$

Equations (7) and (8) are appealing in the sense that, unlike their SAS-normal and SAS-*t* counterparts, they are closed-form expressions and hence their computation requires neither numerical integration nor root finding.

With regard to simulation, consider a uniform random variable on (0, 1), U. Applying the probability integral transform, $L = \log(U/(1-U))$ is a standard logistic random variable and hence

$$X_{\varepsilon,\delta} = \tfrac{1}{2}[e^{\varepsilon/\delta}\{(L^2+1)^{1/2}+L\}^{1/\delta} - e^{-\varepsilon/\delta}\{(L^2+1)^{1/2}+L\}^{-1/\delta}] \tag{10}$$

is a SAS-logistic random variable. Note that the second term between square brackets in (10) is minus the inverse of the first.

3.2 Modality

Although, as Jones and Pewsey (2008) show, all SAS-normal densities are unimodal, more generally sinh-arcsinh densities are not always unimodal. Rosco et al. (2011) show, for instance, that SAS-*t* distributions can be uni- or bimodal. In Section SM1 of the Supplementary Material, we prove that (9) is always *unimodal* if $\delta \leq \sqrt{2}$. It is always *bimodal* if $\delta \geq \delta^{\dagger}$, where, to four decimal places,

$$\delta^{\dagger} = 1/\max_{0<u<u^{\dagger}} h(u) = 5.6876,$$

$u = -\varepsilon + \delta \sinh^{-1}(x)$, $h(u) = \tanh u - \tanh((\sinh u)/2) \cosh u$ and $u^\dagger = \max\{h^{-1}(0)\} = 0.9459$. If $\sqrt{2} < \delta < \delta^\dagger$, then the density is

$$unimodal \text{ if } |\varepsilon| \geq \max_{0<u<u^\dagger} g_\delta(u) \quad \text{and} \quad bimodal \text{ if } |\varepsilon| < \max_{0<u<u^\dagger} g_\delta(u),$$

where $g_\delta(u) = -u + \delta \tanh^{-1}(\delta h(u))$. Graphs of SAS-logistic densities illustrating these results are portrayed in Fig. 1. Its panel (a) displays symmetric cases (with $\varepsilon = 0$) for four δ-values and confirms the unimodality of the density for $\delta \leq \sqrt{2}$ and bimodality for $\delta > \sqrt{2}$. The densities in panel (c)–(e) are all unimodal, corresponding to δ-values of 0.5, 1 and $\sqrt{2}$, respectively. In its panel (f), corresponding to $\delta = 2$, the density is bimodal when $\varepsilon = 0$ but soon becomes unimodal as ε increases. Its panel (g), with $\delta = 6$, illustrates the fact that for $\delta \geq \delta^\dagger$ the density is always bimodal whatever the value of ε. We will comment on the densities in panel (b), obtained as $\varepsilon \to \infty$, in the following subsection.

As mentioned in the Introduction, ε is a parameter that controls skewness and δ controls tailweight. From a consideration of the densities in Fig. 1, it is evident, however, that δ also affects the distribution's dispersion and ε also affects its central location. Clearly, the modality of the distribution depends on the values taken by both parameters. Fig. 2 displays, in black, the $(\delta^*, \varepsilon^*)$ subregion, where $\delta^* = \delta/(1+\delta)$ and $\varepsilon^* = \varepsilon/(1+\varepsilon)$, for which symmetric or positively skewed SAS-logistic densities are bimodal. Its reflection about the horizontal axis is the corresponding region for which symmetric or negatively skewed cases are bimodal.

We view the potential bimodality of SAS-logistic distributions as generally being an unappealing property, primarily because we favour the use of two-piece mixtures of unimodal distributions as a means of modelling bimodality; the interpretation of the latter usually being more straightforward. One might therefore consider the unimodal subfamily associated with the constraints

$$\delta \leq \sqrt{2} \text{ and } \varepsilon \in \mathbb{R} \quad \text{or} \quad \sqrt{2} < \delta < \delta^\dagger \text{ and } |\varepsilon| \geq \max_{0<u<u^\dagger} g_\delta(u).$$

Alternatively, in order to ensure unimodality for any potential value of ε, one could restrict the parameter space still further to $\delta \leq \sqrt{2}$ and $\varepsilon \in \mathbb{R}$. The latter subfamily will always contain the symmetric case associated with any given δ-value. As is clear from a consideration of the densities portrayed in Fig. 1, this restricted unimodal subfamily contains densities displaying wide-ranging levels of skewness and tailweight. Finally, from an applications perspective, there is an argument for not restricting the parameter space at all; with bimodal solutions interpreted as indicating the need to explore the fit of two-piece mixtures of unimodal distributions.

3.3 Limiting densities, tailweight and moments

First, consider the limiting densities as $\varepsilon \to \pm\infty$. Because $f_{-\varepsilon,\delta}(x) = f_{\varepsilon,\delta}(-x)$, it suffices to consider the limiting density as $\varepsilon \to \infty$, denoted by $f_{\infty,\delta}$. After suitable

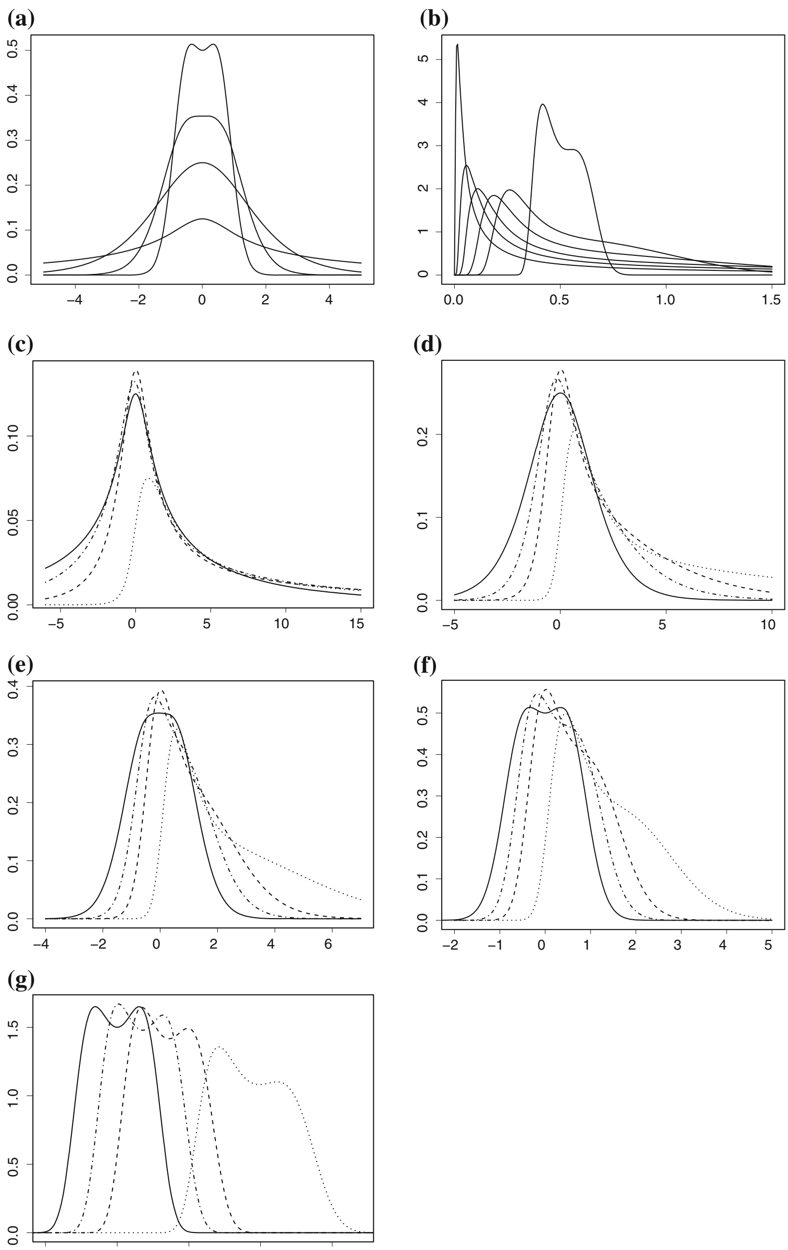

Fig. 1 The SAS-logistic density (9) for: (**a**) $\varepsilon = 0$ and, in order of increasing height at the origin, $\delta = 0.5, 1, \sqrt{2}, 2$; (**b**) $\varepsilon \to \infty$ and, from left to right, $\delta = 0.5, 0.75, 1, \sqrt{2}, 2, 6$; (**c**) $\delta = 0.5$ and, from left to right, $\varepsilon = 0, 0.5, 1, 2$; (**d**) $\delta = 1$ and, from left to right, $\varepsilon = 0, 0.5, 1, 2$; (**e**) $\delta = \sqrt{2}$ and, from left to right, $\varepsilon = 0, 0.5, 1, 2$; (**f**) $\delta = 2$ and, from left to right, $\varepsilon = 0, 0.5, 1, 2$; (**g**) $\delta = 6$ and, from left to right, $\varepsilon = 0, 1, 2, 5$. The densities in panels (**c**)–(**g**) have been drawn with different line types to aid visualisation

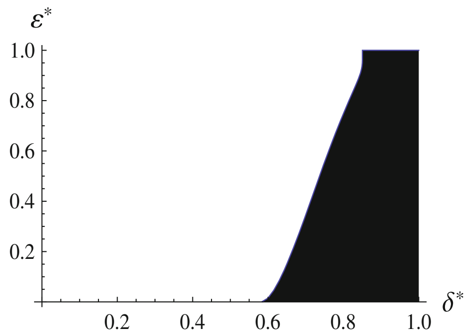

Fig. 2 Subregion, in black, of $(\delta^*, \varepsilon^*)$-values, where $\delta^* = \delta/(1+\delta) \in (0, 1)$ and $\varepsilon^* = \varepsilon/(1+\varepsilon) \in [0, 1)$, for which symmetric ($\varepsilon^* = 0$) or positively skewed ($\varepsilon^* > 0$) SAS-logistic densities are bimodal

standardisation of location and scale, one obtains

$$f_{\infty,\delta}(y) = \frac{\delta}{4y}\cosh(\delta \log 2y)\,\mathrm{sech}^2(\tfrac{1}{2}\sinh(\delta \log 2y)).$$

This is the density of $Y = \exp\{\sinh^{-1}(L)/\delta\}/2$ where L is standard logistic. Examples of this density for a range of δ-values are portrayed in Fig. 1(b). Those that are unimodal are very skew indeed.

The tails of the base logistic distribution are simple exponential. As $|x| \to \infty$, $S_{\varepsilon,\delta}(x) \approx 2^{\delta-1}\mathrm{sgn}(x)\exp\{-\mathrm{sgn}(x)\varepsilon\}|x|^{\delta}$ and $C_{\varepsilon,\delta}(x) \approx 2^{\delta-1}\exp\{-\mathrm{sgn}(x)\varepsilon\}|x|^{\delta}$, where $\mathrm{sgn}(x)$ denotes the sign of x. Making use of these results,

$$f_{\varepsilon,\delta}(|x|) \approx e^{-\mathrm{sgn}(x)\varepsilon}|x|^{\delta-1}\exp\{-e^{-\mathrm{sgn}(x)\varepsilon}|x|^{\delta}\}. \tag{11}$$

These are Weibull-type tails, their relative scales being $e^{\pm\varepsilon}$. Because of this tail behaviour, the moments of $X_{\varepsilon,\delta}$ about the origin, $E(X^k_{\varepsilon,\delta})$, $k = 1, 2, \ldots$, all exist, and are given by

$$E[X^k_{\varepsilon,\delta}] = \int_{-\infty}^{\infty} x^k f_{X_{\varepsilon,\delta}}(x)dx = \int_{-\infty}^{\infty} x^k S'_{\varepsilon,\delta}(x) f_L(S_{\varepsilon,\delta}(x))dx$$
$$= \int_{-\infty}^{\infty} x^k dF_L(S_{\varepsilon,\delta}(x)).$$

Setting $u = F_L(S_{\varepsilon,\delta}(x)), 0 < u < 1$, $F_L^{-1}(u) = Q_L(u) = S_{\varepsilon,\delta}(x)$ and hence $x = S_{-\varepsilon/\delta,1/\delta}(Q_L(u))$. Thus,

$$E[X^k_{\varepsilon,\delta}] = \int_0^1 S^k_{-\varepsilon/\delta,1/\delta}\,(\log\{u/(1-u)\})\,du,$$

which must generally be computed numerically. In Section SM2 of the Supplementary Material, we provide results for the moments of the skew-logistic and symmetric subfamilies corresponding to $\delta = 1$ and $\varepsilon = 0$, respectively.

3.4 Quantile-based measures of location, scale, skewness and kurtosis

Given the simple form of the quantile function (8) and the general lack of analytic expressions for the moments, it is natural to consider measures based on quantiles as summaries of the main characteristics of the SAS-logistic family.

As noted in Section 3.1, the median of any canonical sinh-arcsinhed distribution is $\sinh(\varepsilon/\delta)$. Using the first identity in (2) and the fact that the sinh function is odd, it follows that

$$S_{-\varepsilon/\delta,1/\delta}(x) - S_{-\varepsilon/\delta,1/\delta}(-x) = 2\cosh(\varepsilon/\delta)\sinh(\delta^{-1}\sinh^{-1}(x))$$

and

$$S_{-\varepsilon/\delta,1/\delta}(x) + S_{-\varepsilon/\delta,1/\delta}(-x) = 2\sinh(\varepsilon/\delta)\cosh(\delta^{-1}\sinh^{-1}(x)).$$

Thus, the interquartile range of $X_{\varepsilon,\delta}$ is given by

$$\begin{aligned} Q_{\varepsilon,\delta}(3/4) - Q_{\varepsilon,\delta}(1/4) &= S_{-\varepsilon/\delta,1/\delta}(\log 3) - S_{-\varepsilon/\delta,1/\delta}(-\log 3) \\ &= 2\cosh(\varepsilon/\delta)\sinh(\delta^{-1}\sinh^{-1}(\log 3)), \end{aligned} \tag{12}$$

whilst the quantile-based skewness coefficient of Bowley,

$$B_{\varepsilon,\delta} = \frac{Q_{\varepsilon,\delta}(3/4) - 2Q_{\varepsilon,\delta}(1/2) + Q_{\varepsilon,\delta}(1/4)}{Q_{\varepsilon,\delta}(3/4) - Q_{\varepsilon,\delta}(1/4)},$$

reduces to

$$\begin{aligned} B_{\varepsilon,\delta} &= \frac{S_{-\varepsilon/\delta,1/\delta}(\log 3) + S_{-\varepsilon/\delta,1/\delta}(-\log 3) - 2\sinh(\varepsilon/\delta)}{2\cosh(\varepsilon/\delta)\sinh(\delta^{-1}\sinh^{-1}(\log 3))} \\ &= \frac{\sinh(\varepsilon/\delta)\{\cosh(\delta^{-1}\sinh^{-1}(\log 3)) - 1\}}{\cosh(\varepsilon/\delta)\sinh(\delta^{-1}\sinh^{-1}(\log 3))} \\ &= \tanh(\varepsilon/\delta)\{\coth(\delta^{-1}\sinh^{-1}(\log 3)) - \operatorname{csch}(\delta^{-1}\sinh^{-1}(\log 3))\}. \end{aligned} \tag{13}$$

The limits of $B_{\varepsilon,\delta}$ are:

$$\lim_{\varepsilon\to 0} B_{\varepsilon,\delta} = \lim_{\delta\to\infty} B_{\varepsilon,\delta} = 0;$$
$$\lim_{\delta\to 0} B_{\varepsilon,\delta} = \operatorname{sign}(\varepsilon);$$
$$\lim_{\varepsilon\to\infty} B_{\varepsilon,\delta} = \operatorname{sign}(\delta)(\coth(\delta^{-1}\sinh^{-1}(\log 3)) - \operatorname{csch}(\delta^{-1}\sinh^{-1}(\log 3))).$$

The kurtosis measure of Moors (1988),

$$M_{\varepsilon,\delta} = \frac{Q_{\varepsilon,\delta}(7/8) - Q_{\varepsilon,\delta}(5/8) + Q_{\varepsilon,\delta}(3/8) - Q_{\varepsilon,\delta}(1/8)}{Q_{\varepsilon,\delta}(3/4) - Q_{\varepsilon,\delta}(1/4)},$$

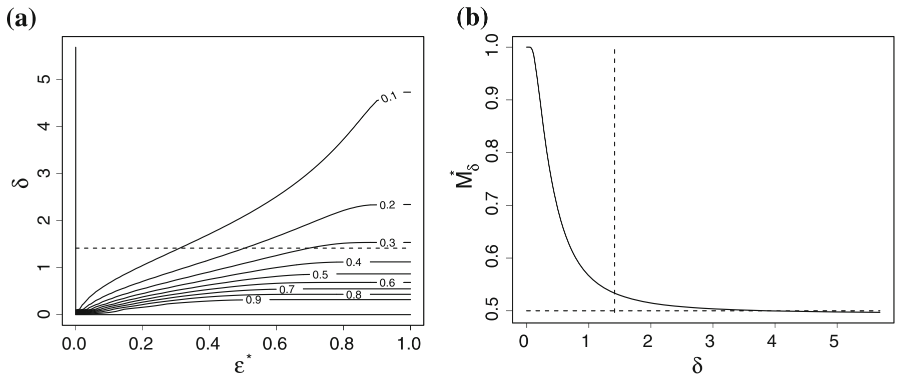

Fig. 3 Plots, for SAS-logistic distributions with density (9) and $0 < \delta < 5.6876$, of: (**a**) contours of Bowley's skewness coefficient (13) as a function of $0 < \varepsilon^* = \varepsilon/(1+\varepsilon) < 1$ and δ; (**b**) the standardised version, $M^*_\delta = M_\delta/(1+M_\delta)$, of Moors' kurtosis measure (14) as a function of δ. In (a), the solid horizontal and vertical lines are the contours for values of 1 and 0 of Bowley's coefficient, respectively. The dashed line delimits $\delta = \sqrt{2}$. In (**b**), the horizontal dashed line delimits the value of the standardised measure for the uniform distribution, 0.5, and the vertical dashed line delimits $\delta = \sqrt{2}$

simplifies to

$$M_\delta = \frac{\sinh(\delta^{-1}\sinh^{-1}(\log 7)) - \sinh(\delta^{-1}\sinh^{-1}(\log(5/3)))}{\sinh(\delta^{-1}\sinh^{-1}(\log 3))}, \tag{14}$$

the latter notation reflecting the fact that Moors' measure only depends on the value taken by the tailweight parameter δ. In fact, as shown in Jones et al. (2011), all quantile-based measures of kurtosis that involve only (possibly scaled) differences between quantile function values of the form $Q(u) - Q(1-u)$ are skewness invariant for *all* sinh-arcsinhed distributions. The limits of M_δ are:

$$\lim_{\delta\to 0} M_\delta = \infty;$$

$$\lim_{\delta\to\infty} M_\delta = \frac{\sinh^{-1}(\log 7) - \sinh^{-1}(\log(5/3))}{\sinh^{-1}(\log 3)} \approx 0.9778.$$

Fig. 3 displays plots, for SAS-logistic distributions with $0 < \delta < \delta^\dagger = 5.6876$, of (13), for $\varepsilon \geq 0$, and a standardised version of (14). As is evident from (13), the effect of ε on $B_{\varepsilon,\delta}$, amplified by small values of the tailweight parameter, δ, enters via the term $\tanh(\varepsilon/\delta)$. The behaviour of this term is reflected in Fig. 3(a), the skewness hardly increasing at all for $\varepsilon^* = \varepsilon/(1+\varepsilon) > 0.7$ ($\varepsilon > 2.33$). Fig. 3(b) illustrates that all possible values of $M^*_\delta = M_\delta/(1+M_\delta)$, and hence M_δ, above that for the logistic distribution ($M^*_1 = 0.5664$, $M_1 = 1.3063$) are attainable. The minimum value reached is $M^*_{\delta^\dagger} = 0.4970$ ($M_{\delta^\dagger} = 0.9881$), marginally *below* that of $M^*_\delta = 0.5$ ($M_\delta = 1$) for the uniform distribution. This might at first seem surprising but remember that the SAS-density with $\delta = \delta^\dagger$ and $\varepsilon = 0$ is bimodal. Overall, the loss in skewness and kurtosis flexibility corresponding to the restriction $\delta \le \sqrt{2}$, ensuring unimodality, appears minor.

3.5 Symmetric and skew-logistic subfamilies

There are two subfamilies that are of particular interest: the symmetric one obtained when $\varepsilon = 0$ and the skew-logistic subfamily corresponding to $\delta = 1$.

Examples of densities from the symmetric subfamily are provided in Fig. 1(a). As proved in Section SM1 of the Supplementary Material, such densities are unimodal if $\delta \le \sqrt{2}$ and are bimodal otherwise. They range from the very heavy-tailed when δ is close to 0, through the logistic when $\delta = 1$, to the wide-bodied, bimodal and light-tailed when δ is large. Whatever the value of δ, the mean, median and mode all equal 0, as are all the odd moments. Results for the even moments are provided in Section SM2 of the Supplementary Material. The interquartile range and Bowley's coefficient of skewness, (12) and (13), reduce to $2\sinh(\delta^{-1}\sinh^{-1}(\log 3))$ and 0, respectively. As Jones and Pewsey (2009) have proved, δ acts as a kurtosis parameter in the sense of van Zwet's (1964) ordering for any symmetric sinh-arcsinhed distribution. As we have seen above, it is also the parameter which most strongly regulates the distribution's modality. When δ is small, $C_{0,\delta}(x) \approx 1$ and $S_{0,\delta}(x) \approx \delta \sinh^{-1}(x)$, and substituting these results in (9), one obtains

$$
\begin{aligned}
f_{0,\delta}(x) &\approx \frac{\delta}{(1+x^2)^{1/2}} \frac{e^{-\delta \sinh^{-1}(x)}}{(1+e^{-\delta \sinh^{-1}(x)})^2} \\
&= \frac{\delta}{(1+x^2)^{1/2}} \frac{\{x+(1+x^2)^{1/2}\}^{\delta}}{[1+\{x+(1+x^2)^{1/2}\}^{\delta}]^2}.
\end{aligned}
$$

This is the density of the symmetric subset of Tadikamalla and Johnson (1982) L_U distribution. Symmetric L_U distributions are those of X_δ, defined via the transformation $L = \delta \sinh^{-1}(X_\delta)$, and have tails that are heavier than those of the logistic.

All the members of the second subfamily, obtained when $\delta = 1$, are unimodal. Applying (3) together with the third equality in (2), $X_{\varepsilon,1} = S_{-\varepsilon,1}(L) = L\cosh\varepsilon + (1+L^2)^{1/2}\sinh\varepsilon$. Such random variables have densities like those displayed in Fig. 1(d). They are true skew-logistic densities in the sense that, as is evident from (11) with $\delta = 1$, both of their tails are simple exponential, like those of the base logistic distribution. This is not the case for the skew-logistic densities of Nadarajah (2009), for which the weight in one of the tails is altered through perturbation by the standard logistic distribution function. Moment results are provided in Section SM2 of the Supplementary Material. The interquartile range and Bowley's skewness coefficient for $X_{\varepsilon,1}$ are $\log(9)\cosh(\varepsilon)$ and $\tanh(\varepsilon)\cosh(\sinh^{-1}(\log 3))/\log 3$, respectively. Because of its skewness invariance, Moors' kurtosis measure is the same as that for the logistic distribution, namely 1.3063 to four decimal places.

4 Likelihood-based inference

4.1 Maximum likelihood estimation and reparametrisation

The log-likelihood for a random sample, $X_1, \ldots, X_n$, drawn from the four-parameter SAS-logistic distribution with density $\eta^{-1} f_{\varepsilon,\delta}((x-\xi)/\eta)$ is

$$\ell(\xi, \eta, \varepsilon, \delta) = n(\log \delta - \log 4 - \log \eta) + \sum_{i=1}^{n} \Big[\log C_{\varepsilon,\delta}(Y_i) - \tfrac{1}{2}\log(1+Y_i^2) - 2\log\cosh(\tfrac{1}{2} S_{\varepsilon,\delta}(Y_i))\Big], \tag{15}$$

where $Y_i = (X_i - \xi)/\eta$. The score equations and the elements of the observed information matrix evaluated at the maximum likelihood solution are given in the Appendix. There are no closed-form expressions for the maximum likelihood estimates and (15) must be maximised numerically. To achieve this, we have made use of both the Nelder-Mead (Nelder and Mead 1965) and L-BFGS-B (Byrd et al. 1995) methods of optimisation available within R's general purpose optimisation function `optim`. Other gradient-based methods of optimisation, as well as simulated annealing, are also available as options of `optim`. Whichever method is employed, it is sensible to use a range of different starting values in an attempt to ensure that the global maximum is identified because the various methods can converge to local maxima. The L-BFGS-B method is particularly attractive because one of its options returns a numerical approximation to the Hessian matrix. The quality of the approximation should be good but can always be checked by comparing the Hessian matrix returned with -1 times the observed information matrix calculated using the results in Appendix C2. The observed information matrix (or its approximation) can be inverted to obtain (an approximation to) the asymptotic covariance matrix, which, in turn, can be transformed into (an approximation to) the asymptotic correlation matrix for the maximum likelihood estimates $\hat{\xi}$, $\hat{\eta}$, $\hat{\varepsilon}$ and $\hat{\delta}$.

Extensive numerical work for both the SAS-normal and SAS-logistic distribution indicates that the asymptotic correlation matrix tends towards singularity as the tailweight parameter, δ, increases; its rank tending to 2. The problem is mainly a consequence of large asymptotic correlations between $\hat{\eta}$ and $\hat{\delta}$ and between $\hat{\xi}$ and $\hat{\varepsilon}$, and kicks in for $\delta \geq 2$. This behaviour is evidence of light-tailed cases of the model being non-identifiable. For low-δ, i.e. heavy-tailed unimodal, cases, singularity of the information matrix, and hence non-identifiability, is not an issue. Using $\eta_\delta = \eta/\delta$ instead of η, as suggested in Jones and Pewsey (2009), results in an asymptotic correlation between $\hat{\eta}_\delta$ and $\hat{\delta}$ that is somewhat lower than that between $\hat{\eta}$ and $\hat{\delta}$, but the asymptotic correlation matrix for $(\hat{\xi}, \hat{\eta}_\delta, \hat{\varepsilon}, \hat{\delta})$ is still close to singular, now with rank 3, for $\delta \geq 2$. In principal, the Gram-Schmidt orthogonalisation process proposed by Rotnitzky et al. (2000) can be employed to obtain an orthogonal reparametrisation. Instead, we sought reparametrisations in which the parameters had clear interpretations. As the singularity problem is due to the correlations between the maximum likelihood estimates of the location and skewness parameters, and between those for the scale and tailweight parameters, we investigated replacing ξ and η by location and scale parameters that were robust to changes in skewness and tailweight, respectively. This logic led us to consider parametrisations in which ξ is replaced by the median $M = \xi + \eta \sinh(\varepsilon/\delta)$, and/or η is replaced by the interquartile range $IQR = 2\eta \cosh(\varepsilon/\delta) \sinh(\delta^{-1} \sinh^{-1}(\log 3))$. In our investigations, we also considered replacing ε and δ by the quantile-based skewness and kurtosis measures $B_{\varepsilon,\delta}$ and M_δ, respectively, but these extra substitutions did not lead to any general improvement in the properties of the asymptotic correlation matrix.

The results from our numerical investigations indicate that no benefit accrues from reparametrisation when samples are drawn from extremely heavy-tailed symmetric SAS-logistic distributions. For samples from heavy- and very heavy-tailed cases, the $(\xi, IQR, \varepsilon, \delta)$ parametrisation should be used. For samples from all other SAS-logistic distributions, with $\delta \geq 1$, the $(M, IQR, \varepsilon, \delta)$ parametrisation should be employed. Note that these recommendations do not lead to full parameter orthogonality but much improved properties of the asymptotic correlation matrices of the maximum likelihood estimates. In particular, the maximum likelihood estimate of the location parameter (ξ or M, as appropriate), which in applications will generally be the parameter of main interest, will often be close to orthogonal to those of the other three parameters. In applications, it may be possible to identify which of the three cases described above applies from a consideration of the tail behaviour manifested in a histogram of the data. However, the ability to do so will heavily depend on the size of the sample being considered. Should there be any doubt over which parametrisation to use, all three can be applied and the resulting (estimated) asymptotic correlation matrices compared. The R code for fitting the SAS-logistic distribution under the original parametrisation is easily amended to accommodate the other two. R code for fitting the family under each one of the parametrisations is available from the first author.

4.2 Likelihood-ratio tests for logisticness and symmetry

During the model reduction stage of data analysis involving a putative SAS-logistic distribution, exploration of the fits of two specific subfamilies will be of particular interest: namely the logistic ($\varepsilon = 0, \delta = 1$) and symmetric ($\varepsilon = 0$) distributions. As both are nested within the wider SAS-logistic family, likelihood-ratio tests (LRTs) can be used to investigate the superiority of the fit of the full SAS-logistic family over those of the two subfamilies. Under standard regularity conditions, the degrees of freedom for the asymptotic chi-squared distribution of the test statistic $-2(\ell_0 - \ell_1)$, where ℓ_0 denotes the maximum of the log-likelihood under the null hypothesis (of logisticness or symmetry) and ℓ_1 denotes its maximum under the full SAS-logistic family, are 2 and 1, respectively. Extensive simulation shows that the asymptotic chi-squared distributions provide reasonable approximations to the sampling distributions of the test statistic for samples sizes of 50 or more for the LRT for symmetry and of 100 or more for the LRT for logisticness. Results for the true size of the latter, as a function of sample size, n, are displayed for a nominal significance level of 5 % in Fig. 4(a).

For smaller sample sizes, the true levels of both tests based on asymptotic chi-squared theory tend to be higher than their nominal levels, i.e. both tests are liberal.

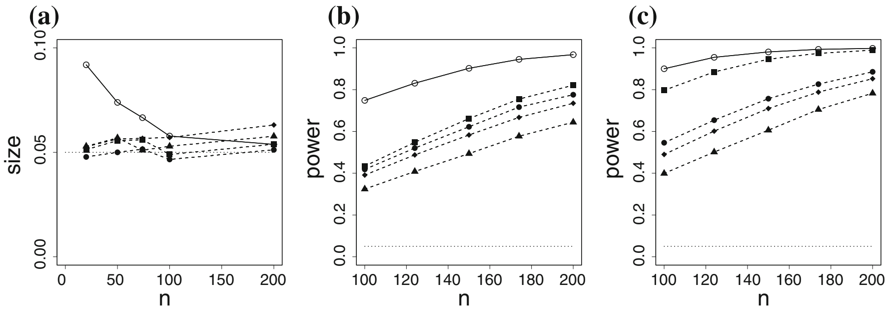

Fig. 4 Proportion of samples for which the null hypothesis of an underlying logistic distribution was rejected in nominally 5 % tests as a function of n calculated using 10000 random samples of size n from the (**a**) logistic, (**b**) Tiku and (**c**) extreme-value distributions. The solid lines connect the results for the likelihood-ratio test, the dashed lines for the Anderson-Darling (*square*), Cramér-von Mises (*solid circle*), Kolmogorov-Smirnov (*triangle*) and Kuiper (*diamond*) tests. The dotted line in each panel delimits the nominal level of 0.05

Calibration of the tests when n is small can be achieved using the bootstrap. In the case of the LRT for symmetry, Efron's (1979) device for symmetrising a sample around an estimate of its centre, such as the mean or median, can be used first to obtain an augmented sample of size $2n$. Bootstrap samples of size n are then randomly drawn from the augmented sample. For the LRT for logisticness, bootstrap samples are drawn from the logistic distribution with parameter values set equal to their maximum likelihood estimates for the original sample.

There are numerous alternative tests against which the LRTs for logisticness and symmetry might be compared. Meintanis (2004) summarises much of the literature on testing for logisticness; Cabilio and Masaro (1996) do likewise for tests of symmetry. Simulation-based results reported in Meintanis (2004) indicate that, of the various tests for logisticness, the Anderson-Darling test studied by Stephens (1979) is generally competitive. Due to their ease of application and similarity of form, in our simulation experiments we compared the performance of the LRT for logisticness with those of the Cramér-von Mises, Watson, Anderson-Darling, Kolmogorov-Smirnov and Kuiper empirical distribution function based tests considered in Stephens (1979). In our Monte Carlo studies, the size and power results obtained for the Cramér-von Mises and Watson tests were found to be identical and so in what follows we refer to their common results as being those of the Cramér-von Mises test. We compared the performance of the LRT for symmetry with the tests of Boos (1982) and Cabilio and Masaro (1996), recommended in the latter paper as omnibus tests of symmetry. The Boos test is calibrated to the logistic distribution, whereas the normal distribution was used when calibrating the test of Cabilio and Masaro (1996).

Unsurprisingly, we found that the two LRTs generally outperform the alternative tests for samples drawn from the SAS-logistic distribution. However, the application of the LRTs need not be restricted to scenarios involving just SAS-logistic distributions. In order to investigate their performance as more general omnibus tests against distributions outside the SAS-logistic family, we conducted a further extensive Monte Carlo experiment. As representative symmetric models we used, ranging from the light- to

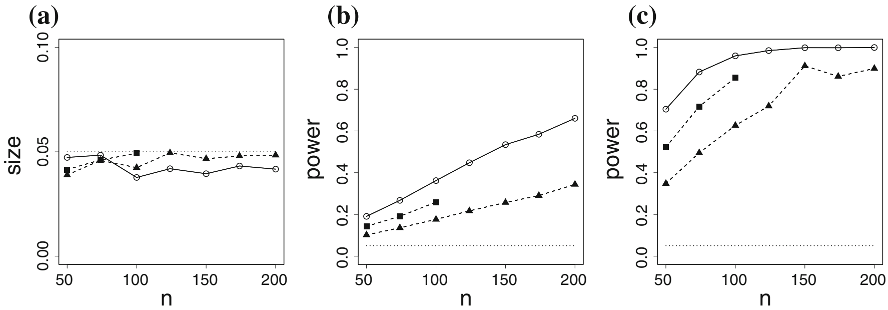

Fig. 5 Proportion of samples for which the null hypothesis of an underlying symmetric distribution was rejected in nominally 5 % tests as a function of n calculated using 10000 random samples of size n from the (**a**) normal $\equiv SN(0)$, (**b**) $SN(2)$ and (**c**) $SN(5)$ distributions. The solid lines connect the results for the likelihood-ratio test, the dashed lines for the Boos (*square*) and Cabilio-Masaro (*triangle*) tests. The dotted line in each panel delimits the nominal level of 0.05. Results for the Boos test are unavailable for $n > 100$ as the storage required to perform it proved prohibitive

the heavy-tailed, the Tiku (Tiku et al. 2001), normal and t_2 distributions. Mudholkar and George (1978) have shown that the logistic distribution is very similar to the t_9 distribution and thus sits between the last two. The asymmetric models used were the extreme-value and a range of Azzalini skew-normal, log F (Jones 2004) and Jones and Faddy (2003) skew-t distributions with varying degrees of asymmetry. The latter two families are order-statistic distributions generated by the logistic and t_2 distributions, respectively.

The LRT was found to be by far the most powerful of the five tests for logisticness against all of the symmetric and asymmetric alternatives considered. In order to illustrate its superior performance, panels (b) and (c) of Fig. 4 portray the powers of the five tests for samples of different size drawn from the Tiku and extreme-value distributions and a nominal significance level of 5 %.

Of the three tests of symmetry considered, the LRT maintains the nominal level best against Tiku's distribution. As Fig. 5(a) portrays, it is slightly conservative against the normal distribution. The Cabilio-Masaro test maintains the nominal level best against the t_2 distribution; the LRT and Boos tests being liberal. The LRT was the most powerful of the tests against all of the asymmetric distributions considered. The results obtained for samples simulated from Azzalini's skew-normal distribution, $SN(\lambda)$, with values of the skewness parameter, λ, of 2 and 5 are presented in panels (b) and (c) of Fig. 5, respectively. The $SN(2)$ distribution is only slightly skewed, whereas the $SN(5)$ displays considerable skewness. Despite the LRT being slightly conservative against the normal, or $SN(0)$, distribution, of the three tests considered it does by far the best job of picking up on the lack of symmetry.

The SAS-logistic-based LRT for symmetry provides an alternative to its SAS-normal-based counterpart proposed by Jones and Pewsey (2009). The former tends to maintain the nominal level better than the latter against the light-tailed Tiku and heavy-tailed t_2 distributions. Of the two, the SAS-normal-based LRT tends to be the slightly more powerful against the various asymmetric distributions considered.

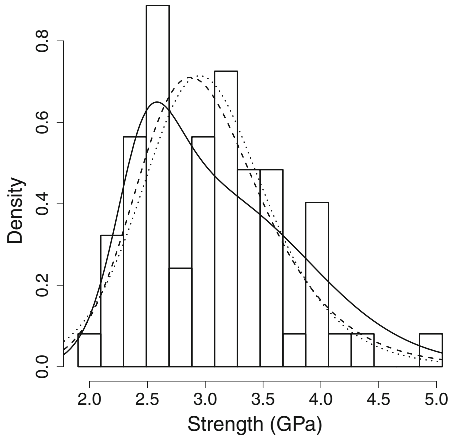

Fig. 6 Histogram of 63 carbon fibre strength measurements with the densities of the maximum likelihood fits for the four-parameter SAS-logistic (*solid*) and three-parameter Azzalini-type skew-logistic (*dotted*) distributions superimposed. The dashed curve is the density of the fit obtained for the three-parameter Type I generalised logistic distribution using L-moment estimation

Given these findings, we recommend both LRTs as omnibus tests for logisticness and symmetry. For small-sized samples, we would expect their bootstrapped versions to provide highly competitive alternatives to existing tests. We illustrate the use of the bootstrap version of the LRT for logisticness in the next section where a sample with a size lower than 100 is analysed.

5 Illustrative example

As our illustrative example, we consider a data set first presented in Badar and Priest (1982) consisting of the strengths of 63 single carbon fibres of length 10mm measured in gigapascals (GPa). Gupta and Kundu (2010) reproduce the data and provide the maximum likelihood fit for the Nadarajah (2009) three-parameter Azzalini-type skew-logistic distribution, as well as what they consider to be the best fit of the three-parameter Type I generalised logistic (or power logistic) distribution, obtained using L-moment estimation. The densities for those two fits are superimposed upon a histogram of the data in Fig. 6, together with the density of the maximum likelihood fit for the four-parameter SAS-logistic family. The latter appears to model the data in the left-hand tail and much of the centre best, whilst the other two densities seem to do a better job in the right-hand tail.

Results for the fits of the SAS-logistic family and its logistic, symmetric and skew-logistic subfamilies are provided in Table 1. With values of $\hat{\delta} = 1.35$ and $\varepsilon = 0$, the best fitting symmetric density is bimodal. The p-values of likelihood-ratio tests for underlying logistic, symmetric and skew-logistic distributions, calculated using asymptotic chi-squared theory, are 0.02, 0.02 and 0.04, respectively. The p-value for the bootstrap version of the likelihood-ratio test for logisticness, computed using 9999 bootstrap samples of size 63 simulated from the logistic distribution with $\xi = 3.02$ and $\eta = 0.35$, was slightly higher at 0.03. Thus, according to these tests, the SAS-logistic family offers a significant improvement in fit over all three subfamilies. The AIC and BIC values, and the p-values for the bootstrap versions of the Anderson-Darling

Table 1 Maximum likelihood estimates, the values of the maximised log-likelihood (MLL), Akaike information criterion (AIC), Bayesian information criterion (BIC) and p-value for bootstrapped versions of the Anderson-Darling goodness-of-fit test (p-value) for the fits to the carbon fibre strengths of the four-parameter SAS-logistic family and its skew-logistic ($\delta = 1$), symmetric ($\varepsilon = 0$) and logistic ($\varepsilon = 0, \delta = 1$) subfamilies

Model	$\hat{\xi}$	$\hat{\eta}$	$\hat{\varepsilon}$	$\hat{\delta}$	MLL	AIC	BIC	p-value
SAS-logistic	2.62	0.60	0.81	1.57	−55.47	118.95	127.52	0.53
Skew-logistic	2.73	0.32	0.48	(1)	−57.56	121.12	127.55	0.03
Symmetric	3.05	0.56	(0)	1.35	−58.30	122.60	129.03	0.17
Logistic	3.02	0.35	(0)	(1)	−59.33	122.66	126.95	0.10

goodness-of-fit test, support the superior fit of the full family. The bootstrap tests were performed by comparing the value of the Anderson-Darling test statistic for the original data and a given fitted model with its values obtained using 9999 parametric bootstrap samples of the same size simulated from the same fitted model. They suggest, in fact, that all but the skew-logistic subfamily provide adequate fits to the data. Any other goodness-of-fit test statistic could have been used with the bootstrap, but we chose the Anderson-Darling statistic because of the competitive power of the Anderson-Darling test for logisticness reported in Section 4.2.

Nominally 95 % confidence intervals for ξ, η, ε and δ, calculated using the inverse of the observed information matrix together with asymptotic normal theory, are $(2.20, 3.04)$, $(0.17, 1.04)$, $(-0.08, 1.70)$ and $(0.74, 2.41)$. The intervals for ε and δ suggest an underlying positively skewed distribution with lighter than logistic tails. The estimated asymptotic correlations between $\hat{\xi}$ and $\hat{\varepsilon}$, and between $\hat{\eta}$ and $\hat{\delta}$, are -0.96 and 0.94, respectively. Under the $(M, IQR, \varepsilon, \delta)$ reparametrisation, the maximum likelihood estimates of the median and interquartile range are $\hat{M} = 2.94$ and $I\hat{Q}R = 0.88$, with nominally 95 % confidence intervals of $(2.78, 3.11)$ and $(0.69, 1.07)$, respectively. For this reparametrisation, the asymptotic correlations between the parameter estimates are all lower than 0.67 in absolute value.

Finally, we also fitted the four-parameter SAS-normal family of Jones and Pewsey (2009) to the data, obtaining a marginally lower maximised log-likelihood value of -55.79. Its density is very similar to that for the SAS-logistic fit in the left-hand tail and lies between the latter and the density of the Type I generalised logistic in the centre and right-hand tail.

6 Multivariate SAS-logistic distributions

With regard to multivariate extensions, consider first a random vector $\mathbf{Z} = (Z_1, Z_2, \ldots, Z_d)$ following the classical d-variate Gumbel–Malik–Abraham logistic distribution (Gumbel 1961; Malik and Abraham 1973) with distribution function

$$F(\mathbf{z}) = 1/\{1 + \sum_{i=1}^{d} \exp(-z_i)\}$$

and density

$$f(\mathbf{z}) = \frac{d!\exp(-\sum_{i=1}^{d} z_i)}{\{1+\sum_{i=1}^{d}\exp(-z_i)\}^{d+1}}.$$

The individual Z_i's are univariate logistic and $\mathrm{corr}(Z_i, Z_j) = 1/2$. Applying the SAS transformation to each of the components of $\mathbf{Z}$ ($Z_i = S_{\varepsilon_i,\delta_i}(X_i)$), its natural SAS-logistic analogue has density

$$f(\mathbf{x}) = \frac{d!\exp\{-\sum_{i=1}^{d} S_{\varepsilon_i,\delta_i}(x_i)\}}{[1+\sum_{i=1}^{d}\exp\{-S_{\varepsilon_i,\delta_i}(x_i)\}]^{d+1}} \prod_{i=1}^{d} \frac{\delta_i C_{\varepsilon_i,\delta_i}(x_i)}{(1+x_i^2)^{1/2}}. \tag{16}$$

By construction, the marginal distribution of X_i is univariate SAS-logistic with parameters $(\varepsilon_i, \delta_i)$. In general, $\mathrm{cov}(X_i, X_j)$ is intractable. However, when $\delta_i = \delta_j = 1$,

$$\begin{aligned}\mathrm{cov}(X_i, X_j) &= \cosh\varepsilon_i \cosh\varepsilon_j E[Z_i Z_j] + \sinh\varepsilon_i \sinh\varepsilon_j \,\mathrm{cov}((1+Z_i^2)^{1/2}, \\ &\quad (1+Z_j^2)^{1/2}) + (\cosh\varepsilon_i \sinh\varepsilon_j + \sinh\varepsilon_i \cosh\varepsilon_j) E[Z_i(1+Z_j^2)^{1/2}] \\ &= \frac{\pi^2}{6}\cosh\varepsilon_i \cosh\varepsilon_j + 0.2881 \sinh\varepsilon_i \sinh\varepsilon_j \\ &\quad -0.4090(\cosh\varepsilon_i \sinh\varepsilon_j + \sinh\varepsilon_i \cosh\varepsilon_j),\end{aligned}$$

with both numerical values quoted to four decimal places.

In applications, the fixed correlation structure between the Z_i's in the Gumbel–Malik–Abraham multivariate logistic distribution will be an important potential limitation to its use, as well as to the use of its SAS-logistic analogue. As an alternative approach, copula theory (Nelsen 2006) provides a means of generating multivariate SAS-logistic distributions with wider ranging dependence structures. Let C denote a copula and $F_{X_i}(x_i) : i = 1, ..., d$ the distribution functions of d univariate SAS-logistic random variables. Then, the joint distribution function of the corresponding multivariate SAS-logistic distribution is obtained as $F(\mathbf{x}) = C(F_{X_1}(x_1), \ldots, F_{X_d}(x_d))$.

Continuing the logistic theme of the paper, appealing copulas in the context of extreme-value theory are the logistic, or Gumbel-Hougaard, copula and its extensions discussed in Gudendorf and Segers (2010). The distribution function obtained using the logistic copula $C(u_1, ..., u_d) = \exp[-\{(-\log u_1)^{1/\theta} + ... + (-\log u_d)^{1/\theta}\}^{\theta}]$ with SAS-logistic marginals is

$$F(\mathbf{x}) = \exp\left\{-\left(\sum_{i=1}^{d}\left[\log\{1+e_{\varepsilon_i,\delta_i}(x_i)\}\right]^{1/\theta}\right)^{\theta}\right\},$$

where $\theta \in (0, 1]$ controls the degree of dependence; $\theta = 1$ for independence and $\theta = 0$ for complete dependence, and $e_{\varepsilon_i,\delta_i}(x_i) = \exp\{-S_{\varepsilon_i,\delta_i}(x_i)\}$. The joint density does not, in general, have a simple form, but reduces to

$$f(\mathbf{x}) = \prod_{i=1}^{d} f_{X_i}(x_i) = \prod_{i=1}^{d} \frac{\delta_i C_{\varepsilon_i,\delta_i}(x_i)}{(1+x_i^2)^{1/2}} \frac{e_{\varepsilon_i,\delta_i}(x_i)}{\{1+e_{\varepsilon_i,\delta_i}(x_i)\}^2}$$

under independence (i.e. when $\theta = 1$). For the bivariate case, the joint density can be represented as

$$f(x_1, x_2) = F(x_1, x_2) \left\{ \mathcal{L}^{2(\theta-1)} + \frac{1-\theta}{\theta} \mathcal{L}^{\theta-2} \right\} \times \prod_{i=1}^{2} \frac{\delta_i C_{\varepsilon_i,\delta_i}(x_i)}{(1+x_i^2)^{1/2}} \frac{e_{\varepsilon_i,\delta_i}(x_i) l_{\varepsilon_i,\delta_i}^{(1-\theta)/\theta}(x_i)}{\{1+e_{\varepsilon_i,\delta_i}(x_i)\}}, \quad (17)$$

where $\mathcal{L} = \sum_{i=1}^{2} l_{\varepsilon_i,\delta_i}^{1/\theta}(x_i)$ and $l_{\varepsilon_i,\delta_i}(x_i) = \log\{1 + e_{\varepsilon_i,\delta_i}(x_i)\}$.

Illustrative contour plots of (17) are provided in Fig. 7. The two panels in the first row have SAS-logistic components that are independent, whilst those in the second row have components that are strongly dependent ($\theta = 0.2$). The effects of the skewness, tailweight and dependence parameters are reflected in the shapes assumed by the contours. Those shapes suggest (17) could well provide a plausible model for bivariate data.

For (17), $\mathrm{cov}(X_i, X_j)$ is, in general, intractable. However, as an alternative, non-moment based, measure of dependence, the local dependence function $\gamma(x_1, x_2) = \partial^2 \log f(x_1, x_2)/\partial x_1 \partial x_2$, proposed as a localised correlation coefficient by Jones (1996), is readily available. It is given by

$$\gamma(x_1, x_2) = (1-\theta) \left(\frac{\mathcal{L}^{\theta-2}}{\theta} + \frac{2\mathcal{L}^{2\theta} + (\theta - 3 + 4/\theta)\mathcal{L}^{\theta} + (1-\theta)(2-\theta)/\theta^2}{\{\theta\mathcal{L}^{\theta+1} + (1-\theta)\mathcal{L}\}^2} \right) \times \prod_{i=1}^{2} \frac{\delta_i C_{\varepsilon_i,\delta_i}(x_i)}{(1+x_i^2)^{1/2}} \frac{e_{\varepsilon_i,\delta_i}(x_i) l_{\varepsilon_i,\delta_i}^{(1-\theta)/\theta}(x_i)}{\{1+e_{\varepsilon_i,\delta_i}(x_i)\}}, \quad (18)$$

which reduces to 0 under independence (when $\theta = 1$).

7 Concluding remarks

In this paper, we have introduced the four-parameter SAS-logistic family of distributions; a flexible extension of the logistic distribution with parameters controlling skewness and tailweight. Its density, distribution function and quantile-based summaries inherit the tractability of the logistic distribution. Its moments, on the other hand, do not have closed-form expressions and must generally be computed using numerical integration. Likelihood-based inference is straightforward for data from all but light-tailed cases of the family. For the latter scenario, we have identified reparametrisations which considerably reduce the magnitudes of the asymptotic correlations

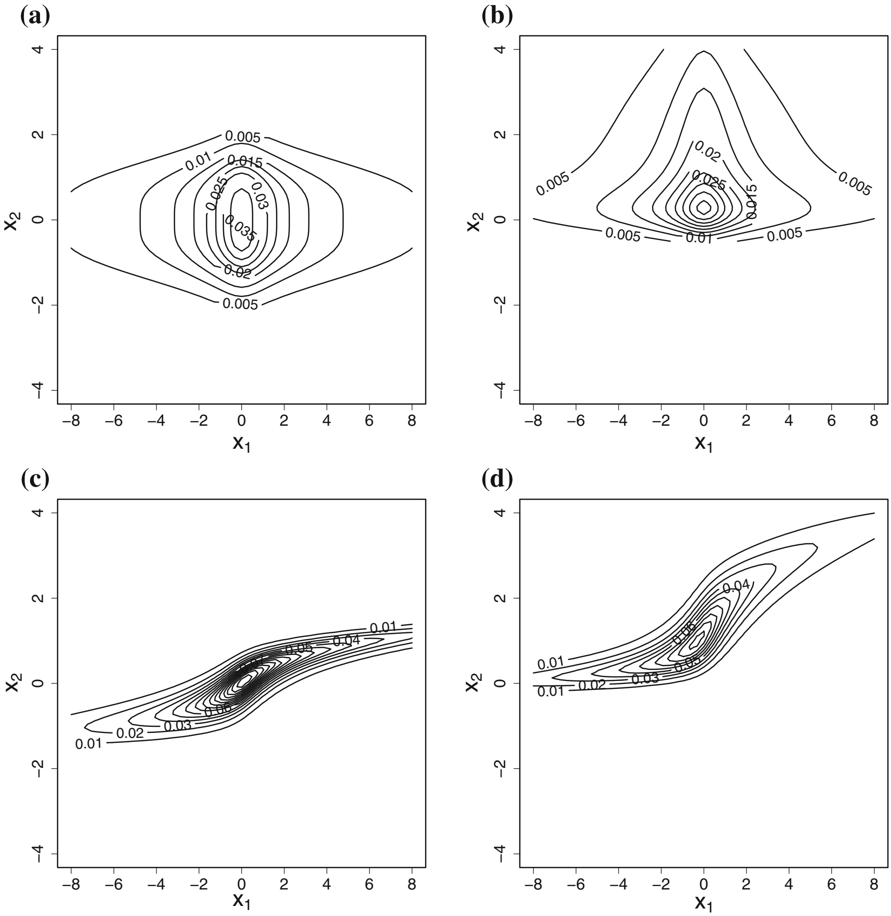

Fig. 7 Contour plots of the bivariate density (17). In the panels in the first row, $\theta = 1$ (corresponding to independence); in the second, $\theta = 0.2$. In the panels in the first column, $\varepsilon_2 = 0$; in the second, $\varepsilon_2 = 1.5$. Throughout, $\delta_1 = 0.5$, $\delta_2 = \sqrt{2}$ and $\varepsilon_1 = 0$

between the maximum likelihood estimators. Here, we have applied the SAS-logistic as a model for univariate data. Other obvious applications involve its use as a model for the error terms in regression problems.

We also proposed two natural multivariate extensions of the family; both with SAS-logistic marginals. The second of the two was derived using the logistic copula. Other multivariate distributions with SAS-logistic marginals can easily be obtained by applying the extensions of the logistic copula discussed in Gudendorf and Segers (2010) or, indeed, any other appealing copula. When discussing the dependence structure of the logistic copula-based bivariate extension, we presented its local dependence function. Inference for the local dependence function is beyond the scope of the present paper. The interested reader is referred to Jones (1996).

8 Appendix: Likelihood results

8.1 C1. Score equations

$$\frac{\partial \ell}{\partial \xi} = \frac{1}{\eta}\left\{\sum_{i=1}^{n}\frac{Y_i}{1+Y_i^2} + \delta \sum_{i=1}^{n}\frac{1}{(1+Y_i^2)^{1/2}}A(Y_i)\right\} = 0,$$

$$\frac{\partial \ell}{\partial \eta} = \frac{1}{\eta}\left\{\sum_{i=1}^{n}\frac{Y_i^2}{1+Y_i^2} + \delta \sum_{i=1}^{n}\frac{Y_i}{(1+Y_i^2)^{1/2}}A(Y_i) - n\right\} = 0,$$

$$\frac{\partial \ell}{\partial \varepsilon} = \sum_{i=1}^{n}A(Y_i) = 0, \quad \frac{\partial \ell}{\partial \delta} = \frac{n}{\delta} - \sum_{i=1}^{n}A(Y_i)\sinh^{-1}(Y_i) = 0,$$

where $A(Y_i) = C_{\varepsilon,\delta}(Y_i)\tanh(\frac{1}{2}S_{\varepsilon,\delta}(Y_i)) - T_{\varepsilon,\delta}(Y_i)$ and $T_{\varepsilon,\delta}(Y_i) = \tanh(\delta\sinh^{-1}(Y_i) - \varepsilon)$.

8.2 C2. Elements of the observed information matrix

The elements of the observed information matrix at the maximum likelihood solution are the following:

$$-\frac{\partial^2 \ell}{\partial \xi^2} = \frac{1}{\eta^2}\left\{\sum_{i=1}^{n}\frac{1-Y_i^2}{(1+Y_i^2)^2} - \delta\sum_{i=1}^{n}\frac{Y_i}{(1+Y_i^2)^{3/2}}A(Y_i) + \delta^2\sum_{i=1}^{n}\frac{1}{1+Y_i^2}B(Y_i)\right\},$$

$$-\frac{\partial^2 \ell}{\partial \xi \partial \eta} = \frac{1}{\eta^2}\left\{\sum_{i=1}^{n}\frac{Y_i(1-Y_i^2)}{(1+Y_i^2)^2} - \delta\sum_{i=1}^{n}\frac{Y_i^2}{(1+Y_i^2)^{3/2}}A(Y_i) + \delta^2\sum_{i=1}^{n}\frac{Y_i}{1+Y_i^2}B(Y_i)\right\},$$

$$-\frac{\partial^2 \ell}{\partial \xi \partial \varepsilon} = \frac{\delta}{\eta}\sum_{i=1}^{n}\frac{1}{(1+Y_i^2)^{1/2}}B(Y_i),$$

$$-\frac{\partial^2 \ell}{\partial \xi \partial \delta} = -\frac{1}{\eta}\sum_{i=1}^{n}\frac{1}{(1+Y_i^2)^{1/2}}\{A(Y_i) + \delta B(Y_i)\sinh^{-1}(Y_i)\},$$

$$-\frac{\partial^2 \ell}{\partial \eta^2} = \frac{1}{\eta^2}\left\{\sum_{i=1}^{n}\frac{2Y_i^2}{(1+Y_i^2)^2} + \delta\sum_{i=1}^{n}\frac{Y_i}{(1+Y_i^2)^{3/2}}A(Y_i) + \delta^2\sum_{i=1}^{n}\frac{Y_i^2}{1+Y_i^2}B(Y_i)\right\},$$

$$-\frac{\partial^2 \ell}{\partial \eta \partial \varepsilon} = \frac{\delta}{\eta}\sum_{i=1}^{n}\frac{Y_i}{(1+Y_i^2)^{1/2}}B(Y_i),$$

$$-\frac{\partial^2 \ell}{\partial \eta \partial \delta} = -\frac{1}{\eta}\sum_{i=1}^{n}\frac{Y_i}{(1+Y_i^2)^{1/2}}\{A(Y_i) + \delta B(Y_i)\sinh^{-1}(Y_i)\},$$

$$-\frac{\partial^2 \ell}{\partial \varepsilon^2} = \sum_{i=1}^{n} B(Y_i), \qquad -\frac{\partial^2 \ell}{\partial \varepsilon \partial \delta} = -\sum_{i=1}^{n} B(Y_i) \sinh^{-1}(Y_i),$$

$$-\frac{\partial^2 \ell}{\partial \delta^2} = \frac{n}{\delta^2} + \sum_{i=1}^{n} B(Y_i)\{\sinh^{-1}(Y_i)\}^2,$$

where $B(Y_i) = S_{\varepsilon,\delta}(Y_i)\tanh(S_{\varepsilon,\delta}(Y_i)) + \frac{1}{2}C^2_{\varepsilon,\delta}(Y_i)\text{sech}(\frac{1}{2}S_{\varepsilon,\delta}(Y_i)) - 1/C^2_{\varepsilon,\delta}(Y_i)$.

Acknowledgments We are most grateful to an Associate Editor and two referees for their careful reading of a previous draft of the paper and suggestions towards improving its content and presentation. Financial support for the research which led to the production of this paper was received by the first author in the form of grant MTM2010-16845 from the Spanish Ministry of Science and Education and grant GR10064 from the Junta de Extremadura.

References

Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, *12*, 171–178.

Badar, M. G., Priest, A. M. (1982). Statistical aspects of fiber and bundle strength in hybrid composites. In T. Hayashi, K. Kawata, S. Umekawa (Eds.), *Progress in science and engineering composites* (pp. 1129–1136). Tokyo: ICCM-IV.

Balakrishnan, N. (1992). *Handbook of the logistic distribution*. New York: Dekker.

Boos, D. D. (1982). A test for asymmetry associated with the Hodges–Lehmann estimator. *Journal of the American Statistical Association*, *77*, 647–651.

Byrd, R. H., Lu, P., Nocedal, J., Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, *16*, 1190–1208.

Cabilio, P., Masaro, J. (1996). A simple test of symmetry about an unknown median. *Canadian Journal of Statistics*, *24*, 349–361.

Efron, B. (1979). Bootstrap methods: another look at the jackknife. *Annals of Statistics*, *7*, 1–26.

Gudendorf, G., Segers, J. (2010). Extreme-value copulas. In P. Jaworski, F. Durante, K. Jaworska, T. Rychlik (Eds.), *Copula theory and its applications* (pp. 127–145). Berlin: Springer.

Gumbel, E. J. (1961). Bivariate logistic distributions. *Journal of the American Statistical Association*, *55*, 335–349.

Gupta, R. D., Kundu, D. (2010). Generalized logistic distributions. *Journal of Applied Statistical Sciences*, *18*, 51–66.

Johnson, N. L., Tadikamalla, P. R. (1992). Translated families of distributions. In N. Balakrishnan (Ed.), *Handbook of the logistic distribution* (pp. 189–208). New York: Dekker.

Johnson, N. L., Kotz, S., Balakrishnan, N. (1994). *Continuous univariate distributions*. (2nd ed., vol. 2). New York: Wiley.

Jones, M. C. (1996). The local dependence function. *Biometrika*, *83*, 899–904.

Jones, M. C. (2004). Families of distributions arising from distributions of order statistics. *Test*, *13*, 1–43.

Jones, M. C., Faddy, M. J. (2003). A skew extension of the *t* distribution, with applications. *Journal of the Royal Statistical Society, Series B*, *65*, 159–174.

Jones, M. C., Pewsey, A. (2008). Sinh-arcsinh distributions: a broad family giving rise to powerful tests of normality and symmetry. Technical Report 08/06, The Open University, Department of Statistics. http://statistics.open.ac.uk/TechnicalReports/TechnicalReportsIntro.htm.

Jones, M. C., Pewsey, A. (2009). Sinh-arcsinh distributions. *Biometrika*, *96*, 761–780.

Jones, M. C., Rosco, J. F., Pewsey, A. (2011). Skewness-invariant measures of kurtosis. *American Statistician*, *65*, 89–95.

Malik, H. J., Abraham, B. (1973). Multivariate logistic distributions. *Annals of Statistics*, *1*, 588–590.

Meintanis, S. G. (2004). Goodness-of-fit tests for the logistic distribution based on empirical transforms. *Sankhyā*, *66*, 306–326.

Moors, J. J. A. (1988). A quantile alternative for kurtosis. *Journal of the Royal Statistical Society, Series D*, *37*, 25–32.
Mudholkar, G. S., George, E. O. (1978). A remark on the shape of the logistic distribution. *Biometrika*, *65*, 667–668.
Nadarajah, S. (2009). The skew logistic distribution. *Advances in Statistical Analysis*, *93*, 187–203.
Nelder, J. A., Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, *7*, 308–313.
Nelsen, R. B. (2006). *An introduction to copulas* (2nd ed.). New York: Springer.
Rosco, J. F., Jones, M. C., Pewsey, A. (2011). Skew t distributions via the sinh-arcsinh transformation. *Test*, *20*, 630–652.
Rotnitzky, A., Cox, D. R., Bottai, M., Robins, J. (2000). Likelihood-based inference with singular information matrix. *Bernoulli*, *6*, 243–284.
Stephens, M. A. (1979). Tests of fit for the logistic distribution based on the empirical distribution function. *Biometrika*, *66*, 591–595.
Tadikamalla, P. R., Johnson, N. L. (1982). Systems of frequency curves generated by transformations of logistic variables. *Biometrika*, *69*, 461–465.
Tiku, M. L., Islam, M. Q., Selcuk, A. S. (2001). Nonnormal regression. II. Symmetric distributions. *Communications in Statistics—Theory and Methods*, *30*, 1021–1045.
van Zwet, W. R. (1964). *Convex transformations of random variables*. Amsterdam: Mathematisch Centrum.