

<<http://www.amstat.org/publications/jse>>

## The Weight of Euro Coins: Its Distribution Might Not Be As Normal As You Would Expect

Ziv Shkedy, Marc Aerts, and Herman Callaert  
Hasselt University

Journal of Statistics Education Volume 14, Number 2 (2006),  
[www.amstat.org/publications/jse/v14n2/datasets.aerts.html](http://www.amstat.org/publications/jse/v14n2/datasets.aerts.html)

Copyright © 2006 by Ziv Shkedy, Marc Aerts, and Herman Callaert, all rights reserved. This text may be freely shared among individuals, but it may not be republished in any medium without express written consent from the authors and advance notification of the editor.

---

\*Key Words: \*Normal mixture; Normal probability plot; Outlier; Skew-Normal distribution; Truncation.

### Abstract

Classical regression models, ANOVA models and linear mixed models are just three examples (out of many) in which the normal distribution of the response is an essential assumption of the model. In this paper we use a dataset of 2000 euro coins containing information (up to the milligram) about the weight of each coin, to illustrate that the normality assumption might be incorrect. As the physical coin production process is subject to a multitude of (very small) variability sources, it seems reasonable to expect that the empirical distribution of the weight of euro coins does agree with the normal distribution. Goodness of fit tests however show that this is not the case. Moreover, some outliers complicate the analysis. As alternative approaches, mixtures of normal distributions and skew normal distributions are fitted to the data and reveal that the distribution of the weight of euro coins is not as normal as expected.

### 1. Introduction

The last couple of decades have seen a reform in the K-12 mathematics curriculum in the United States. In the same period, the undergraduate curriculum in collegiate mathematics has undergone a major review.

The recommendations towards a significant pedagogical reform include emphasis on so-called ♦data-driven♦ mathematics education, and by extension, on data-driven statistics education. In the K-12 mathematics curriculum, statistics (referred to as ♦Data Analysis and Probability♦ in the NCTM 2000 Principles and Standards for School Mathematics) has a much more prominent role than ever before. The recent introduction of the AP (Advanced Placement) Examination in Statistics focused attention on the importance of statistics even more.

In ♦The Case for Undergraduate Statistics♦ Richard L. Scheaffer and Carl Lee write: ♦Many now realize that statistical thinking and statistical concepts are essential skills for all academic disciplines as well as for life-long learning. As a result of this thinking, spurred on by the influx of freshmen who have already seen basic statistics by the time they arrive at the college door, statistics is receiving more attention at two-year colleges, four-year colleges, and universities." Also here,

the increased attention given to statistics goes hand in hand with an increased emphasis on "statistical thinking" and "data and concepts", as recommended by the ASA-MAA Joint Curriculum Committee.

The Guidelines for Assessment and Instruction in Statistics Education (GAISE) were endorsed by the ASA on May 17, 2005. The GAISE reports contain excellent recommendations for statistics education, as well in PreK-12 years as at the introductory college level.

This change in the teaching of statistics is not confined to the U.S. alone, but it takes place in many countries all over the world, and at different pace and intensity. In Flanders for example (the Flemish part of Belgium), statistics played only a minor role in mathematics education at school. A recent curriculum reform (compulsory for all schools) imposes a significant increase of statistics in grades eight through twelve.

At the university level, the recent introduction of a European-wide bachelor-master system has forced almost all institutions to restructure their curricula. Flemish universities have turned this burden into an opportunity for reshaping the statistics education, starting with an introductory course where data and concepts prevail over formal mathematical proofs.

In the light of these reforms, several real life projects are being developed, providing examples of what can be done in class. Not so long ago, the introduction of the Euro has had a major impact in the daily life of everyone (many people still have to count back into the old currency to make sure that they rightly value the Euro price!). It therefore was considered a good idea to look a bit closer at the new coin itself, and to consider some of its physical properties. After all, everybody in Belgium (and in the many countries of the European Union) has to use such coins now. Here, the choice was made to study the one Euro coin. This coin has one side that is the same everywhere (with the number 1 on it), but the other side has a representation that is country dependent. Therefore, the choice was made to restrict the study to the Belgian 1 Euro coin.

The paper is organized as follows. The data are introduced in Section 2 [<v14n2/datasets.aerts.html#Sect2>](#) and a first analysis, using the Kolmogorov-Smirnov and Shapiro-Wilk goodness of fit tests, is presented in Section 3 [<v14n2/datasets.aerts.html#Sect3>](#). A second analysis, based on normal mixture models, is discussed in Section 4 [<v14n2/datasets.aerts.html#Sect4>](#). As it will contribute to a better understanding, the details of the data collection procedure are given in Section 5 [<v14n2/datasets.aerts.html#Sect5>](#). In a third and final analysis, in Section 6 [<v14n2/datasets.aerts.html#Sect6>](#), we do not assume that the distribution of the euro coins is symmetric. The family of skew normal distributions, a class of distributions which includes the normal ones (Azzalini 1985 [<v14n2/datasets.aerts.html#Azzalini1985>](#)), offers an interesting alternative to model the weight of euro coins. A possible classroom use of the data is discussed in Section 7 [<v14n2/datasets.aerts.html#Sect7>](#).

## 2. The Dataset

According to information from the National Bank of Belgium the 1 Euro coin weighs 7.5 grams. It was anticipated that the weight of this coin would be normally distributed with mean 7.5 g. Students could try this out in class, and this type of data collection is fortunately not very complicated. Most schools possess precise weighing scales in their physics or chemistry labs, and asking each student to bring a couple of

1 Euro coins from home looks trivial. A variety of statistical techniques could be demonstrated by the teacher, going from histograms to  $z$ -scores, from sample averages to confidence intervals, and so on. Since we didn't want to take our assumptions for granted, we started out to construct a kind of reference set, to be made available to all participating schools. A local bank was kind enough to let us borrow 2000 coins. Then, two of our assistants, Sofie Bogaerts and Saskia Litière, weighted those coins one by one. They had reserved a small chemistry lab where a precise digital weighing scale of the type Sartorius BP 310s was available. This scale gives a reading of up to a thousandth of a gram. With the help of the technicians from the lab, the scale was put in place and well calibrated. They then started the following standard procedure. Saskia made sure that the scale gave a reading of 0.000 g. She then put one coin on the scale and waited until the reading was stable. She then told the resulting number to Sofie, who wrote it down in a spreadsheet of her laptop. The two assistants were the only persons in the room, the doors were closed and the temperature was kept constant. The total procedure took 7 hours and 22 minutes. At the end, we had a database consisting of 2000 weights (in milligrams) of Belgian 1 Euro coins.

### 3. A First Analysis

Figure 1 shows a normal probability plot for the weight of the 2000 euro coins and reveals three outliers (all marked by arrows). Let  $Y_i$  be the weight of a euro coin. Based on the sample  $Y_1, \dots, Y_{2000}$ , we wish to test the hypotheses

$$\begin{aligned} H_0 &: Y \text{ is normally distributed,} \\ H_1 &: Y \text{ is not normally distributed.} \end{aligned} \quad (1)$$

The two-sided Kolmogorov-Smirnov (KS, see e.g. Section 6.5.2 in Shao 1999) test statistic equals 0.0234 resulting in a  $p$ -value of 0.0131, indicating that the null hypothesis can be rejected at a significance level of 5%. The Shapiro-Wilk (SW) statistic is equal to 0.975 ( $p$ -value  $< 0.0001$ ). This latter test is designed specifically for testing normality and is therefore more powerful than the KS test. The SW statistic is also attractive because it has a simple, graphical interpretation as an approximate measure of the correlation in the normal probability plot in Figure 1 (see e.g. Stuart, Ord, and Kendall 1991).

Figure 1

Figure 1: Normal probability plot for the weight of 2000 euro coins.

It is interesting to examine the influence of the three outliers 7.201 g, 7.656 g and 7.752 g on the analysis. Figure 2 shows the corresponding normal probability plot after excluding these outliers. Now, the KS test statistic equals 0.0246 ( $p$ -value=0.0071) and the SW statistic is 0.9974 ( $p$ -value=0.0022), indicating, once again, that the null hypothesis of normally distributed weights can be rejected. This shows that non-normality of the weight variable  $Y$  is not simply a matter of excluding outliers but that there are more fundamental features of the data that are not in line with a normal distribution.

-----  
Figure 2 <v14n2/datasets.aerts\_figure2.gif>

Figure 2: Normal probability plot for the weight of 1997 Euro coins, excluding the three outliers 7.201 /g/, 7.656 /g/ and 7.752 /g/ from the original dataset.

-----

As the observed deviation from normality can be induced by heterogeneity in the mean and variance parameter of the normal distribution, an interesting option is to extend a single normal distribution to a hierarchical normal mixture model. The plausibility of this first extension is discussed in the next section.

#### 4. A Second Analysis: Normal Mixtures

The analyses performed in the previous section assume that all Euro coins were sampled from the same normal population. However, if the packages were samples from different normal subpopulations the distribution of the 2000 Euro coins is not expected to be normal but rather a mixture of normal distributions (Gelman, Carlin, Stern, and Rubin 1995 <v14n2/datasets.aerts.html#Gelman1995>; Gilks, Richardson, and Spiegelhalter 1996; <v14n2/datasets.aerts.html#Gilks1996> Congdon 2003 <v14n2/datasets.aerts.html#Congdon2003>). Let  $f(y)$  be the density function of  $y$ . A finite mixture distribution has the form of

where  $\pi_j$  are called the mixture components,  $\mu_j$  are the mixture probabilities and  $\sigma_j^2$  are the parameters to be estimated. In what follows, we focus on a mixture of two normal populations with possibly different mean and variance parameters and, following the approach of Congdon (2003) <v14n2/datasets.aerts.html#Congdon2003> we formulate the mixture model in terms of an hierarchical model using a latent indicator variable.

##### 4.1 A Two Components Normal Mixture With Known Mixture Probability

Let  $G_i$  be an indicator variable that takes the value of 1 if the individual belongs to the first group and 0 otherwise. Hence, the likelihood in the first stage of the hierarchical model is given by

$$(2)$$

The variable  $G_i$  can be seen as a latent classification random variable for which we assume a Bernoulli distribution

$$(3)$$

In this stage we temporarily assume that the mixing probability is known. To complete the specification of the hierarchical model, the following (non informative) prior and hyperprior distributions are assumed

$$(4)$$

Here,  $\text{gamma}(\cdot)$  denotes a gamma distribution with mean equal to  $\mu$  and variance  $\mu^2$ . Hence, by choosing  $\mu = 0.0001$ , the prior (and hyperprior) distribution for the variance parameters is a gamma distribution with mean 1 and variance 10000, which reflects our uncertainty about the true value of the parameters (see, for example Gilks, et al. 1996 <v14n2/datasets.aerts.html#Gilks1996>). Three different values for  $\mu$  were examined,  $\mu = 0.95$ ,  $0.99$ , or  $0.995$ . Table 1

[v14n2/datasets.aerts.html#Table1](#) presents the posterior means for the mean and variance parameters of the two normal mixture components. The upper panels of Figure 3 [v14n2/datasets.aerts.html#Figure3](#) show the posterior means of  $G_i$  while the lower panels show the posterior medians for the three values of  $\theta$ . Note that since  $G_i$  is a binary variable, the median and the mode of the posterior distribution are equal. Based on the two sets of plots the coins can be classified into the two components of the mixture. For  $\theta = 0.995$  (right panels), 4 coins are classified into the second component of the mixture. This group of coins does contain the three outliers. So the two component normal mixture is able to cope with the three outliers. Note that as the value of  $\theta$  decreases the number of coins which are classified into the second component of the mixture increases (5 and 11 for equal to 0.99 and 0.95, respectively).

---

Table 1: Posterior mean for  $\mu_j$ ,  $\sigma_j^2$ ,  $j=1, 2$ , for a two component normal mixture with known mixture probability  $\theta$ .  $N_2$  is the number of observations assigned to group 2.

	$\mu_1$	$\sigma_1^2$	$\mu_2$	$\sigma_2^2$	$N_2$
0.995	7.521	0.001087	7.504	0.06362	4
0.99	7.521	0.001072	7.509	0.01987	5
0.95	7.521	0.001016	7.52	0.004859	11

---

Figure 3 [v14n2/datasets.aerts\\_figure3.gif](#)

Figure 3: Upper panels: posterior mean for  $G_i$  with  $\theta = 0.95$  (left), 0.99 (middle) and 0.995 (right). Lower panels: posterior medians for  $G_i$  with  $\theta = 0.95$  (left), 0.99 (middle) and 0.995 (right).

---

#### 4.2 A Mixture of Two Components With Unknown Mixture Probability

In the previous section the mixture probability was assumed to be known. In this section we relax this assumption and specify a prior distribution for  $\theta$ . Hence, we modify the model in Equation (3) [v14n2/datasets.aerts.html#Eqn3](#) to

$$(5)$$

The likelihood in Equation (2) [v14n2/datasets.aerts.html#Eqn2](#) and the hyperprior distributions in Equation (4) [v14n2/datasets.aerts.html#Eqn4](#) remain the same.

A density estimate for the posterior distribution of  $\theta$  is shown in Figure 4 [v14n2/datasets.aerts.html#Figure4](#) and Table 2 [v14n2/datasets.aerts.html#Table2](#) presents the posterior means for the mean and the variance parameter of both mixture components. The posterior mean of  $\theta$  equals 0.9964. Figure 5 [v14n2/datasets.aerts.html#Figure5](#) shows that, according to the posterior mean and mode of  $G_i$ , in this case only two coins are classified into the second component. These two coins correspond to the two extreme outliers (7.201  $\mu/g$  and 7.752  $\mu/g$ ).

---

Figure 4 [v14n2/datasets.aerts\\_figure4.gif](#)

Figure 4: Density estimate for the posterior distribution of .

Table 2: Posterior mean for and ,  $i = 1, 2$ .  $N_{i2}$  is the number of observations assigned to group 2.

Parameter	Estimate
	7.521
	7.431
	0.001093
	3.928
posterior mean	0.9964
$N_{i2}$	2

Figure 5 <v14n2/datasets.aerts\_figure5.gif>

Figure 5: Left (right) panel: posterior mean (median) for  $G_i$  / with = 0.9964.

The two component normal mixture approach is able to attribute the outlying values to a separate normal component. But since our first analysis showed that the normal distribution hypothesis could not be retained after excluding the outliers, this approach is still not giving us a satisfactory insight in which way the weight distribution deviates from a normal one. Of course many other distributions could be fit to these data. But to better understand the reason for this non-normality, we investigated the nature of the data and its collection in more detail.

## 5. The Data Collection

The rather surprising result of our first data analysis made us suspicious about the data collection process. The design of the study was set up so that the 2000 coins would make up a simple random sample (SRS) from the collection of all Belgian 1 euro coins. Since the euro has been introduced for some time now, it was anticipated that the coins would be well mixed by now. Hence, it was expected that collecting coins from any bank would result in a sample, which, for all practical purposes, could be considered as an SRS.

But what happened in reality? From our side, we had asked the cooperation of the university administration (section finances) since they have good contacts with local banks. They contacted a bank and sent an employee to go and collect the coins. But when the employee asked for  $\text{€}$ Belgian only  $\text{€}$  coins, he found out that banks do not store Belgian euros separated from other euros. Fortunately, he was told that the  $\text{€}$ National Bank of Belgium  $\text{€}$  possesses  $\text{€}$ Belgian only  $\text{€}$  coins. So, he went to that place and came back with the 2000 coins.

We further asked our two assistants about the way in which they  $\text{€}$ physically  $\text{€}$  received those coins. Here came the big surprise. The  $\text{€}$ National Bank of Belgium  $\text{€}$ , which is the distributor of money to other banks, had given 8 packages of 250 brand new coins each. Our assistants explained further that they had worked systematically, package by package, and that the weights were entered into the spreadsheet in that order. This gives additional information, in the sense that the weight of each coin can also be classified according to the package to which it

belongs.

Knowing that the 2000 coins came from 8 different packages made us perform a further analysis, as described in the next paragraph. Of course, this analysis now should be seen in the context of the peculiar way in which the data were collected, and any resulting conclusion should be handled with caution.

## 6. A Third Analysis: Testing Normality by Package

In this section we redo the first analysis, by package, both including and excluding the three outliers (one outlier from package 3 (7.201 /g/), one from package 5 (7.656 /g/) and one from package 8 (7.752 /g/)). Figure 6 <v14n2/datasets.aerts.html#Figure6> shows the normal probability plots for each package (after excluding the outliers) and Table 3 <v14n2/datasets.aerts.html#Table3> presents the values of the KS and the SW test statistics together with sample means and variances. An overall comparison of the two goodness of fit tests confirms the conservative nature of the KS test and the better power characteristics of the SW test. When including the outliers, the null hypothesis of a normal distribution is clearly rejected for package 3, 6 and 8. For the fifth package, and with or without the outlying euro coin with weight 7.656 /g/, the null hypothesis cannot be rejected. When the three outliers were excluded from the analysis (last three rows in Table 3 <v14n2/datasets.aerts.html#Table3>), the null hypothesis of a normal distribution cannot be rejected for packages 3 and 8, at the significance level of 5% (borderline for package 3). So, after excluding the outliers, the packages with the smallest /p/-values for the SW test are: package 6 (0.0068), package 3 (0.0565) and package 8 (0.0833); all other packages result in a /p/-value about 10% (package 7 is next with /p/-value 0.1120). Recall that when the coins from the 8 packages were pooled together, excluding the 3 outliers, the null hypothesis was rejected.

-----  
 Figure 6 <v14n2/datasets.aerts\_figure6.gif>

Figure 6: Normal probability plots by package (outliers excluded). Upper plots: packages 1 to 4, from left to right; lower plots: packages 5 to 8, from left to right.

-----  
 Table 3: KS and SW two-sided tests for normality, by package.  
 First 8 rows: all data, last three rows: excluding the outliers.  
 Here, SD = standard deviation, /N/ = size of package.

Package	KS (/p/-value)	SW (/p/-value)	Mean	SD	/N/
1	0.038 (0.500)	0.995 (0.6830)	7.520	0.034	250
2	0.033 (0.500)	0.991 (0.1219)	7.523	0.035	250
3	0.078 (0.001)	0.863 (<0.0001)	7.510	0.037	250
4	0.045 (0.500)	0.995 (0.6827)	7.531	0.029	250
5	0.035 (0.500)	0.991 (0.1290)	7.531	0.030	250
6	0.055 (0.063)	0.984 (0.0068)	7.515	0.033	250
7	0.042 (0.500)	0.990 (0.1120)	7.523	0.033	250
8	0.070 (0.005)	0.936 (<0.0001)	7.517	0.036	250
3	0.047 (0.500)	0.989 (0.0565)	7.511	0.031	249
5	0.045 (0.500)	0.993 (0.4204)	7.531	0.029	249
8	0.053 (0.081)	0.989 (0.0833)	7.516	0.033	249

Some of the normal probability plots in Figure 6 [Figure6](#) reveal a clear pattern of departure from normality: package 3, 6 and 8 (U-shaped about the straight line, indicating some skewness to the right), and to some lesser extent, package 7 (U-shaped below the straight line, indicating some skewness to the left). This was the motivation to consider a non-symmetric extension of the normal distribution: the skew-normal distribution as introduced by Azzalini (1985) [Azzalini1985](#).

## 7. A Fourth Analysis: Skew-Normal Distributions

The univariate-skew normal distribution, proposed by Azzalini (1985) [Azzalini1985](#), is an extension of the normal distribution with three parameters: the mean the variance and the shape parameter  $\alpha$ . The normal distribution is a special case, taking the shape parameter equal to zero. Formally, the density function of the skew normal distribution  $SN(\mu, \sigma, \alpha)$  is given by

$$(6)$$

where  $\phi(z)$  is a  $z$ -score, is the standard normal density function,  $\Phi(z)$  is the standard normal distribution function and  $\alpha$  is the shape parameter. Note that for  $\alpha = 0$ , the skew-normal density reduces to the standard normal density. Figure 7 [Figure7](#) shows some graphs of the skew-normal density functions for several values of  $\alpha$ . Note that as  $\alpha$  goes to infinity the skew-normal distribution converges to the so called half-normal distribution.

Figure 7 [aerts\\_figure7.gif](#)

Figure 7: Skew-normal densities for different values of the shape parameter  $\alpha$ .

For a random variable  $Z \sim SN(\mu, \sigma, \alpha)$ , it can be shown that  $E(Z) = \mu + \sigma \alpha \phi(\alpha/\sigma) / \Phi(\alpha/\sigma)$  and  $\text{Var}(Z) = \sigma^2 [1 - (\alpha/\sigma)^2 \phi(\alpha/\sigma)^2 / \Phi(\alpha/\sigma)^2]$ , where  $\phi$  and  $\Phi$  are the standard normal density and distribution functions, respectively. Using a linear transformation  $X = \mu + \sigma Z$ , we get a random variable with mean  $\mu$ , variance  $\sigma^2$  and shape parameter  $\alpha/\sigma$ . For more details on moments and cumulants, see Azzalini (1985) [Azzalini1985](#).

The shape parameter was estimated for the data in each package, as shown in Table 4 [Table4](#). In the sequel, we work with the reduced data, excluding the three outliers. Using maximum likelihood estimation, the shape parameter in package 6 is estimated to be 0.512, with estimated standard error 0.143. This indicates a right skewed distribution. Within this three-parameter family of distributions the normality hypothesis is equivalent to the hypothesis

$$H_0 : \alpha = 0. \quad (7)$$

Using the Wald test

$$(8)$$

the normality hypothesis (7) [Eqn7](#) can be tested using a chi-square null distribution (see e.g. Section 6.4.2 in Shao 1999 [Shao1999](#)). For package 6 the value



of  $W$  equals 12.85 which corresponds to a  $p$ -value of about 0.0003 leading us to reject normality. So, this confirms the previous analysis of this particular package. Moreover the fitted skew-normal distribution can now be used to model the package 6 data.

Table 4: Estimates of the shape parameter and corresponding estimated standard errors for the skew-normal distribution, by package.

Package		
1	0.161	0.176
2	-0.142	0.136
3	0.225	0.145
4	0.000	0.025
5	-0.072	0.175
6	0.512	0.143
7	-0.189	0.164
8	0.390	0.147

Also for package 8 the estimated lambda parameter and its estimated standard error leads to a rejection of the normality hypothesis ( $\lambda = 0.39, \text{SE} = 0.147$ ), and the value of  $W$  equals 7.04 with  $p$ -value 0.008. These are the only packages for which the normality hypothesis ( $H_0$ ) can be rejected at the 5% significance level (which is in line with the results in the previous analysis). Also for other packages, like package 3 and 7 and even package 2, the magnitude of relative to its agrees with the  $p$ -values in Table 3.

Histograms of the weight of the euro coins by package, overlaid with the estimated skew normal density (solid line), are shown in Figure 8. The 95% confidence intervals for  $\lambda$ , by package, are shown in Figure 9. As expected, the confidence intervals for package 6 and 8 do not cover the value of zero.

Figure 8

Figure 8: Histogram, fitted skew normal density (solid line) and fitted normal density (dashed line), by package. Upper row: packages 1 to 4, from left to right; lower row: packages 5 to 8, from left to right.

Figure 9

Figure 9: Estimated shape parameter with 95% confidence intervals, by package.

## 8. Classroom Use And Final Discussion

Initial exploration of the dataset might start with methods from exploratory data analysis (EDA). A simple normal probability plot like in Figure 1 reveals the existence of three probable outliers. A deviation in the tails might hint at

skewness. When the 3 outliers are excluded, a new normal probability plot in Figure 2 shows the same behavior in the tails. One then could go one step further and perform a formal test. Shapiro-Wilk confirms that the null hypothesis about the normality of the weights can be rejected. A more advanced course in statistics then could concentrate on finding a model that better fits the data. A first idea could come from the class of normal mixtures, leading to the discussion in Section 4 [<v14n2/datasets.aerts.html#Sect4>](#).

A second aspect to be discussed in class is certainly as important as the exploration of the dataset. It is about the design of the study and about the many practical issues that occur  $\diamond$ in real life $\diamond$ . A big lesson can be learned here. Even a well-designed study can produce data, which are collected in such a way that one has to be cautious about the conclusions of the further statistical analysis. Apart from the design, one has to look at the size and the complications of the data collection process. If one has to rely on collaborators for collecting the data, one should monitor also this process carefully.

Within the restrictions of what has been said above, one could continue and go for a further  $\diamond$ tentative $\diamond$  analysis. Fortunately, we were able to identify which coin belonged to which batch, so that a study by batch was possible. This is carried out in Section 6 [<v14n2/datasets.aerts.html#Sect6>](#), using normal probability plots as well as formal tests. Several of the plots seem to point in the direction of some skewness. In a more advanced course in statistics, one could take the opportunity here to study an interesting class of distributions, having a shape parameter  $\lambda$ , but reducing to the symmetrical normal when  $\lambda = 0$ . This class of skew-normal distributions, introduced by Azzalini (1985) [<v14n2/datasets.aerts.html#Azzalini1985>](#), is studied in Section 7 [<v14n2/datasets.aerts.html#Sect7>](#) for fitting the data of the 8 packages of coins.

## 9. Getting the Data

The file `euroweight.dat` [<datasets/euroweight.dat>](#) is a text file containing the raw data. The file `euroweight.txt` [<datasets/euroweight.txt>](#) is a documentation file containing a brief description of the dataset.

---

## Acknowledgement

The authors are grateful to the Datasets and Stories Editor and two anonymous referees for their helpful comments and suggestions that improved this paper.

---

## Appendix: Software

### Normal Probability Plots and KS and SW Test

We used the `trellis` figures library in S-PLUS to produce the normal probability plots discussed above. Figure 1 [<v14n2/datasets.aerts.html#Figure1>](#) was produced using the following code (note that `y` is an S-PLUS object containing the weights of the 2000 euro coins).

```

>y<-(euro$MASSA..G)
>upy<- max(y)
>loy<- min(y)
qqmath( ~ y , data = euro,ylim=c(loy,upy),
  prepanel = prepanel.qqmathline,
  panel = function(x, y) {
    panel.grid()
    panel.qqmathline(y, distribution = qnorm)
    panel.qqmath(x, y)},
  layout = c(1, 1), aspect = 1,
  xlab = "Unit Normal Quantile", ylab = "Weights (gr)")

```

For the KS two sided test we used

```
>ks.gof(y,alternative = "two.sided", distribution = "normal")
```

and for the SW test:

```
>shapiro.test(y)
```

### Normal Mixture with Two Components

The mixture models discussed in Section 4 were fitted in WinBUGS 1.3. The posterior means for the parameters in each model are based on a Markov Chain Monte Carlo (Gilks, et al. 1996 <[v14n2/datasets.aerts.html#Gilks1996](#)>) simulation with 10000 iterations. The first 1000 were discarded from the analysis (treated as the burn-in period). In order to avoid autocorrelation, the chains were monitored every 10 iterations. The following code was used for the model with the unknown mixture probabilities  $N_{sub}$  is the number of euro coins and  $N_{mix}$  is the number of components in the mixture). For a more elaborate discussion on hierarchical mixture models, we refer to Congdon (2003) <[v14n2/datasets.aerts.html#Congdon2003](#)>.

```

model
{
  for( i in 1 : Nsub ) {
    Yi[i]~dnorm(mu.y[T[i]],tau.i[T[i]]) #likelihood
    T[i]~dcat(P[]) #prior for G_i
    ti[i]<-T[i]-1
  }
  for(k in 1:Nmix){
    tau.i[k] ~ dgamma(0.0001, 0.0001)
    mu.y[k] ~ dnorm(0,tau.mu)
    sigma[k]<-1/tau.i[k]
  }
  tau.mu ~ dgamma(0.0001, 0.0001)
  P[1]~dunif(0,1) #prior for pi
  P[2]<-1-P[1]
}

```

### Skew-Normal Distribution

The skew-normal distributions were fitted using the S-PLUS library `sn`, which is available on Professor Azzalini's website [azzalini.stat.unipd.it/SN/index.html#lib-sn](http://azzalini.stat.unipd.it/SN/index.html#lib-sn) <[v14n2/datasets.aerts\\_link1.html](#)>. Maximum likelihood estimates for the mean, variance and shape parameter were obtained using the function `sn.mle()`. A call to this function has the form

>sn.mle(x1,plotit=T)

---

## References

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

Congdon, P. (2003), /Applied Bayesian Modelling/, Chichester, England: Wiley.

GAISE Reports (2005). Electronic text at [it.stlawu.edu/~rlock/gaise/v14n2/datasets.aerts\\_link2.html](http://it.stlawu.edu/~rlock/gaise/v14n2/datasets.aerts_link2.html)

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995), /Bayesian Data Analysis/, London: Chapman and Hall.

Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (1996), /Markov Chain Monte Carlo in Practice/, London: Chapman and Hall.

Scheaffer, R. L. and Lee, C. (2000), /The Case for Undergraduate Statistics/ Electronic text at [www.amstat.org/meetings/jsm/2000/usei/case.html](http://www.amstat.org/meetings/jsm/2000/usei/case.html)  
<v14n2/datasets.aerts\_link3.html>.

Shao, J. (1999), /Mathematical Statistics/, New York: Springer.

Stuart, A., Ord, J. K., and Kendall, M. (1991), /Kendall's Advanced Theory of Statistics, Volume 2/, London: Arnold.

---

Ziv Shkedy  
Center for Statistics  
Hasselt University  
Diepenbeek  
Belgium  
[/ziv.shkedy@uhasselt.be/](mailto:ziv.shkedy@uhasselt.be) <<mailto:ziv.shkedy@uhasselt.be>>

Marc Aerts  
Center for Statistics  
Hasselt University  
Diepenbeek  
Belgium  
[/marc.aerts@uhasselt.be/](mailto:marc.aerts@uhasselt.be) <<mailto:marc.aerts@uhasselt.be>>

Herman Callaert  
Center for Statistics  
Hasselt University  
Diepenbeek  
Belgium  
[/herman.callaert@uhasselt.be/](mailto:herman.callaert@uhasselt.be) <<mailto:herman.callaert@uhasselt.be>>

---

Volume 14 (2006) <[contents\\_2006.html](#)> | Archive <[jse\\_archive.html](#)> |  
Index <[jse\\_index.html](#)> | Data Archive <[jse\\_data\\_archive.html](#)> |  
Information Service <[jse\\_info\\_service.html](#)> | Editorial Board  
<[jse\\_board.html](#)> | Guidelines for Authors <[jse\\_author\\_info.html](#)> |  
Guidelines for Data Contributors <[jse\\_data\\_contributor\\_info.html](#)> | Home  
Page <./> | Contact JSE <<mailto:journals@amstat.org>> | ASA Publications

</publications/>