

## Chapter 6

# Comparative graphical analysis

### 6.1 KDEs and Histograms

The view adopted here is that if comparison between just a small number of films is required this is better pursued using graphical methods, with numerical descriptive statistics in a supplementary role. That is, given two films to compare and for example, it may not be that interesting to compare their ASLs or median SLs, even if accompanied by measures of scale and shape. It is likely that anything to be inferred from such comparisons will be evident from graphical analysis, and they may miss a lot if the SL distributions are at all ‘non-standard’.<sup>1</sup>

The view is also taken that the comparison of SL distributions using separate histograms is unwieldy; they are not easily overlaid, for example. Comparison using KDEs is more effective. The skewed nature of SL distributions also makes comparison on the untransformed scale quite awkward, so many of the illustrations below will be based on log-transformed SLs. The code illustrated is easily adapted to untransformed SLs if this is what is really needed. For initial illustration a late Hitchcock silent, *Easy Virtue* (1928), and early ‘talkie’ *The Skin Game* (1931). This is for comparison with Redfern (2009), an early advocate for some of these techniques in the cinematics literature.

Figure 6.1 compares the histograms on the untransformed SL scale. The code used was

```
par(mfrow = c(2,1))
hist(SL.Easy_Virtue, 50, xlim = c(0,175), xlab = "SL", main = "Easy Virtue")
hist(SL.The_Skin_Game, 200, xlim = c(0,175), xlab = "SL", main = "The Skin Game")
par(mfrow = c(1,1))
```

The maximum SL across the two films is 174.7 and `xlim = c(0, 175)` defines the range of the scale to go just beyond this for both histograms. The numbers of bins, 50 and 200 were selected, with a bit of experimentation, to ensure a reasonable appearance with equal bin-widths of 1 second. This is all in the interests of comparability. It is possible to program R to do this automatically.

The most obvious features are that the bulk of SLs are concentrated in much the same region for the two films, and *The Skin Game* has a much longer tail. A lot of finer detail is not that obvious, as later analyses will show.

The same comparison is undertaken using KDEs in Figure 6.2. This is neater, and more readily interpretable since the two KDEs are easily overlaid on the same graph. The code used was

```
plot(density(SL.The_Skin_Game), ylim = c(0, .14), lwd = 2, col = "blue",
xlab = "SL", main = "")
```

---

<sup>1</sup>In dealing with a large body of films, summary statistics are useful for getting an overview of possible patterns in the data. The decline in ASLs since the 1950s is a commonly cited example of this sort of use. Where films have an approximate lognormal distribution it is possible to summarise them well using just two statistics, but the lognormality needs to be demonstrated first.

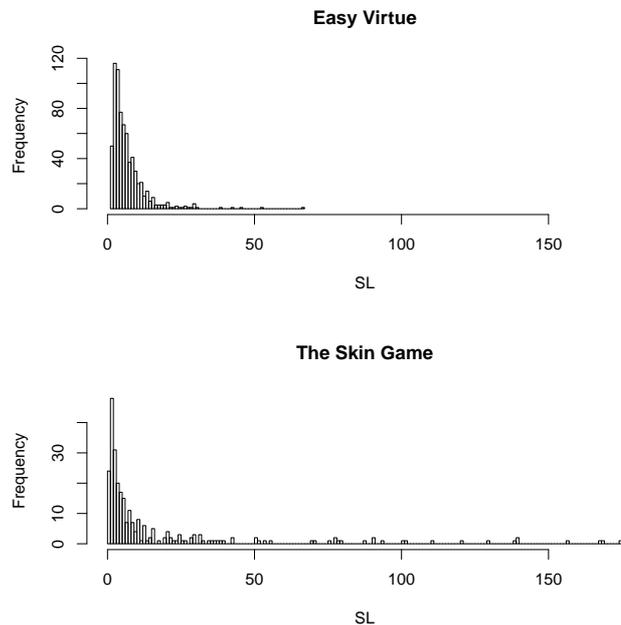


Figure 6.1: A comparison of SL histograms for *Easy Virtue* and *The Skin Game*.

```
lines(density(SL.Easy_Virtue), lty = 2, lwd = 2, col = "red")

legend("topright", legend = c("The Skin Game", "Easy Virtue"), lty = c(1,2),
      lwd = c(2,2), col = c("blue", "red"), bty = "n")
```

The effect of the `lines` command is to overlay whatever is specified on the structure dictated by the `plot` command. Most of the arguments are for presentational purposes and can be omitted in exploratory analysis. The `lwd` argument, for example, controls the line-width, and is chiefly there to ensure better visibility when looking at the graph either on screen or the printed page.

The horizontal scale of the plot is dictated by the film named in the `plot` command; *The Skin Game* has the greater range of SLs and its placement here avoids, in this instance, having to use the `xlim` argument. The `ylim` argument is needed to control the vertical scale and accommodate the KDE for *Easy Virtue*; you only discover this after an initial analysis and need to experiment, or look at the two KDEs separately first, to get suitable limits.

The `legend` command for – no surprise – adding legends is introduced here. They can be a little tiresome to construct; `?legend` in R, or examples in texts such as Murrell (2006), are helpful guides.

As far as interpretation goes the greater density of *Easy Virtue* at the shorter SLs is more apparent than with the histogram. The rather different tail behaviour for SLs greater than about 20 is also more apparent. This becomes even more obvious, along with other feature, once KDEs on a log-scale are compared.

Figure 6.3 is similar to the previous figure except that a logarithmic scale is used. Commands identical to those for that figure were used, except that film names were replaced by their logarithms (e.g., `log(SL.Easy_Virtue)`) and the vertical scale was determined by `ylim = c(0, .55)`. Bandwidths of 0.17 and 0.30 were used for *Easy Virtue* and *The Skin Game* respectively

Visually the plot for *Easy Virtue* doesn't look too unlike a normal distribution. That for *The Skin Game* is clearly non-normal, and the rather lumpy nature of the tail should caution against any simple summary in terms of either the median SL or ASL. Apart from the different distributions the different tail behaviour is evident with SLs for *The Skin Game* extending beyond those for *Easy Virtue* in either direction. This can be seen in other figures but is much more

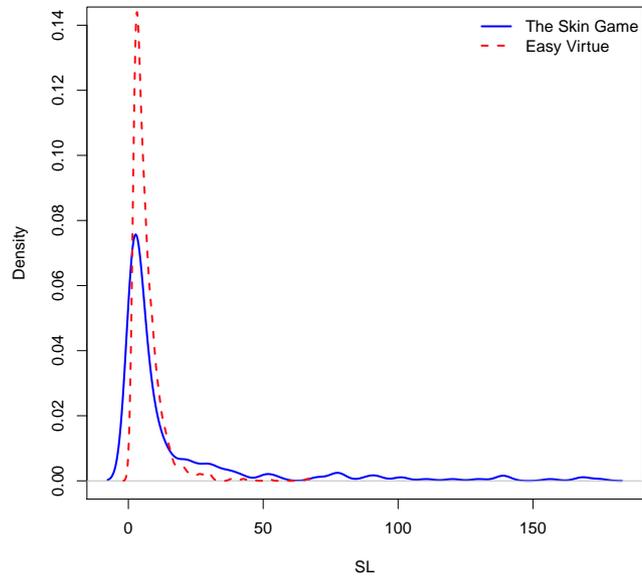


Figure 6.2: A comparison of KDEs for the SLs of Easy Virtue and The Skin Game.

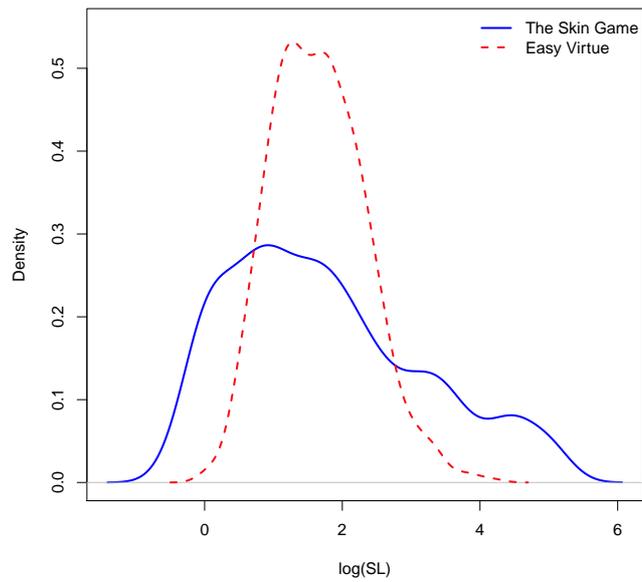


Figure 6.3: A comparison of KDEs for the logged SLs of Easy Virtue and The Skin Game.

prominently displayed for the shorter SLs on the log-scale.

This analysis can be taken further if desired. The KDEs for the two films intersect to the left at about 0.7, corresponding to an SL of 2. For *The Skin Game* 24.6% of the SLs are less than this, compared to 5.5% for *Easy Virtue*. To the right the intersection is at about 2.77, corresponding to an SL of 16; 22.8% of SLs for *The Skin Game* exceed this, compared to 5.2% for *Easy Virtue*.<sup>2</sup>

## 6.2 Boxplots and violin plots

For comparing SLs for two films, as in the previous section, graphical display may be all that is necessary. Graphs can be accompanied with numerical summaries if these are thought helpful, but this may be gilding the lily, and possibly misleading if SLs are at all non-lognormal. Redfern (2012d) has advocated the use of five-number summaries, which are enshrined in the boxplot (Section 5.3).

Comparative boxplots and violin plots, on a log-scale, are illustrated in Figure 6.4. The boxplots show, well-enough, the similar medians for the two films and the greater spread of *The Skin Game*. The rather lumpy nature of this film is not captured, so the plot is not as effective a method of comparison as the KDEs. To correct for this violin plots might be used. The default is not especially useful and in the figure the degree of smoothing, controlled by  $h = .2$ , was obtained after some experimentation. It shows the lumpier distribution of *The Skin Game* better, the appearance of the plot for *Easy Virtue* being little affected by the choice. It is a happy accident, and meaningless, that the plot for *The Skin Game* more resembles a violin.

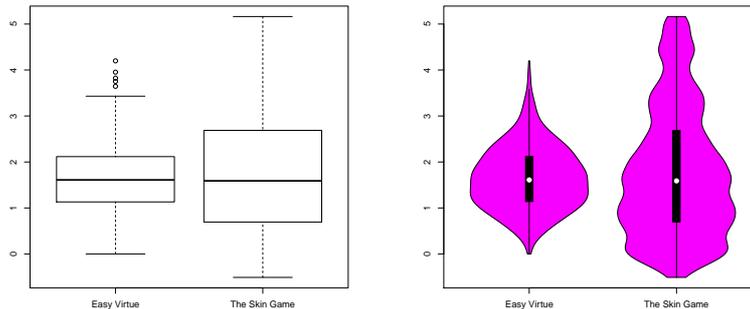


Figure 6.4: *Comparative boxplots and violin plots, on a log scale, for Easy Virtue and The Skin Game.*

The code used for the figures is

```
boxplot(log(SL.Easy_Virtue), log(SL.The_Skin_Game), names = c("Easy Virtue", "The Skin Game"))

library(vioplot)
vioplot(log(SL.Easy_Virtue), log(SL.The_Skin_Game), h = 0.2, names = c("Easy Virtue", "The Skin Game"))
```

In addition to the use of the `h` argument, note the use of `names` to control labelling.

## 6.3 Cumulative frequency diagrams

Empirical *cumulative frequency diagrams* (ECFDs) are another way of comparing distributions. They plot, in the examples used here, the *proportion* of SLs less than or equal to an SL against

<sup>2</sup>In R, with a graph of the KDEs on the screen, vertical lines can be placed on the plot using, for example, `abline(v = 0.7)`. Play around with the value of `v` to locate positions of interest. Replace `v` with `h` if horizontal lines are of interest. Given a logged SL, 0.7 say, it can be converted to an SL using `exp(0.7)`.

SL. If interest centers on comparing SLs for films in terms of the proportion of shots above/below some value, they are a more natural tool to use than histograms or KDEs, though I suspect non-statistical users find them more difficult to read. I have seen little use of them in the cinematics literature; Redfern (2009) and elsewhere is an exception; the plots in Figure 6.5 emulate those in <http://nickredfern.wordpress.com/2012/04/12/using-the-ecdf-to-analyse-film-style/>

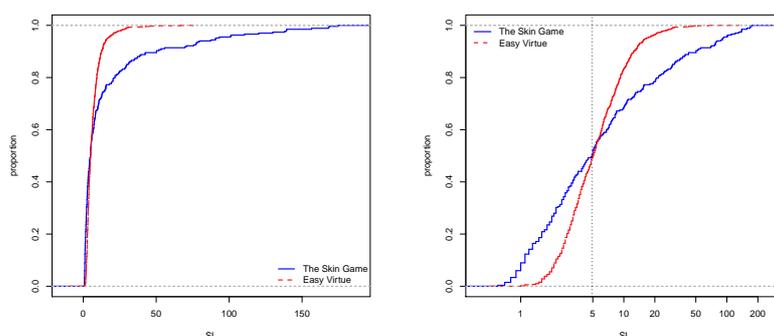


Figure 6.5: A comparison of ECDFs for the SLs of *Easy Virtue* and *The Skin Game* to the left and logged SLs to the right.

Other than emphasizing what we already know, that *The Skin Game* has a longer tail than *Easy Virtue*, the plot for the untransformed SLs hides a lot of detail. A log-transformation is indicated once again, shown to the right of the figure. The median, which happens to be the same for both films is indicated by the dotted vertical line. The effect of the log-transformation is to ‘stretch out’ the lower tail while compressing the upper tail. As with the KDEs the difference in SLs at the lower values is much more obvious.

The plot for the logged data was obtained using

```
lSG <- log(SL.The_Skin_Game)
lEV <- log(SL.Easy_Virtue)

plot(ecdf(lSG), verticals = T, lwd = 2, col = "blue", pch = "", xaxt = "n", xlab = "SL",
     ylab = "proportion", main = "")

lab <- c(1,5,10,20,50,100,200)
axis(1, at = log(lab), labels = lab)
lines(ecdf(lEV), verticals = T, lwd = 2, lty = 2, col = "red", pch = "")
abline(v = median(lSG), lty = 3, lwd = 1.5)

legend("topleft", inset = c(0, .02), legend = c("The Skin Game", "Easy Virtue"), lty = c(1,2),
      lwd = c(2,2), col = c("blue", "red"), bty = "n")
```

where `xaxt = "n"` removes the default x-axis, and the `lab` and `axis` command replace it with an axis where the labelling is on the untransformed rather than SL scale <sup>3</sup>.

<sup>3</sup>This is just to show it can be done if you want to; it’s not needed if labelling with the logged values is acceptable. The plot is on the log-scale; `lab` selects those SLs you want to show as labels on the axis; in `axis` the 1 specifies the x-axis as the axis to draw; `at = log(lab)` generates the tick marks for the logged values at which the labels are to be placed; and `labels = lab` adds the labels specified in `lab`. This can be done in other ways, probably more neatly, but it worked here. In `legend`, placed by the first argument in the top-left corner, `inset = c(0, .02)` shifts it down slightly to avoid overlap with the guideline at the top; `?legend` provides the detail.

## 6.4 Comparison with reference distributions

### 6.4.1 Comparisons with the lognormal distribution

It is sometimes of interest to see how well an SL distribution conforms to a lognormal distribution and, perhaps more interestingly, to see what any departure from lognormality looks like. Lognormality does not imply that this is what the makers of a film consciously aim for. For a variety of reasons SLs are expected to vary (e.g., shorter ASLs for ‘action’ scenes, longer ASLs for ‘romantic’ interludes is mentioned from time to time). Patterns of SLs may vary more or less rhythmically within a film (Chapter 7) but this kind of variation is lost if one looks at the ‘external’ patterning of SLs, as we have been doing so far. It is worth quoting Salt (2011) at length, on this.

The Lognormal distribution results when the quantity under consideration, in our case shot length, is determined as a result of the probabilities associated with a large number of independent causative factors being multiplied together. . . . the amount of conscious decision-making in making films may be less than people think, and many film-makers may often be doing what they do unthinkingly, because that is what everybody does. In films what is presumably concerned in determining the length of a shot is the simultaneous interaction of such factors in the scene being filmed as how the actor moves in the shot with respect to the closeness of the camera, the length of the lines he speaks, and how the other actors react, and so on. The fact that different individuals are usually responsible for these various components of a film, from the scriptwriter to the director to the editor, assists the independence of these causes. However, once in a while a film-maker may decide to do something unorthodox on purpose (this is the aesthetic impulse, after all), and this will upset the regularity of features like the Lognormal distribution.

Elsewhere, in his contribution to the Cinematics debate, and with specific reference to films made in the early sound era, Salt cites technological factors that might be expected to militate against the appearance of lognormality. The extent to which films do have lognormal distribution has been disputed. This issue is not addressed here and not that important for what is to follow in this section. It is, I think, indisputable that a lot of films have qualitatively similar and reasonably regular SL distributions, and the lognormal distribution is a useful approximate base, even if strictly wrong, for judging serious departures from regularity that invite explanation.

Lognormality implies normality of the log-transformed SLs. It is easier to visually assess normality than lognormality and the following exploration of the use of probability plots and KDEs is based on log-transformed data where the normal distribution can be used as a reference distribution. A brief excursion into the normal distribution is provided first (also, see the Appendix).

### 6.4.2 Aspects of the normal distribution

Call the SLs  $X$  and denote the particular SLs for a film with  $n$  shots as  $(x_1, x_2, \dots, x_n)$ . Unless it is important to identify a specific shot it is simplest to ignore the subscript and refer to observed SLs as  $x$ .

In the same vein let  $Y = \log X$  with  $y = \log x$  the observed logged values. Interest for the purposes of this section lies in whether or not the  $y$  look reasonably like a sample from a normal distribution. A normal distribution is completely defined by knowledge of its mean,  $\mu$  and standard deviation (SD),  $\sigma$ . This is conventionally written as

$$Y \sim N(\mu, \sigma^2)$$

read as  $Y$  is normally distributed with mean,  $\mu$  and variance,  $\sigma^2$ , the variance being the square of the SD.

It is often convenient to *standardize* the data so it has a mean of zero and variance (and standard deviation) of 1. This is done by defining

$$z = \frac{y - \mu}{\sigma}.$$

If the data have a normal distribution then  $z \sim N(0, 1)$  (the *standard normal distribution*) but

this is not assumed here as the question of whether or not the transformed data are normal is that at issue.<sup>4</sup>

In the context of SL distributions, it is possible for two films with lognormal distributions to have rather different ASLs and scales. Even though qualitatively similar they will look different on the untransformed scale; after log-transformation they will have the same normal ‘bell-shape’, but if placed on the same graph will have different locations and scales (spread). What the  $z$ -transform does is remove location and scale differences so that, if normal, the plots of the distribution will be identical. Any differences between films in terms of departures from lognormality may be more apparent if compared in this way and, in particular, they can be compared to the standard normal distribution to see where any differences from the normal lie.

In practice, it is necessary both to assess whether observed differences are of potential substantive significance, to see how these are manifest on the untransformed scale and, ultimately, to see if useful explanations can be offered for departures from lognormality (or, more generally, any commonality of distribution exhibited by SL distributions). These, sometimes complex, issues are largely avoided in the discussion to follow of how to effect comparisons.

### 6.4.3 Normal probability plots

Plots of the kind discussed so far are good at identifying clear departures from (log)normality in the middle of an SL distribution, or for identifying obvious outliers. More subtle features of tail behaviour, particularly those associated with the smaller SLs, are harder to discern, and probability plots may be better for this purpose (though they can be difficult to interpret).

One way of thinking about this is that the *ordered* values of a sample are plotted against values they are expected to have if they follow a normal distribution. If they do, an approximate straight line should result. It’s simplest to think in terms of standardized data. The median of the observed data (the 50th *percentile*) should plot against the median, zero, of the standard normal; the lower quantile of the data;  $Q_1$ , or the 25th percentile, should plot against that for the standardized normal, -0.67, and so on. To illustrate *A Night at the Opera* is used. The SLs are given by `z`, `scale(log(z))` standardizes the logged data, and `qqnorm(scale(log(z)))` produces Figure 6.6.

The plot is of the ordered data against the values they should have if they are exactly normal. The dashed line shows exactly where the points should lie in this case. It is obtained by using `abline(a = 0, b = 1)` which generates a line with zero intercept and unit slope. The solid line is obtained from `qqline(scale(log(z)))` and serves a similar function except that it is designed to pass through the upper and lower quartiles. The way the `qqline` is defined avoids the effects of any outliers in the data. (The `lty`, `lwd` and `col` arguments that control the appearance of the lines have been omitted in the above commands.)

As applied here `abline` is only useful if the data are standardized. If not, the plot looks exactly the same but the scale on the vertical axis is different and this line is not appropriate. In this latter case if a line is fitted to the plot the intercept and slope can be used as estimates of the mean and SD.

The curved nature of the plot is a clear indicator of non-normality, related to the kind of skewness evident in Figures 5.1 and 5.5. The obvious outlier previously detected stands apart in the upper-right of the plot. Interpretation of these kinds of plot is not always straightforward; further examples are provided in the next sections.

### 6.4.4 Using KDEs for SL comparisons

To illustrate some of the issues concerning SL comparisons with KDEs, touched on in Section 6.4.1, *Exodus* (1960) and *Aristocats* (1970) will be used. They have been chosen for two reasons. One is that in Redfern’s (2012a) study of lognormality they were among the few films accepted by his

---

<sup>4</sup>Unfortunately the term *normalization* is sometimes used for the transformation; the term *standardization* is preferred here as it avoids the implication that the data are converted to normality. If the distribution is non-normal to begin with it will be non-normal after transformation to  $z$ .

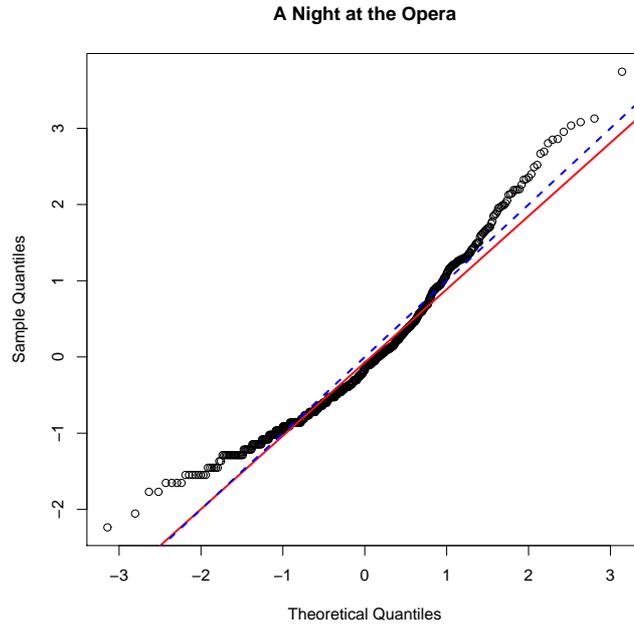


Figure 6.6: A normal probability plot for standardized logged SL data for a A Night at the Opera. See the text for an explanation of the reference lines.

(quite stringent) criteria as having approximately lognormal distributions so this can, perhaps, be taken as given. The other reason is that they have markedly different ASLs (22,4 and 4.0). and medians (13.8 and 3.3), so that comparing their underlying distributions presents challenges.

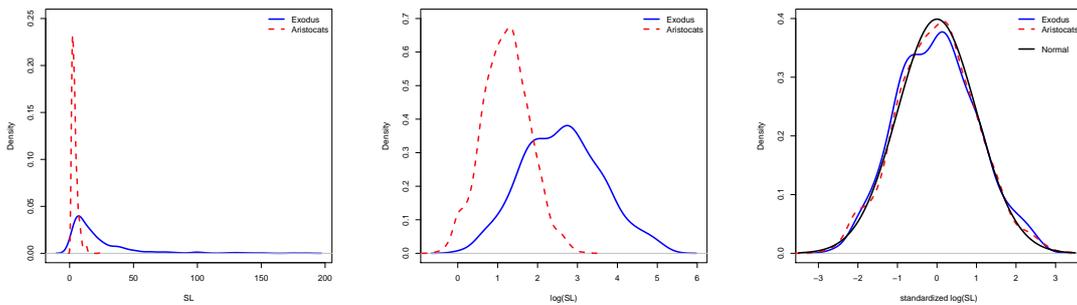


Figure 6.7: Different comparisons of SLs for Exodus and Aristocats using KDEs. See the text for discussion.

This is illustrated in the left-hand plot of Figure 6.7, using untransformed SL data. The differences in location and scale and shape make assessment of the underlying distributional similarity difficult. After log-transformation the middle plot removes this problem to some extent, but location and scale differences remain. By standardizing the log-transformed data, as in the final figure, their distributional similarity to each other and to the reference standardized normal distribution (hence implying lognormality) is apparent.

It should be emphasised that this way of comparing SL distributions focuses on only one aspect, the underlying probability distribution (if any). In other contexts location and scale differences, of the kind evident in the left-hand plot, may be of more direct interest.

Code to obtain the plots is of the basic form

```
plot(density(f1))
lines(density(f2))
```

where `f1` and `f2` are the appropriate transformations for the films. That is, successively, `f1` is `SL.Exodus`, `log(SL.Exodus)`, `scale(log(SL.Exodus))`, with `Aristocats` replacing `Exodus` for `f2`. Use the film with the largest SL for `f1`. Arguments governing appearance of the plots have been omitted; `ylim` may need experimenting with to fit everything on the plot; thus `ylim = c(0, 0.7)` and `ylim = c(0, 0.4)` were needed for the first two plots.

The normal curve is added using

```
x <- seq(-3.5, 3.5, .01) # Plotting positions
y <- dnorm(x)             # Normal densities
lines(x,y, lwd = 3)      # Add to KDE plot
```

where 3.5 in the definition of `x` will normally be adequate, but can be increased if the range of the KDEs exceeds these limits.

## 6.4.5 Examples

### *Walk the Line* (2005)

Cutting ([http://www.cinematics.lv/cutting\\_on\\_salt.php](http://www.cinematics.lv/cutting_on_salt.php)) in a debate with Barry Salt about various aspects of Cutting *et al.* (2010) presents a histogram very similar to that to the left of Figure 6.8, for *Walk the Line* (2005). It is used in support of the contention that many films, over a period of about 70 years and on a log-scale have distributions that ‘look normal’. Of the 134 films assessed for lognormality in Redfern (2012a), a subset of the 150 used in Cutting *et al.* (2010), *Walk the Line* is one of the *least* (log)normal according to several criteria used.

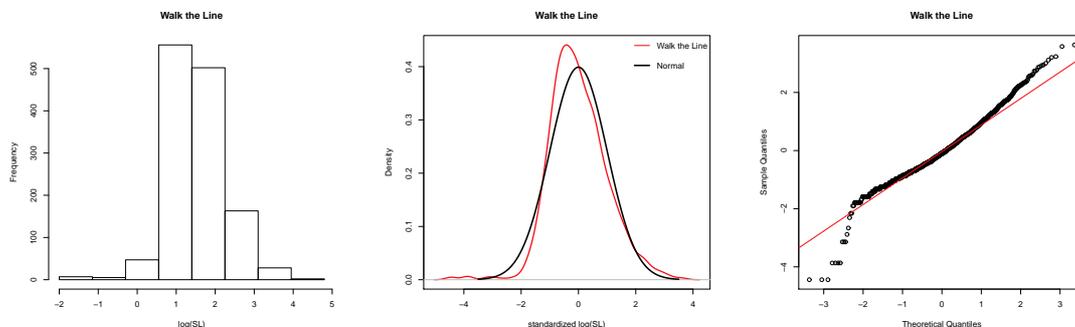


Figure 6.8: A histogram, KDE and normal probability plot for the logged SLs of *Walk the Line*, also standardized for the last two plots.

The KDE indicates why. There are several aspects of it that are clearly non-normal, most noticeably the left tail behaviour, but also a degree of skewness. These are emphasised even more in the probability plot where, apart from the curvature over much of the range, the aberrant left tail behaviour is clearly emphasised. The transformed data are clearly not normal and, contrary to the implication of Cutting’s claim, the film does not have a lognormal SL distribution.

The histogram is a blunt instrument for assessing normality at best. The problem is compounded here by the choice of large bin-widths, producing a rather uninformative histogram where the tail behaviour is nearly invisible. The main point of the example, though, is to illustrate how sensitive the probability plot can be to distributional features of the tails, indicative of non-(log)normality.

Whether or not departures from lognormality are of any importance in whatever one considers to be the ‘stylistic’ aspects of a film is an entirely separate judgment. The probability plot shows that about 15 SLs (of 1310) are noticeably smaller than expected if lognormality applies. They are shots measured as having a duration of 1 second or less, down to 2 deci-seconds.

### Comparing probability plots

Probability plots can be difficult to interpret, and direct comparison of plots for different films, on the same graph, may be useful in getting some feel for what patterns and differences tell you. Figure 6.9 attempts such a comparison for *Walk the Line* and *A Night at the Opera*, mainly to show how this might be done. It is necessary to use standardized logged SLs for this plot.

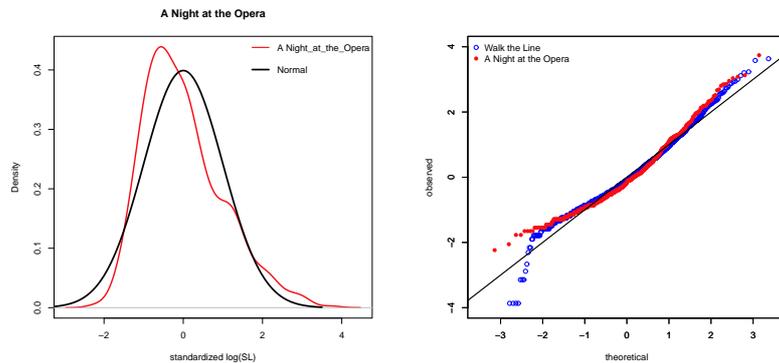


Figure 6.9: To the left a KDE of standardized logged SL data for *A Night at the Opera*; to the right probability plots for this film and *Walk the Line* compared.

The left-hand plot showing the KDE for standardized logged SL data for *A Night at the Opera* may be compared with the comparable plot for *Walk the Line* Figure 6.8. The point to note for present purposes is that the left tail of the KDE for the latter film trails beyond that for the standardized normal distribution; the opposite is the case for the former film. You need to look carefully to see this – it doesn’t ‘hit you in the face’. The difference is starkly evident in the tail behaviour exhibited in the probability plots to the right. These were obtained from the following code.

```
slWTL <- scale(log(SL.Walk_the_Line))
slNATO <- scale(log(SL.Night_at_the_Opera))
x1 <- qqnorm(slWTL)$x
y1 <- qqnorm(slWTL)$y
x2 <- qqnorm(slNATO)$x
y2 <- qqnorm(slNATO)$y

plot(x1, y1, col = "blue", xlim = c(-3.5, 3.5), ylim = c(-4, 4), xlab = "theoretical",
     ylab = "observed", main = "")

par(new = T)

plot(x2, y2, pch = 16, col = "red", xlim = c(-3.5, 3.5), ylim = c(-4, 4), xlab = "", ylab = "",
     main = "")

abline(0,1, lwd = 2)

legend("topleft", legend = c("Walk the Line", "A Night at the Opera"), pch = c(1,16),
     col = c("blue", "red"), bty = "n")
```

The objects `x1` to `y2` hold the coordinates for plotting the two probability plots, executed using the `plot` commands; `par(new = T)` overlays the plots, with the `xlim` and `ylim` arguments

set to ensure the plot limits are the same. The `pch` argument determines the plotting symbol, the default (1) being an open circle. Googling ‘R pch’ is enough to see what the choice is, and the symbols are listed in books such as Venables and Ripley (2002) and Murrell (2006).

### *Comparing more than two KDEs*

One obvious advantage of KDEs over histograms is the ease with which they can be overlaid, for comparative purposes. How many plots can be usefully overlaid depends a bit on what they look like and what you might want to say about them. The next example started out in a ‘for fun’ spirit to see how far things might be pushed but turns out to embody a possibly useful idea.

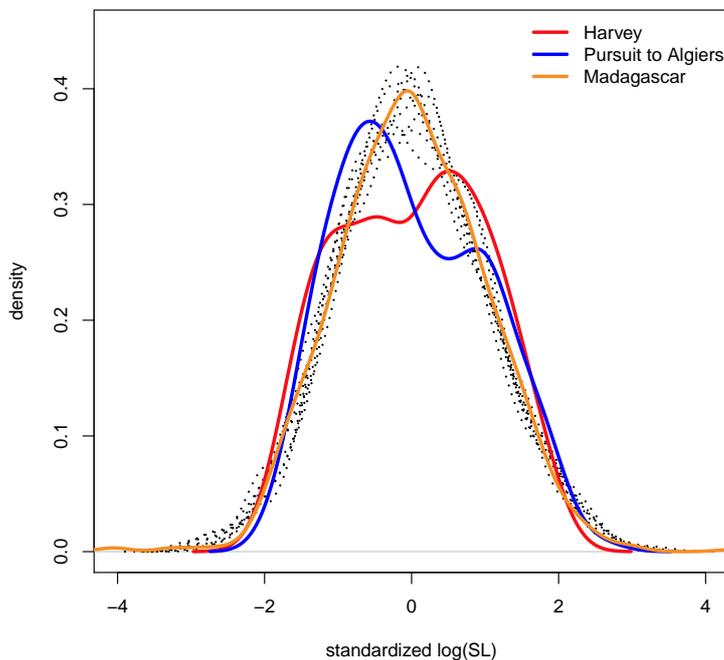


Figure 6.10: A comparison KDEs for the logged SLs of three named films against a background of eight others with approximately lognormal distributions.

In Redfern’s (2012a) study of the lognormality of SL distributions only 8/134 films satisfied the statistical hypothesis tests (of normality of log-transformed data) he applied. In chronological order these were *Captain Blood* (1935), *Fantasia* (1940), *The Great Dictator* (1940), *Brief Encounter* (1945), *Night of the Hunter* 1955, *Exodus* (1960), *Aristocats* (1970) and *Barry Lyndon* (1975). Since others are more optimistic about the prevalence of lognormality they might be taken as a standard against which the lognormality of other films can be compared. Their KDEs are plotted as ‘background’ in Fig 6.10, as dotted lines and without labelling to play down the visual impact.

There is a problem with Redfern’s hypothesis testing methodology in that results are dependent on the number of shots in the film. *Exodus* has 547 shots; if you watch the film twice and treat it as a single film the number of shots is doubled without affecting the SL distribution. Application of Redfern’s methodology now clearly rejects the hypothesis of normality, though nothing but the sample size has changed. This implies that films with a larger number of SLs - say greater than 1000, which includes a majority of the post-1960 films in Redfern’s sample and all but four from 1980 on - are more likely to be rejected as lognormal because of their sample size. This motivates

the use of more informal graphical methods of assessing (log)normality.

Fig 6.10 includes KDEs for the logged SLs of *Harvey* (1950) and *Pursuit to Algiers* (1945), about which I think there can be no argument about their non-normality. *Madagascar* (2005), with 1131 shots, looks reasonably normal to my eye, more so, in fact, than some films in the ‘reference’ collection. Close inspection and further analysis shows that the clear rejection of normality in Redfern (2012a) is attributable to two outliers in the left tail attributable to the log-transformation, and omitting these results in a judgment of normality more in accordance with the visual impact.<sup>5</sup>

That this kind of visual assessment, that (log)normality looks reasonable, does need further checking. Logged SLs can look fairly normal with KDEs even when the tail behaviour, as with *Walk the Line*, is not. Using probability plots in conjunction with KDEs is, therefore, advisable.

---

<sup>5</sup>One of the tests used in Redfern (2012a) is the Shapiro-Francia test of normality, and the decision rule used is a 5% significance level. The unmodified data for *Madagascar* is, on the log scale, rejected as normal at the 1% level also. On omitting the two outliers, associated with SLs of 1 deci-second, normality is acceptable at the 3% level.