Evolution in Hollywood editing patterns?

Mike Baxter[1]

16 Lady Bay Road, West Bridgford, Nottingham, NG2 5BJ, U.K.

(e-mail: michaelj.baxter@btconnect.com)

Working Paper 1: April, 2013

[1]Emeritus Professor of Statistical Archaeology, Nottingham Trent University, U.K.

Abstract *Cutting, Delong and Nothelfer (2010) use statistical methods to investigate the evolution of shot-length patterns in popular film. They argue, using what they call a 'modified autoregressive index' (mAR), that patterns are becoming increasingly clustered and also evolving towards 1/f structure, a pattern described in later publication as 'like those that our minds may naturally generate'. This paper shows that the interpretation of the mAR index is wrong. It is also shown that that the results concerning 1/f patterns can be interpreted in an equally plausible and much less 'exciting' way. That is, although there are undoubtedly interesting temporal patterns in the shot length structure, they can't be interpreted in terms of the 'evolution of Hollywood film' in the sense intended in the original paper.*

1 Introduction

The idea that aspects of film style can be both quantified and modelled using statistical methods dates back to Salt (1974). Application of such methods for modelling the internal patterning of shot length (SL) structure in films is of more recent origin. Software on the *Cinemetrics* website [2], which exploits some of the more basic methods, dates to the mid-2000s (see also Salt, 2006, pp.395-6) but there is little before that.

An important recent application of statistical time-series methodology for modelling internal SL patterns is Cutting, DeLong and Nothelfer (2010) (henceforth CDN). Methodologically this is one of the most ambitious and demanding applications of statistical methodology to cinemetric data I have seen. Using a database of 150 Hollywood films (1935-2005), indices of SL patterning are developed that, it is claimed, show that SLs within films have become increasingly clustered into packets of shots of similar length. This is described as 'evolution'; in the words of CDN, 'film editors and directors have incrementally increased their control over the visual momentum of their narratives, making the relations among shot lengths more coherent over a 70-year span'.

This analysis is linked with claims that 'the shot structure in film has been evolving toward 1/f spectra'. The claim is repeated in Smith *et al.* (2012, p.119) who explain this as meaning that 'whole films are characterized by rhythmic fluctuations that appear to guide viewer attention' and that 'the shot-duration patterns of popular film might increasingly be like those that our minds may naturally generate'.

These claims have attracted attention; the *New Scientist* ran an article entitled "Solved: The mathematics of the Hollywood blockbuster", and a web search based on the terms 'science', 'Hollywood' and 'blockbuster' will bring up several similar reviews. That CDN have a more sober and nuanced view of their work is evident in the debate between Barry Salt and James Cutting on the *Cinemetrics* website [3]. The claims are, nevertheless, sufficiently striking to merit careful critical analysis. Such an analysis has recently been provided by Redfern (2012), for the claims based on SL clustering patterns. He introduces some modifications into the procedure used in the the CDN analysis, that lead him to the conclusion

[2] http://www.cinemetrics.lv/

[3] http://www.cinemetrics.lv/articles.php

that the claims about the evolution of SL clustering cannot be sustained. Baxter (2012) suggested that the results presented in the supplement to CDN could, without any methodological modifications, be reinterpreted to suggest that evolution was not occurring. An attemp to reproduce the CDN and Redfern led to the realisation that the interpretation placed on a statistic fundamental to both papers was wrong.

The present paper has two purposes. Firstly, the reasons why previous interpretation is wrong are explained. That the internal SL structure of film has changed over time can be demonstrated clearly but has a much simpler explanation than that which has been previously and incorrectly advanced. Both CDN and Redfern (2012) rely on identifying the order of time-series autoregression (AR) models used to summarise temporal variation in SL patterns, interpreted as indicative of the degree of clustering of shots of similar lengths. The patterning identified in the papers is an artefact of the statistical methods used to derive them. It can be usefully interpreted in terms of aspects of film structure, but not in the way that has been advanced.

Secondly, there is a mathematical connection between AR modelling and the spectral analysis used to investigate $1/f$ ratios that might lead one to expect similar conclusions to be derived from either approach. The reinterpreted AR indices show little evidence of the evolution previously claimed, and it is suggested, using the results of CDN and a slightly different interpretive approach, that the same is true for their $1/f$ analysis.

To understand the arguments involved a little more statistical background than that given in either paper is provided in the next section. The methodology developed by CDN and adopted by Redfern (2012) is described and illustrated here. This is followed by a statistical critique of the methodology, including a simpler way of obtaining the same patterns as those previously observed. The reasons why previous analyses have uncovered these patterns is also explored. The concluding section looks at interpretational issues, including the validity of the explanations previously offered for the observed patterns

2 Methodology

A lot of the interest in applied time-series analysis, in other areas of application, lies in predicting future observations from past observations. This is often done in terms of statistical models that postulate a systematic relationship between a current observation and past observations, to which is added a random (or error, or disturbance) term.

There are many ways of specifying the systematic and random components and the way they are combined, but commonly linear models are used. There is also considerable choice among these; one of the simplest, and the only one of concern here – it is the one used in CDN – is the autoregression (AR) model, which assumes that an observation can be predicted from a linear combination of the observed values of previous observations. The order, or index, of the AR model is the number of past terms needed for satisfactory prediction.

Pairs of adjacent observations are separated by a lag of $h = 1$; the autocorrelation coefficient at lag 1, ρ_1, is the correlation between all such pairs; ρ_2 is the correlation between pairs separated by one intervening observation, and so on; $\rho_0 = 1$. The plot of ρ_h against h is the *autocorrelation function* (ACF).

The *partial autocorrelation coefficient* at lag h, α_h, is the correlation between pairs of observations after factoring out the effect of other intervening terms. It can be thought of as measuring the 'predictive power' that still remains after removing the effect of observations closer to the value to be predicted. A variety of algorithms exist for calculating α_h; $\alpha_1 = \rho_1$, and at greater lags the partial autocorrelations can be defined recursively in terms of the autocorrelations, but other methods can be used. The plot of α_h against h for $h \geq 1$ is the *partial autocorrelation function* (PACF).

Faced with real data and a plethora of models that might be used to describe them, choosing an appropriate model is not an easy task. The ACF and PACF are useful diagnostic tools. The PACF is widely used to assess whether an AR model is appropriate. The idea is that so long as observations at lag h (and those with smaller lags) retain predictive power the calculated α_h will be 'significantly' large and will drop sharply to non-significant values once they no longer have predictive value. The lag after which the drop occurs defines the order of the AR model.

Identification of where the 'drop occurs' is frequently not this obvious and statistical aids can be used. In both CDN) and Redfern (2012) their initial analyses examine α_h sequentially to see if the value exceeds $2/\sqrt{n}$ and identify the order of the AR model at the lag beyond which partial autocorrelations cease to be significant. The cut-off criterion, where n is the number of shots in the film, is based on the idea that (given appropropriate theoretical conditions) observed $\alpha_h > 2\sqrt{n}$ differ from zero at the 5% level of significance.

One other idea needs to be explained before examining the methodology of the two papers in detail, and that is *detrending*. Methods of the kind described above for model identification can be invalidated if there are strong trends in the data. One way to avoid problems posed by this is to detrend the data by first fitting an appropriate model of trend and subtracting the values it predicts for the series before proceeding further. Unlike CDN, Redfern does this and his approac is emulated in our analyses.

In cinemetric studies prediction is not of interest; rather, the hope is that the order of the AR models identified will provide an insight into the structure of temporal variation of SLs *within* films, and the way structure has changed over a period of time. Specifically, the order of an AR process has been equated with the extent to which shots tend to be clustered with others of similar length, with higher orders associated with greater clustering[4].

Apart from detrending, the important difference in the two papers lies in the way correlations are measured. CDN measure correlations using the (Pearson's product-moment) correlation coefficient whereas Redfern uses Spearman's rank-order correlation coefficient. This is equivalent to transforming the (detrended) SL data to ranks and is undertaken in the interests of 'robustness', to guard against problems caused by the skewed nature of SL distributions and the presence of outliers. It is argued that will lead to underestimation of

[4]For a discussion of the issue of clustering and reasons for it see the debate between Salt and Cutting on the *Cinemetrics* website, previously referenced, particularly Salt's contribution.

the 'true' order of an AR model[5]. Our analysis also emulates Redfern in using only 134 of the 150 films studied in DN, to avoid problems posed by SLs recorded as zero or negative.

Using the `pacf` function in `R`, Figures 1 and 2 show some results for *The Informer* (1935) and *The Grapes of Wrath* (1940). The plot to the left of Figure 1 is for the unmodified SLs, similar to those that form the basis of analysis in CDN. The plot to the right, using rank-transformed data after linear detrending, illustrates Redfern's (2012) approach. Figure 2 is similar, except that the left-hand plot shows results after detrending the SLs and before log-transformation.

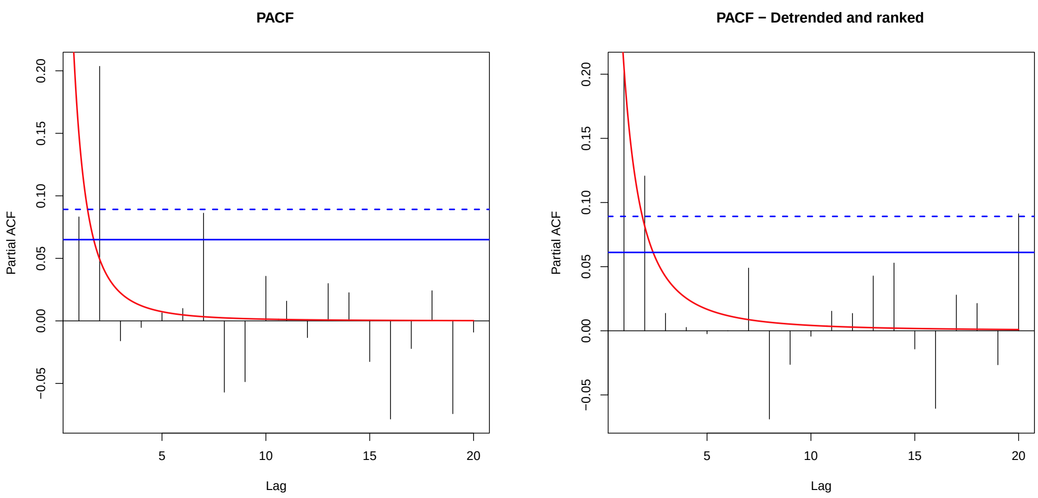

Figure 1: *PACFs for* The Informer *before linearly detrending the SL data and after linear detrending and a subsequent rank transformation.*

The horizontal dashed lines in all plots are at $2/\sqrt{n}$, so for *The Informer* would identify AR models of order 0 (to the left) and 2, following the methodologies just described. For *The Grapes of Wrath* orders of 1 and 3 are indicated. Suppressing any concern about these interpretations for a moment, it is sufficient to note that the use of rank-transformation does make a difference to the results.

CDN express dissatisfaction with this approach for several reasons and develop what they call a modified AR index (mAR). One motivation for this was a feeling that films with smaller values of n, and therefore larger $2/\sqrt{n}$, were 'penalized' in the sense that this 'generated smaller AR indices'. Another motivation, of greater consequence, was based on the observation that 'there can be much noise in partial-autocorrelation functions'. This arguably suggests that for many films an AR model may be inappropriate; however CDN deal with this isuue by using an 'exponential' model to smooth the PACF. This was coupled with a fixed rather than variable bound used to determine, via its intersection with the fitted curve, the mAR index. This approach is adopted in Redfern's analysis[6].

[5]It may be remarked that log-transformation of the data is likely to achieve the same ends.

[6]In Cutting *et al.* the mean of n in their sample of films replaces n in $2/\sqrt{n}$; Redfern prefers the median

This replaces the discrete AR index with a continuous mAR index and introduces the fundamental problem with which much of this paper is concerned..

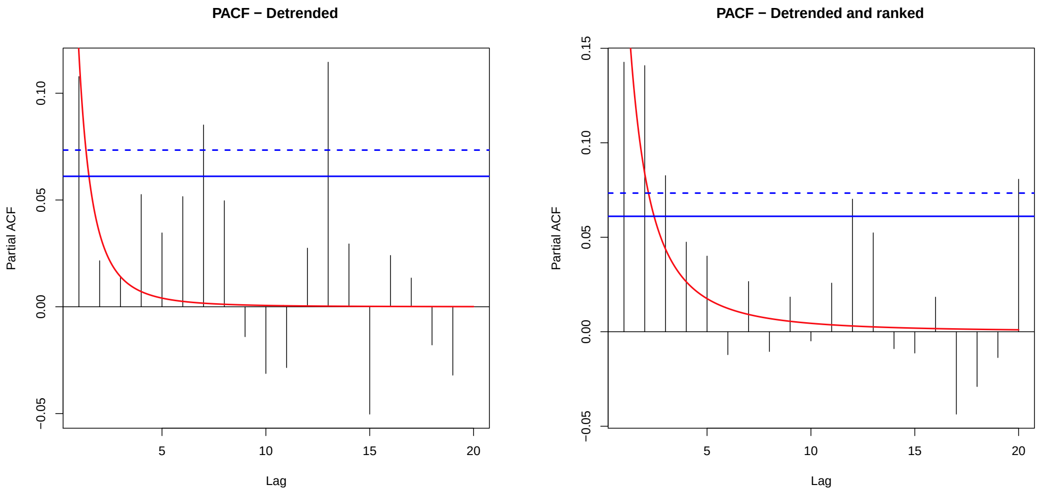

Figure 2: *PACFs for* The Grapes of Wrath *after linearly detrending the SL data and after a subsequent rank transformation. Horizontal lines are bounds, used to identify alternative AR indices from a plot.*

The fitted lines in both papers, and shown on the graphs, are obtained from an 'exponential' model of the form

$$\alpha_h = (1+h)^{-\beta}.$$

The `nls` function in `R` was used to fit the models in this paper, with results identical to those in Redfern (2012).

3 Commentary

This section addresses some purely statistical concerns about the merits of the mAR index. Interpretational issues are considered in the next section.

1. The appropriateness of an AR model is questionable for many of films. The PACF often bears little resemblance to what might be expected from an 'ideal' AR model(e.g., the left-hand panels of Figures 1 and 2). The ranked data, for the examples shown, are better and this is more generally, though by no means universally, the case.

2. CDNs comment about 'much noise in partial-autocorrelation functions', appears to acknowledged this. Estimating an mAR via the fitting of a smoothed 'exponential'

to the mean. These fixed cut-off criteria are represented by the solid horizontal lines in the figures. There is little difference, 0.065 using the mean and 0.061 using the median.

curve disguises rather than deals with the problem of inappropriateness in the first instance. An exponential model is often not appropriate, given the patterns in the data in many cases.

This can be approached both empirically and theoretically. Most simply, the exponential model is suited to data that, allowing for random variation, exhibit a pattern of decay. None of the plots in the figures really do this. Of the three PACFs shown for illustration in Figure 1 of CDN, fitting an exponential to *King Kong* looks reasonable, but not so *Detour*. If an AR model really is appropriate, then a PACF pattern with a fairly obvious 'drop' is expected, and in this case an exponential model should not be expected to be appropriate. The PACF for *Ordinary People* in Figure 1 of CDN is a case in point; it might be fairly interpreted as consistent with an AR index of 2, but can then be argued that then fitting an inappropriate exponential model to get an mAR of 2.47 is unnecessary.

3. A perhaps minor point is that mAR estimates will be affected by the sample of films used. The mean average shot length in films has been declining over the last 40 years or more, with the consequence that modern films will tend to have rather more shots than older ones from, say, the 1930s and 1940s. This means that while the PACF function for a film will remain unchanged the mAR calculation will depend on the sample of which it is a part, being somewhat larger for a sample of mainly modern films than one spread evenly over a much longer period.

What's being claimed here is that many film SL patterns are not adequately described by an AR model. If they are not, it is difficult to see what the mAR is measuring. If patterns can be adequately described by an AR model the mAR would seem to be redundant. The problme with mARs is, however, worse than this..

4 Alternative methods

The implication of the above is that interpretations offered for indices for any particular film, and patterns in them over time, can be called into question. The analysis to be described was intended to exploit the information in the α_h without imposing on them an interpretation in terms of an underlying AR model, or manipulating them in several stages to get an index of uncertain interpretability. The approach, based on principal component analysis (PCA), led to a rather simple conclusion that allows the PCA methodology that revealed itto be jettisoned[7]. The idea was to apply PCA to the unstandardised α_h, and see if there was any pattern in plots of the PCA scores, against each other, or external

[7]PCA is a standard technique described in any good text on multivariate analysis. It produces new variables – principal components or PCs, that are linear combinations of the α_h, that are uncorrelated and explain successively decreasing amounts of variation in the data. It is hoped that some of the PCs are 'interpretable'. Only the first PC was, and it was largely determined by α_1 so patterns can be explored using the latter, forgetting about the mechanism used to reach this conclusion.

factors such as date or genre of the films. This is in the same spirit in which the various AR indices have been used.

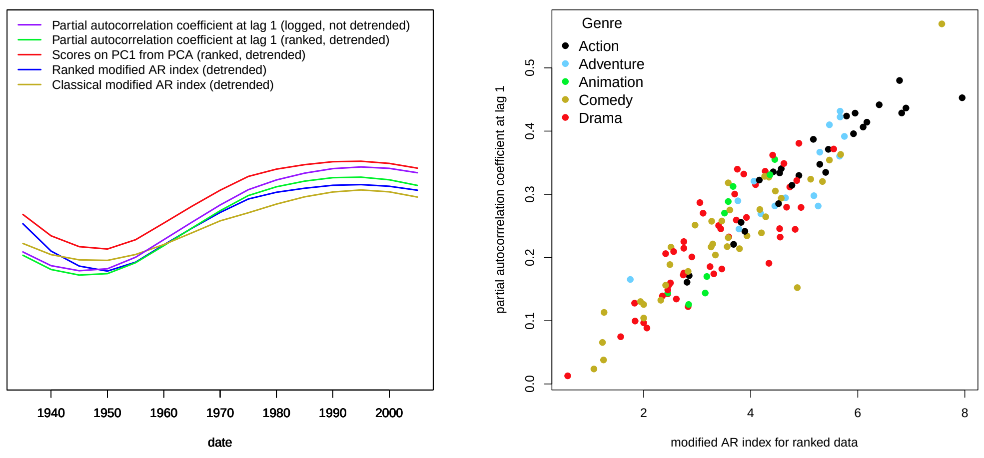

Figure 3: *To the left plots against date of different 'indices' intended to measure aspects of SL structure in films, for different data treatments and methodologies. Plots are based on loess smooths with a span of 3/4 using localised robust quadratic regression fitting. The right-hand plot is of* α_1 *and mAR for ranked detrended SLs.*

Only the first principal component (PC1) seemed obviously interpretable; in particular a plot of scores on PC1 against date showed a pattern remarkably similar to that of Figure 3(b) in Redfern (2012) for his mAR index based on ranked SLs. The pattern doesn't depend on the number of lags used in the PCA. The first PC is largely determined by the value of α_1 and this is all that's needed to obtain the same pattern as in Redfern's analysis.

This is illustrated in the left-hand panel of Figure 3 which contrasts the pattern for the original PCA analysis, the first partial autocorrelation coefficient (of ranked, linearly detrended data), Redfern's (2012) mAR index for ranked SLs and, for good measure, his mAR index for the detrended but unranked data. Scales for the ordinates are not, for the most part, comparable so have been omitted; plots are overlaid to emphasise the similarity of pattern, but nothing is to be read into the distance between lines[8]. The results for logged SLs, subsequently detrended, are almost identical to those for ranked detrended data and not shown. The curve shown for α_1 for logged data does not involve trend removal of any kind.

An obvious question to ask is why are the results so similar. It suffices to concentrate on a comparison between the first partial autocorrelation coefficient for the ranked detrended data, and the mAR index for the ranked data. The strength of the relationship is illustrated in the right-hand plot of Figure 3, where the genre of a film is also indicated.

[8]The lines are loess smooths using a span of 0.75, and robust quadratic regression models for the localised fitting involved.

Superficially the methodology used to arrive at these very similar results looks rather different. Use of α_1 for showing changing structure over time is pehaps an obvious thought anyway, but was indicated here by the results of the PCA. The PCs are linear combinations of the α_h only the first of which is important here. From the exponetial model previously given, and for a fixed cut-off point, c,

$$\text{mAR} = c^{\beta} - 1$$

so that any variation in the mAR depends on variation in β which is estimated by non-linear least squares methods and is a function of all the 20 lags used. The equivalence between the results suggests that this estimate is very strongly dominated by the magnitude of α_1 but this is less transparent than for the PCA.

Very heuristically, it can be suggested that the larger α_1 is to begin with, the 'longer' the fitted curve is likely to take to fall to its limit of zero, and the larger the mAR, where the curve intersects the fixed cut-off boundary, will be. That this happens can be confirmed via simulation using the `arima.sim` function in `R`, simulating 1000 observations from an AR model of order 1 for different values of α_1, several times for each value. Thus for $\alpha_1 = 0.5$ and ten simulations mAR indices in the range about 3.45-4.00 are obtained – call a typical value 3.7. Reduce α_1 to 0.1 in steps of 0.1 and typical values are 3.1, 2.5, 1.7 and 1.2, with 0.8 for $\alpha_1 = 0.05$.

This is a very small experiment and there is a lot of variation about the typical values quoted, but the pattern is clear enough. If 134 such analyses were to be performed, varying α_1 between 0.05 and 0.50, and assigning 'dates' to each simulation with later dates corresponding to higher α_1, plots like those to the left of Figure 3 and more linear might be expected. This would show an 'evolution' in mARs that, if the interpretation of CDN was accepted, would be interpreted as an increase in clustering into packets of shots of similar length.

It cannot be so interpreted. In CDN the AR index is replaced by the mAR to deal with the fact that the real data often does not look like an AR process. By contrast, here the simulated data, by construction, are from an AR process, so there is no need for an mAR and no 'evolution' since the processes are all of the same order. The suggestion of 'evolution' in the simulated example is entirely an artefact of the methodology used to construct mARs, whereby larger α_1 generate larger mARs, regardless of the true order of the underlying process.

The above is based on the simulation of AR models of order 1. It is more tedious to investigate models of higher order, as a lot more variation in the parameters that define a model is possible. I have not looked at this systematically, but it is easy enough, using simulated data, to show that mARs more nearly reflect the values of the larger partial autocorrelations present than the order of the underlying process. As a rough illustration of the kind of variation to expect, using detrended and ranked data, there are 49 films with an AR index of 2, whose mARs range from 1.94 to 7.57, and 21 films with an AR index of 3 whose mARs vary between 2.45 and 5.79.

Only a brief discussion of $1/f$ analyses is attempted, using the results reported in the supplement to CDN. They smooth the power spectrum for the (standardized) SLs of a film

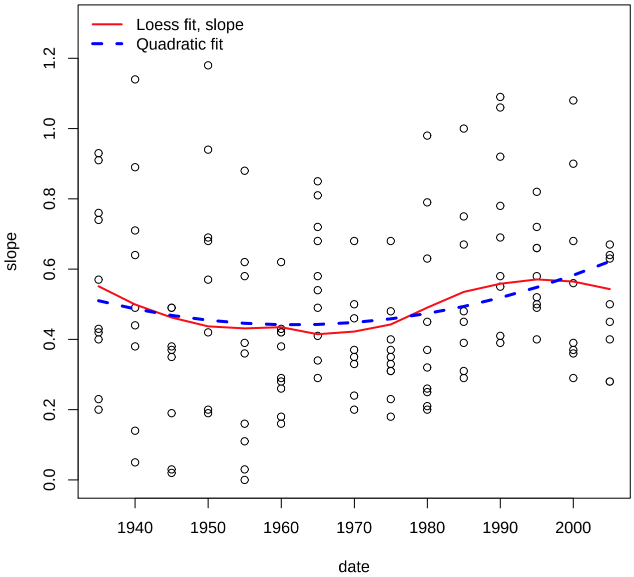

Figure 4: *Plots of α ('slope') against date using the model of $1/f$ patterning developed in Cutting* et al. *(2010) and a loess smooth of the same data. See the text for detail.*

using 7-10 estimated points, to which a line determined by a parameter, α, is fitted. The closer α is to 1 the closer the film is to exhibiting $1/f$ patterning. The estimated values of α are plotted against date, a line through the points fitted with a quadratic regression model, and 'evolution' towards $1/f$ patterning claimed on the basis that the right 'arm' of the fitted model is moving smoothly upward towards 1. This is illustrated in Figure 4 along with a loess smooth to the same data, of the kind used in the previous figure[9].

It can be remarked in passing that, on the basis of Figure 3 in CDN, the propriety of fitting a linear model to estimate α for some films can be questioned; the data do not have the shape to merit such a fit. Even so, and ignoring this, the loess smooth to the data in Figure 4 shows no sign of 'evolution', the short period beween about 1970 and 1990 being the only one when anything of the sort was happening. The question has been asked (by Salt) as to why, if there has been an evolution to $1/f$ patterning since 1935, we aren't closer to it. A prefectly plausible answer is that no such evolution is taking place. I am not claiming that the loess smooth is 'right' and the quadratic smooth 'wrong' – there is too much variation in the data for either to be regarded as compelling evidence for a 'trend' of any kind. That the data can tolerate this kind of variation in the fitted lines, though, is reasonably compelling evidence that claims for an evolution to $1/f$ patterning cannot be strongly pressed, if at all.

[9]For consistency the 134 films analysed elsewhere in this paper have been used. There are slight differences between the figure and Figure 2(c) in CDN as a consequence, but if all 150 are used the loess curve suggests, even more strongly, that 'evolution' is not occurring.

5 Discussion

To summarise the conclusions of this paper:

1. Looking at the patterns in SL 'structure' it is quite possible to reinterpret, without any strain, the results in CDN. The claim is that evolution is taking place in clustering into packets of shots of equal length, based on a linear fit of mAR indices against date. If linearity is not *assumed*, and a non-parametric fit applied that allows for modest departures from linearity, it is clear that changes have occurred over the 1935-2005 period, but the results are consistent with not much having changed in the last 30 years or so. The same is true of claims about evolution towards a $1/f$ pattern, except that they are based on the almost linear 'arm' to the right of a quadratic fit to the data (Figure 4).

2. Neither CDN nor Redfern (2012) query the underlying assumption that fitting AR models to the data is a reasonable thing to do. Film-by-film inspection of the PACFs shows that a significant number bear limited resemblance to what might be expected from an ideal AR process, not even approximately so in some cases. The observation in CDN that PACFs can exhibit a lot of 'noise' can be read as a tacit admission of this. The mAR is intended to be a continuous measure of the order of an AR process, but if an AR model is inappropriate in the first place it does not correct for this.

3. Putting aside this last objection, and even if an AR model is valid, the mAR does not measure what is claimed for it. It is supposed to be a continuous measure of the order of an AR process, interpreted in turn as a measure of the degree of clustering. What it actually appears to be, a mathematical/statistical artefact of the way it is constructed, is a surrogate measure for the (partial) autocorrelation at the first lag. Assuming an AR model is valid this autocorrelation only provides information on whether the process is of order zero or not.

   To put this another way, processes of the same order can give rise to entirely different mARs, and processes of different orders can give rise to the same mAR. The mAR is of no value as a diagnostic tool for the order of an AR process.

4. What this means is that results can be interpreted in terms of the autocorrelation at lag 1, but not in terms of clustering as understood in CDN and Redfern. Depending a little on the way the data are treated, the general pattern is of a decline between 1935 to about 1950, a steady rise to about 1980 and not much changing since then.

   This pretty much imitates (inversely), up to the early 1990s, the pattern of mean ASL against year for US feature films in Salt (2009, p.378). There was some debate, in the Salt/Cutting interchange on the *Cinemetrics* website about the relationship between changes in ASLs and the kind of patterns that were reported in CDN. Cutting asserts, correctly I think, that a causal relationship between ASLs and clustering/correlations is unlikely, but Salt advances what I think is a plausible heuristic reason for expecting an association (which does not imply causality) between clustering/correlations and ASLs, the former increasing as the latter decreases.

If Salt's argument 'works' it would appear to do so up to about 1990, and for ASLs down to about 6.5. Mean ASLs hover around 7 to 6.5, with a slight decline from about the mid-1970s to 1990 then show a sharp decrease. This sharp decrease is not exhibited by any of the patterns displayed in Figure 3.

5. It is noticeable in the right-hand plot of Figure 3 that the larger partial autocorrelations and mAR indices are dominated by action and adventure films, mainly absent for the lower values, which are drama and comedy. CDN do not claim that their sample is random in any sense – quite the opposite in fact – and genres are not equally distributed across the years of their sample, action and (to a lesser extent) adventure being over-represented in the later years and drama in the earlier years. This selection may or may not reflect trends in the numbers of each genre being made, or in audience tastes (films were selected to be high-grossing, or otherwise indicated as 'popular') – the matter is not discussed in detail.

   The temporal patterns exhibited by α_1 and the other indices in the left-hand plot of Figure 3 may thus possibly be associated with genre selection. Without larger and statistically representative samples it is impossible to say whether such patterns as are observed reflect a tendency common to all genres, to a subset of genres or, and it is obviously related, trends in genre popularity. Some debate on this topic takes place in the Salt/Cutting interchange, where Salt's heuristic arguments noted above were particularly directed to modern action films.

The scope and ambition of the work initiated in CDN is impressive, technically demanding at the statistical level, and leads to some 'eye-catching' conclusions. The analyses presented here show that, on the basis of the numerical results reported in CDN, somewhat more prosaic interpretations are equally plausible. The additional statistical arguments advanced here suggest, beyond this, that the original interpretations are also unsustainable.

A little care is needed here. The possibility that there is increased clustering into packets of shots of similar lengths still exists; it's just that analyses based on the mAR don't show this, and other ways of demonstrating the phenomenon, if it is real, are needed. There undoubtedly are patterns in the data that invite explanation, and these are to do with changing patterns of the first order autocorrelation coefficient over time. The challenge is to explain why these rose between about 1950 and 1980 and seem to have stabilised, not, in the last 20 years or so, reflecting changes in ASL patterns.

References

Baxter, M. 2012: *Cinemetrics Data Analysis*. http://www.mikemetrics.com/

Cutting J.E., DeLong J.E. and Nothelfer C.E. 2010: Attention and the evolution of Hollywood film. *Psychological Science***21**, 440-447.

Redfern, N 2012: Robust estimation of the modified autoregressive index for high grossing films at the US box office, 1935 to 2005.
http://nickredfern.files.wordpress.com/2012/11/nick-redfern-the-mar-index-for-hollywood-films1.pdf

Salt, B. 1974: Statistical style analysis of motion pictures. *Film Quarterly* **28**, 13-22.

Salt, B. 2006: *Moving Into Pictures*. London: Starword.

Salt, B. 2009: *Film Style & Technology: History & Analysis (3rd edition)*. London: Starword.

Smith, T.J., Levin, D. and Cutting, J.E. 2012: A window on reality: perceiving edited reality. *Current Directions in Psychological Science* **21**, 107-113.