



Lancaster University
MANAGEMENT SCHOOL

Lancaster University Management School
Working Paper
2005/052

**Habit, aggregation and long memory: Evidence from
television audience data**

David Byers, David Peel and Dennis Thomas

The Department of Economics
Lancaster University Management School
Lancaster LA1 4YX
UK

© David Byers, David Peel and Dennis Thomas
All rights reserved. Short sections of text, not to exceed
two paragraphs, may be quoted without explicit permission,
provided that full acknowledgement is given.

The LUMS Working Papers series can be accessed at <http://www.lums.lancs.ac.uk/publications/>
LUMS home page: <http://www.lums.lancs.ac.uk/>

**HABIT, AGGREGATION AND LONG MEMORY: EVIDENCE FROM
TELEVISION AUDIENCE DATA**

J.D. Byers

University of Cardiff

D.A. Peel

University of Lancaster

and

D.A Thomas

University of Wales, Aberystwyth

Address for Correspondence:

J.D.Byers
Cardiff Business School
Colum Drive
Cardiff CF1 3EU
Wales
UK

Tel: +44 2920 875276
Email: byersjd@cardiff.ac.uk

July 2005

Abstract

Many economic outcomes appear to be influenced by habit or commitment, giving rise to persistence. In cases where the decision is binary and persistent, the aggregation of individual time series can result in a fractionally integrated process for the aggregate data. Certain television programmes appear to engender commitment on the part of viewers and the decision to watch or not is clearly binary. We report an empirical analysis of television audience data and show that these series can be modelled as $I(d)$ processes. We also investigate the proposition that temporal aggregation of a fractionally-integrated series leaves the value of d unchanged.

Keywords: Long Memory, Fractional Processes, Aggregation, Habit.

JEL Classification: C22., D12

HABIT, AGGREGATION AND LONG MEMORY: EVIDENCE FROM TELEVISION AUDIENCE DATA

Introduction

Decisions which have to be repeated often display a considerable degree of persistence which appears to reflect habit or commitment. Examples include attendance at religious ceremonies, support for political parties, attendance at sporting and theatrical events and consumption of tobacco, alcohol and other drugs. Typically, we do not have data at the level of the individual but must work with aggregate data and the two may exhibit quite different properties. One possibility is 'short memory' at the individual level but 'long memory' at the aggregate level .

In this paper we report an analysis of several data series generated by a particular set of choices which are often believed to be strongly affected by habit - the viewing of television programmes. We present evidence that - consistent with a cross-sectional aggregation result which we discuss below - the temporal aggregate of individual choices exhibit long memory as represented by a fractional differencing parameter, d , greater than zero. We also investigate the proposition that temporal aggregation of a fractionally-integrated series leaves the value of d unchanged - an implication of the self-similar behaviour of long-memory processes.

A characteristic of the data is that, at the individual level, it consists of a sequence of values which are either 1 or 0 corresponding to the decision to partake in the activity or not. Since the standard choice problem in economics involves selecting from a continuum of possible outcomes, it may appear that the relevance of the result is limited. Consider, however, the mundane example of purchasing baked beans.

Although the rate of consumption can vary continuously, the purchaser also has to

select a brand from the range of brands available. The decision to buy the same brand as last time or to try some other is a 0/1 decision, one that firms expend substantial sums to influence. It would be easy to multiply examples but a little thought suggests that this aspect of choice is widespread.

Theoretical Issues

At the individual level, it would appear plausible that the inertia resulting from habit will lead to an autoregressive component in behaviour. This may be justified by appeal to either 'myopic' or 'rational' theories of addiction (Becker and Murphy, 1988, Becker, Grossman and Murphy, 1994). The intertemporal structure of *aggregate* outcomes, however, may not be so straightforward since, to take a well-known example, the sum of n independent $AR(1)$ processes with different parameters will, in general, be $ARMA(n, n-1)$. Granger (1980) pointed out that the panel aggregate of a wide range of time series processes may display the long memory property of a fractionally integrated process. Specifically, suppose that the choice over the level of some variable, x , of the j^{th} agent can be represented by the process,

$$x_{j,t} = \alpha_j x_{j,t-1} + y_{j,t} + \beta_j W_t + \varepsilon_{j,t} \quad (1)$$

where $y_{j,t}$, W_t and ε_t are independent processes with the latter white noise, then the

aggregate series, $x_t = \sum_{j=1}^J x_{j,t}$, derived by summing these dynamic microrelationships

over J agents will be integrated of order d if the α_j are drawn from a Beta

distribution over the range $(0, 1)$ with a probability mass of zero at unity. The size of d

is affected by the order of integration of the y and W processes but x_t is fractionally

integrated even if these are $I(0)$. Byers, Davidson and Peel (1997, 2000, 2002) apply

this result to support for political parties in the UK and a number of other countries.

An approach based on the Granger aggregation result is not practicable in the current contextⁱ but an alternative can be found in Willinger *et al* (1995) and Taqqu *et al* (1997) who consider a stochastic mechanism which generates a sequence of ones - the source is switched on - and zeros - the source is switched off. An 'ON' period is a sequence of 1s, and an 'OFF' period is a sequence of 0s. Taqqu *et al* analyse the case where ON and OFF periods alternate with i.i.d. lengths, though with possibly different distributions, and show that the superposition of a large number of such series may be fractionally integrated. The crucial condition is that the upper tails of the distribution functions, $F_i(x)$ $i = ON, OFF$, of the ON and OFF periods decline in accordance with a 'power law', i.e. $1 - F_i(x)$ behaves like $x^{-\alpha_i}$ as $x \rightarrow \infty$ with $1 < \alpha_i < 2$. As a result there are non-negligible probabilities of very long ON and OFF periods. If all the sources are identical, the resulting aggregate time series will have a differencing parameter given by $d = 1 - \alpha_{\min} / 2$. If the sources are heterogeneous the source with the smallest α dominates as $T \rightarrow \infty$. Notice that $\alpha_{\min} = 2$ implies an $I(0)$ series.

At the individual level, we have a nonlinear, 'regime-shifting', process but, at the aggregate level, the observed process has linear characteristics. As Granger and Terasvirta (1999) demonstrate, simple nonlinear processes can exhibit 'misleading' linear properties. They show that the autocorrelations of the non-linear, regime-shifting, process:

$$x_t = \text{sgn } x_{t-1} + e_t \tag{2}$$

where e_t is *nid*, can exhibit hyperbolic decay, indicating long memory. This is, of course, not an aggregation result but $\text{sgn } x_{t-1}$ is an ON/OFF series and the apparent long memory emerges as the probability of a change in regime declines.

A Model of Television Viewing

Consider the following informal model of television viewing. On day (or week) t an individual has to decide whether to watch a programme or not. The Rational Addiction approach suggests that this decision is taken in the light of previous choices and expected future utility. In particular, by watching a programme an individual builds up a stock of 'consumption capital' which raises the current utility from watching the programme but which depreciates over time. Given the current capital stock, an individual will choose to watch a programme if the expected utility from doing so exceeds the utility of not watching. Viewing the programme increases the stock of consumption capital and makes it more likely that the programme will be watched again. The model allows for random shocks to utility with 'positive' shocks raising current utility and making it more likely that the programme will be watched again and 'negative' shocks reducing enjoyment and, perhaps, resulting in crossing a threshold which means that the next broadcast will not be watched. The effects of depreciation may then result in the programme not being watched again regularly until a sufficiently large positive shock occurs.

Suppose we record a value of 1 when the programme is watched and 0 when it is not watched. Habitual viewers will exhibit long sequences of 1 s with occasional 0 s and uncommitted viewers only occasional 1 s among many 0 s. For really committed viewers the $0N$ period approaches infinity, i.e. they always watch. Other viewers, however, may become attached for a number of periods, then, for some reason, not watch for a time and later become a regular viewer once more. Some individuals may have extremely long sequences of zeros – in effect, they never watch. The implication of the model is, however, that for each individual there exists an alternating sequence

of ON/OFF periods in which the lengths of the ON and OFF periods are random variables.

Aggregate data on viewing numbers is collected by various agencies for television companies. If individual households behave in accordance with the model outlined above then these aggregate series may exhibit long memory.

The Data

The data which we analyse consist of daily viewing figures for three programmes broadcast by Sianel 4 Cymru (S4C) - the UK's Welsh-language TV serviceⁱⁱ. The programmes are Heno, a general interest programme covering topical issues, Pobol Y Cwm, a soap opera, and Newyddionddion, the evening news programme. Each lasts for half an hour and the three are broadcast in sequence between 6:30 pm and 8:00 pm from Monday to Friday. For Heno and Newyddionddion the series are nearly continuous over the period from 2 January 1995 to May 22, 2000, being broken only by occasional public holidays - especially around Christmas. The series for Pobol Y Cwm runs continuously only from 11 September 1995 since the programme was not broadcast during the summer months of that year. Breaks in the series after this time are mainly around Christmas. The occasional missing values in each series are interpolated by the simple expedient of taking the average of the same day in the previous and succeeding weeksⁱⁱⁱ.

The viewing figures for each of the three programmes vary over the days of the week and over the weeks of the year. Heno and Pobol Y Cwm, in particular, display strong annual cycles. Since there is a danger of confusing the pattern of low-order autocorrelations generated by this cycle with any long memory in the data, we remove the former using a set of trigonometric seasonals in the form of $S/2$ pairs of sines and cosines where S is the number of seasons and the typical term is

$$\begin{aligned}
& A_s \cos(\lambda_s t) + B_s \sin(\lambda_s t) \\
& \lambda_s = 2\pi s / S \quad s = 1, \dots, S / 2
\end{aligned} \tag{3}$$

In particular, $s = 1$ corresponds to the annual cycle of 260 days, $s = 13$ corresponds to the monthly cycle and $s = 52$ corresponds to the weekly cycle. Harvey (1989) discusses the use of this method and points out that, in practice, most of the seasonal variation in a series be accounted for by the initial few sine/cosine pairs plus a small number of other frequencies. All three series were adjusted by removing the first eight trigonometric seasonals plus those corresponding to the weekly cycle and its first harmonic in addition to a constant and trend.

The adjusted series exhibit remaining stochastic seasonality at the weekly frequency. We deal with this in two distinct ways: firstly by allowing seasonal autoregressive and moving average terms to enter the model for each adjusted series and, secondly, by analysing weekly series obtained from the original series by separate consideration of the different days of the week. This procedure has the advantage of allowing us to compare the estimates of the long memory parameter at two different levels of aggregation since theory suggests (Chambers, 1998) that the fractional differencing parameter will be the same both for an original series and for an 'aggregate' series obtained by selecting every n^{th} value from the original series. The weekly series are adjusted for the annual cycle by removing the first two trigonometric seasonals

Empirical Results

For each of the adjusted daily series the sample autocorrelations decline slowly and are positive for a substantial number of lags. The autocorrelations of Heno are positive to lag 34, those of Newyddion are positive to lag 44 while the first negative autocorrelation of Pobol Y Cwm occurs as late as lag 100. Lo's Modified Rescaled

Range test (Lo, 1991) provides a formal test of long memory. The test rejects the null hypothesis of $d = 0$ against $0 < d < 0.5$ in the upper tail and involves a correction for short-term dependence in the data using the Newey-West variance estimator to scale the range of the process. We report results for lags of 0, 1, 3 and a value, q , determined by Andrew's (1984) data dependent rule for AR(1) processes. The latter would appear to be an appropriate simple alternative model for the aggregate data. Table 1 reports the values of the test for the three daily series with the number of lags in the data-dependent test given in parentheses. The test clearly rejects $d = 0$ against $d > 0$ for Heno and Pobol Y Cwm but is less clear in the case of Newyddion.

We obtain a potential model for each series by comparing the values of the Schwartz Information Criterion (SIC) for a range of ARFIMA processes. Given the evidence of stochastic seasonality in the data, we conduct the search over a range which includes first order seasonal ARMA components. Specifically, we assume multiplicative weekly seasonality and compare models of the form $ARFIMA(p, d, q)(p_s, 0, q_s)$ for $2 \geq p, q \geq 0$ and where p_s and q_s can be 0 or 1.

Weeding out models with implausible parameter values or other obvious deficiencies we arrived at the models which are reported in Table 2. In each case the fractional differencing parameter is significantly different from zero. The cross-sectional aggregation result implies that each aggregate series is a linear, Gaussian, $I(d)$ process for $d \geq 0$. To test this we report statistics on skewness and kurtosis along with tests for serial correlation and ARCH effects. Notice firstly, that the ARFIMA models do a remarkable job of providing a clean set of residuals. Pobol-y-Cwm is something of an exception since the Ljung-Box Q statistics are significant from Q(10). Estimating a higher-order seasonal AR fails to solve this problem but also leaves the estimate of d

largely unchanged. We look at this issue below when examining the weekly series.

All three series appear to be free of low-order ARCH effects.

The normality tests are less satisfactory. Skewness and kurtosis tests are notoriously sensitive but these three tell a consistent story of positive skewness and excess kurtosis. This is borne out by examination of the histograms. It is, however, noticeable that while the skewness and kurtosis statistics for Heno and Newyddion are similar to each other, the residuals for Pobol Y Cwm appear to be much closer to normality with an insignificant Sk statistic. One possible explanation is that the residual series are subject to outliers which, arguably, would often be positive because viewers would be much more likely to watch, say, the national or international news if some major event had taken place and much more likely to watch a soap opera if it were known - as it often is - that a major plot development was about to take place. The percentages of residuals which are more than 1.96 (2.65) standard deviations away from zero for Heno, Newyddion and Pobol Y Cwm are, respectively, 1.54 (4.77), 1.19 (4.98) and 1.12 (5.44). Adopting the crude process of replacing those residuals which are more than 2.65 standard deviations large by a value of zero reduces the skewness statistics for Heno and Newyddion by 60% and renders the kurtosis statistics insignificant. The skewness coefficient of the Pobol Y Cwm series falls further and the kurtosis statistic now becomes insignificantly different from zero. In the latter case the Jarque-Bera test is also insignificant at 0.828. There is, therefore, evidence that outliers are responsible for the rejections of normality in the daily series.

Table 3 gives the estimated models for the five weekly series for each programme.

Here the SIC procedure chose a pure fractional model in each of the 15 cases. In all but one case (the Wednesday series for Newyddion) the estimated d is significantly different from zero, often at the 1% level. As with the daily data, the chosen models

provide residuals which are almost entirely free of serial correlation and ARCH effects. These results also show that the problem with the daily Pobol Y Cwm series can be attributed to the observations for Wednesday. Only three ARCH statistics are significant, indicating first order ARCH effects in two of the series. The skewness and kurtosis statistics are, again, often significantly different from zero but less so in the case of Pobol Y Cwm than the other two series. The Jarque-Bera test rejects normality at 5% for 10 of the 15 series, the exceptions being the Monday and Thursday series for Newyddion and the Tuesday, Wednesday and Friday series for Pobol Y Cwm. However, setting the 'large' (> 2.65 standard deviations) residuals equal to their expected value of zero renders all of the kurtosis coefficients and all but one of the skewness coefficients insignificant at 5%. The Jarque-Bera statistics for these trimmed residuals are all insignificant at 5%.

The lack of normality in the residuals for both daily and weekly models is the main inconsistency with the predictions of the Taqqu *et al* theorem. Removing this problem by removing 'large' residuals seems arbitrary, though one should bear in mind that the procedure removes both large negative as well as large positive values so that it will not obviously eliminate the positive skewness in the untrimmed residuals. One possible explanation lies in the sample size. The Taqqu *et al* theorem is a large sample result and it could simply be that the sample used here - 420 - is not large enough and that viewers who watch only occasionally are overrepresented. In this context it is interesting to note that non-normality is less of a problem with Pobol Y Cwm, the series where one would expect habit to play the most important role, with fewer people watching just the occasional episode, as opposed to current affairs and news programmes which may occasionally attract surprisingly large audiences. The

optimistic interpretation is then that the aggregation effect is strong enough to generate long memory even in quite small panels.

Finally, as mentioned above, the temporal aggregation involved in moving from daily to weekly series should leave the differencing parameter unchanged. The average d value for Heno is 0.145 compared to the estimated value for daily data of 0.207, the average d value for Newyddion is 0.138 which is very close to the estimated d for the daily data of 0.131 and the average d value for Pobol Y Cwm at 0.206 is close to the estimated d for daily data of 0.185. Thus in two of the three data series the temporal aggregation result appears to carry through. The exceptional case, Heno, is the only series which includes a non-seasonal AR coefficient. Estimating a pure fractional model for this series results in a d of 0.146 which is much closer to the average d derived from the weekly data, though the residuals fail the Q test. Hence, it would appear that the inclusion of an autoregressive component has had the effect of boosting the value of the differencing parameter for this series. It should also be noted that the estimated d 's for purely fractional process for Newyddion and Pobol Y Cwm are 0.146 and 0.197 respectively.

Conclusion

Nothing in the basic theory requires that aggregate data on television viewing will be fractionally integrated. Both the Granger and Taqqu *et al* approaches to aggregation require, in effect, an appropriate set of 'weights' and some restrictions on a behavioural parameter to deliver a d between zero and unity rather than a d of zero or unity. Nevertheless, such outcomes do appear to arise quite frequently in practice. In this paper we appealed to habit or commitment in a context where behaviour can be characterised in such a way as to apply the Taqqu *et al* result. We find that all three daily series of viewing figures, once adjusted for seasonality, behave broadly in line

with the Taqqu *et al* result. In addition, all of the weekly series appear to be adequately modelled as purely fractional process with Gaussian noise after the removal of an annual cycle. The fractional differencing parameters are quite small, but significantly different from zero, and there is no strong evidence of further structure, linear or nonlinear. By considering both daily and weekly data we were able also to investigate the effects of temporal aggregation and show that the different frequencies share similar values of d .

References

- Andrews, D., 1984. Heteroscedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica*, **59**: 817-858
- Becker, G.S., and K.M. Murphy, 1988. A Theory of Rational Addiction, *Journal of Political Economy*, **96**: 677-700.
- Becker, G.S., Grossman M., and K.M. Murphy, 1994. An Empirical Analysis of Cigarette Addiction, *American Economic Review*, **84**: 396-418.
- Byers D., Davidson J, and D. Peel, 1997. Modelling Political Popularity: a Correction. *Journal of the Royal Statistical Society A*, **160**: 471-490.
- Byers D, Davidson J. and D. Peel, 2000. The Dynamics of Aggregate Political Popularity: Evidence from Eight Countries, *Electoral Studies*, **19**: 49-62.
- Byers D., Davidson J, and D. Peel, 2002. Modelling Political Popularity: an Analysis of Long Range Dependence in Opinion Poll Series. *Journal of the Royal Statistical Society A*, **165**: 187-189.
- Chambers, M.J., 1998. Long Memory and Aggregation in Macroeconomic Time Series, *International Economic Review*, **39**: 1053-1072.
- Granger CWJ. 1980. Long memory relationships and the aggregation of dynamic models, *Journal of Econometrics* **14**: 227-238.
- Granger ,C.W.J. and Terasvirta, T. 1999. A Simple Nonlinear Time Series Model with Misleading Linear Properties. *Economics Letters*, **62**: 161-165.
- Harvey, A.C. 1989. *Forecasting, Structural Models and the Kalman Filter*, Cambridge, CUP.
- Lo AW. 1991. Long-term memory in stock market prices', *Econometrica* **59**: 1279-1313.
- Taqqu, M.S., Willinger, W. and R. Sherman, 1997. Proof of a Fundamental Result in Self-Similar Traffic Modelling. *Computer Communications Review*, **27**: 5-23.
- Willinger, W., Taqqu, M.S., Leland, W.E. and D.V. Wilson, 1995. Self-Similarity in High-Speed Internet Traffic Measurements. *Statistical Science*, **10**: 67-85.

Table 1				
MRR Tests				
	Truncation Lag			
	0	1	3	Andrews' q
Heno	2.247	2.114	1.901	1.733(5)
Newyddionddion	1.819	1.696	1.564	1.454(5)
Pobol Y Cwm	5.305	4.754	4.219	3.427(8)

Upper Tail Critical Values: 5%, 1.747, 10%, 1.620

Table 2**ARFIMA Models: Daily Data**

	<i>d</i>	AR1	SMA	Sk	Ku	SC(1)	Q(5)	Q(70)	ARCH(1)	ARCH(5)
Heno	0.207** (0.024)	-0.132** (0.036)	0.059* (0.026)	0.585**	1.364**	0.112	7.659	89.282	2.507	5.226
Newyddionddio n	0.131** (0.020)		0.093** (0.026)	0.626**	1.155**	0.160	2.923	69.863	0.281	3.796
Pobol Y Cwm	0.185** (0.016)		0.099** (0.025)	0.127	0.562**	0.267	6.435	131.174**	0.026	4.192

Notes: Standard Errors in parentheses, * (**) denotes significance at 5% (1%). Sk and Ku are the coefficients of skewness and excess kurtosis, SC(1) is the Lagrange Multiplier test for correlation at lag 1, Q(5) and Q(70) are Ljung-Box portmanteau tests using 5 and 70 lags respectively, ARCH(1) and ARCH(5) are tests for autoregressive conditional heteroscedasticity at lags 1 and 5 respectively.

Table 3

ARFIMA Models: Weekly Data

Heno		d	s.e.	sk	ku	SC(1)	Q(5)	Q(30)	ARCH(1)	ARCH(5)
Heno	Mon	0.143*	0.056	0.204	0.187	0.436	2.451	32.229	12.756*	15.739*
	Tue	0.141**	0.054	0.620**	1.353**	0.294	1.737	23.338	0.406	1.686
	Wed	0.166**	0.043	0.382**	-0.055	0.199	4.551	30.732	0.021	4.632
	Thu	0.136*	0.054	0.287*	0.222	0.369	2.435	19.730	2.004	7.881
	Fri	0.137*	0.060	0.898	3.056	1.350	10.001	38.774	0.031	1.841
	Mean	0.145								
Newyddion ddion	Mon	0.144**	0.055	0.464**	0.573	0.104	6.166	35.224	0.013	18.548**
	Tue	0.125*	0.058	0.713**	0.954**	0.805	3.787	34.168	0.292	3.516
	Wed	0.089	0.065	0.742**	1.233**	0.054	7.590	30.619	0.018	1.045
	Thu	0.180**	0.043	1.240**	4.415**	0.034	3.183	36.589	0.112	4.137
	Fri	0.152*	0.069	0.422**	0.241	0.153	4.166	28.579	0.175	5.886
	Mean	0.138								
Pobol Y Cwm	Mon	0.228**	0.053	-0.099	0.844*	0.183	1.448	31.141	0.169	1.974
	Tue	0.179**	0.046	0.016	-0.058	1.009	5.042	27.488	0.746	1.831
	Wed	0.189**	0.043	0.112	-0.064	1.759	7.049	100.946**	0.501	3.887
	Thu	0.185**	0.069	0.492**	0.881**	0.004	2.820	24.919	0.306	8.874
	Fri	0.250**	0.047	-0.012	0.544	0.021	3.188	26.261	0.189	1.345
	Mean	0.206								

Notes: See Notes to Table 2

ⁱ Although the individual decision to watch or not watch a given programme - like the decision to support or not support a given political party - is binary in nature, we do not have an estimate of the proportion of the relevant population engaging in the activity. Should one, for instance, use the population of all viewers at a given time, or all potential viewers? Since our data is for programmes in Welsh, we could also use the population of Welsh speakers or of Welsh-speakers watching television at the time.

ⁱⁱ We are very grateful to S4C for making these figures available to us.

ⁱⁱⁱ There are, in total, 28 missing values for Heno, 8 for Newyddion and 20 for Pobol Y Cwm