TEMPORAL AGGREGATION AND THE POWER OF TESTS FOR A UNIT ROOT*

R. G. PIERSE
London Business School, London NW1 4SA, UK

A. J. SNELL
University of Edinburgh, Edinburgh EH8 9YL, UK

Final revision, September 1993

Abstract

The asymptotic local power of unit root tests with the same data span is shown to be independent of sampling frequency. A measure of the power trade-off between sampling frequency and time span for distinct alternatives is derived using an approximate slopes approach. Only small span increases are generally required to maintain power when reducing sampling frequency. Monte Carlo results support the asymptotic analysis for finite samples. An application is made to a consumption function for the UK. Cointegration of consumption and wealth is rejected with quarterly data but convincingly accepted with a longer span of annual data.

Key Words: Unit roots; Power; Sampling; Temporal aggregation
JEL Classification: C12; C15; C22

1 Introduction and summary

The effects of sampling frequency on estimation and inference in economic models have been explored in the literature in a number of different ways. Nijman and Palm (1990), for example, analyse the gain in forecasting efficiency obtained by using monthly instead of quarterly models to forecast monthly series and weigh this against the cost of collecting monthly National Accounts data. Lippi and Reichlin (1991) assess the effect of temporal aggregation on estimated measures of persistence in the context of trade cycle analysis while Hotta et al. (1993) look at problems arising from overlapping aggregation caused by commonly used smoothing procedures.

This paper focuses on the effect of temporal aggregation on the power of tests for a unit root in simple time series models, following the line of previous work by Shiller and Perron (1985) and Perron (1989, 1991). The topic of testing for the presence of a unit root in economic time series has received much attention in the literature following the influential paper by Nelson and Plosser (1982), and various test statistics have been proposed by Dickey and Fuller (1979, 1981), Phillips (1987a) and Phillips and Perron (1988) among others. Unit root tests also form the basis of the tests for cointegration between economic time series proposed by Engle and Granger (1987) and the power of these tests to reject the null hypothesis of a unit root (implying no cointegration) is of considerable importance. In this paper the relationship between power and sampling frequency is explored. This is an issue of direct practical relevance to the applied researcher who often has the option of choosing between data sets of different frequencies covering different time spans.[1] The macroeconomist, for example, typically has available either quarterly observations since the war or annual data since the beginning of the century. By contrast, in the study of financial markets, it is possible to collect some data on an almost continuous basis over a short period of time or use the published sources which cover a much longer time span but at weekly or monthly intervals.

*Published in *Journal of Econometrics* (1995), Vol. 65, pp. 333–345. The authors would like to thank Len Gill and Steve Satchell for helpful comments. Financial support from ESRC grant no. R000232912 is gratefully acknowledged.

[1] Harvey and Pierse (1984) show how data sets of different frequencies can sometimes be combined using the Kalman filter. However, this requires knowledge of the order of integration of the series which must be determined in some way, such as by a unit root test.

There is a widely held view in empirical work that long data spans are important for identifying mean reversion in slowly decaying processes (see, for example, Diebold *et al.* 1991)). This view is supported by the Monte Carlo results of Shiller and Perron (1985) and Perron (1989) who find that 'over a substantial range of parameter values ... [power depends more] on the span of the data rather than the number of observations' (*op. cit.* p.381). At the analytical level, Perron (1989) finds that with *fixed* alternatives, tests for a unit root are only consistent when the time span rises with the number of observations.

In Section 2 of the paper, the view that it is the data span rather than the number of observations that affects power in the context of testing for unit roots is expressed more forecfully than previously by showing that, when comparing models with the same data span, the asymptotic local power of a one-sided unit root test is *independent* of the frequency of sampling. In Section 3, we turn to lookat power against fixed alternatives using an approach due to Geweke (1981). Using the almost sure limit of the ratio of competing test statistics (comparing a model in 'basic' time units with a corresponding temporally aggregated model) as a measure of relative power, a theorem is stated that allows this ratio to be interpreted as the increase in span required to maintain power fixed at some arbitrary level when moving from the basic model to the aggregated model. This provides a measure of the trade-off between sampling frequency and time span when the interest is in power against a distinct rather than a local alternative. In most cases it is found that the increase in span that is required to keep power fixed is very small relative to the order of temporal aggregation. In order to help the investigator assess the absolute power of a unit root test given a particular data set, an approximation to the power function is proposed that is easy to compute. Section 4 presents the results of some Monte Carlo simulations for some simple models. These show that the asymptotic analytical results of Sections 2 and 3 perform well in finite samples. Finally, in Section 5 the results of the paper are applied to a model of consumer behaviour estimated by Molana (1991). Although Molana was unable to reject the null hypothesis that consumption and wealth are not cointegrated, using the short span of available quarterly data, when his model is reestimated using annual data over a longer span, the null hypothesis is convincingly rejected.

2 Asymptotic local power and temporal aggregation

Let y_t be a variable generated by the discrete first order process[2]

$$y_t = \rho y_{t-1} + u_t \quad t = 1, 2, \ldots, n_b \tag{1}$$

where n_b is the number of observations in 'basic' time units and u_t is an error which is assumed to follow a finite order stationary $ARMA(p, q)$ process whose largest AR root has modulus less than ρ.

Eq. (1) represents a model in 'basic' time units which is to be compared with a counterpart aggregated over a time interval of m (where m is a finite integer). The form of the aggregation will depend on whether y_t is a stock or a flow counterpart. Aggregating over a time interval of mh (where m is a finite integer) gives us the mth order temporally aggregated variable. In the former case the aggregated variable, denoted as y_t^* is simply y_t observed at the points $t = m, 2m, 3m$, and so on. In the case of a flow variable,

$$y_t^* = (1 + L + L^2 + \cdots + L^{m-1})y_t,$$

where L is the lag operator.

In either case, the aggregated form of model (1) is given by

$$y_t^* = \rho^* y_{t-1}^* + u_t^*, \quad t = m, 2m, \ldots, n_a m, \tag{2}$$

where u_t^* is a finite-order stationary $ARMA(p, q^*)$ process, $n_a = [n_b/m]$ is the number of temporally aggregated observations (where $[x]$ denotes the integer part of x), and

$$\rho^* = \rho^m. \tag{3}$$

Amemiya and Wu (1972) and Brewer (1973) show that, for the stock variable case, $q^* = [((p+1)(m-1)+q)/m]$, and for the flow variable case, $q^* = [((p+2)(m-1)+q)/m]$. As $m \to \infty$, then for the stock case, $q^* \to p+1$ for $q \geq p+1$, otherwise $q^* \to p$. IN the flow case, $q^* \to p+2$ for $q \geq p+2$, otherwise $q^* \to p+1$.

[2] It is possible to derive (1) from an underlying model formulated in continuous time by integrating over some time interval (see Bergstrom (1984) for details). Whether the underlying economic decision process being modelled is more appropriately viewed as discrete or continuous is an open question and does not affect the validity of the discrete representation (1).

Consider testing the unit root hypothesis

$$H_0 : \rho = 1 \quad \text{against} \quad H_1 : \rho = e^{-c/n_b}, \tag{4}$$

using some appropriate statistic, t. H_1 represents a one-sided local alternative to the null hypothesis for $c > 0$.[3]

Let the statistic calculated using n_b observations of the basic sampling frequency be denoted by t_b^n and that using n_a observations of the aggregated sampling frequency by t_a^n. Now consider a sequence of tests for increasing values of s, $s = 1, 2, \ldots$, where

$$n_a = s \quad \text{and} \quad n_b = sm. \tag{5}$$

Note that s is the time span measured in temporally aggregated time units. We now prove that the asymptotic local power of a class of unit root tests is independent of the degree of temporal aggregation.

Proposition 1. *Any test of (4) that is asymptotically independent of nuisance parameters under both H_0 and H_1 has a limiting distribution under both null and local alternative that is independent of the frequency of sampling, m.*

Proof: The null[4] and alternative for the temporally aggregated model are

$$H_0 : \rho^* = 1 \quad \text{against} \quad H_1 : \rho^* (= \rho^m) = e^{-c/n_a} . \tag{6}$$

t_b^n and t_a^n are similar tests so that under a common null they have the same limiting distribution. On the alternative hypotheses H_1 in (4) and (6), consider a sequence of tests for increasing values of s where n_a and n_b are defined y (5). Since by assumption the statistics t_b^n and t_a^n are asymptotically independent of the parameters of their respective error processes, it follows that as $s \to \infty$, t_b^n and t_a^n have the same limiting distribution and therefore the same asymptotic local power. **Q.E.D.**

We note that the conditions of Proposition 1 are weak and are met by most of the various different tests for unit roots[5] in the literature such as those proposed by Dickey and Fuller (1979, 1981), Phillips (1987a) and Phillips and Perron (1988). In practice, provided that data spans are long enough, Proposition 1 implies that the power of a unit root test against alternatives close to unity is going to be largely unaffected by the frequency of observation. The intuition behind the result is that the power loss from discarding $((m-1)/m)$ths of the data points is made up by the increased separation of H_0 from H_1 resulting from temporal aggregation.

3 Power against fixed alternatives

The previous section looked at power against a sequence of local alternatives. When we turn to consider fixed alternatives then in order to maintain the same power with a temporally aggregated model, the span of the data has to increase. Here an approach of Geweke (1981) is followed to derive a large-sample approximation to the required increase in span. First some definitions are needed.

Let t_b be a statistic based on the basic data $y_t, t = 1, 2h, \ldots, n_b$ and let t_a be a statistic based on the temporally aggregated data $y_t, t = m, 2m, \ldots, n_a m$. Let t_i^* be a critical value such that H_0 is rejected in favour of H_1 if $t_i > t_i^*$, $(i = a, b)$, and let $\beta(n_i, t^*)$, be the Type II error associated with a common critical value, t^* with sample size n_i. Finally, let $\overline{n}_i$ be the smallest integer for which $\beta(\overline{n}_i, t^*) < \bar{\beta}$, for some Type II error $\bar{\beta} \in [0, 1]$.

[3] In the terminology of Phillips (1987b), y_t is said to be near integrated under H_1.

[4] It is clear that, for *even* values of m, H_0 in (6) is also consistent with the hypothesis that $\rho = -1$. There is thus, in principle, an identification problem here. In practice, however, the possibility of negative unit roots in economic time series is virtually never entertained, so that it is reasonable to restrict consideration to $\rho \geq 0$.

[5] This still applies when the tests are being used on residuals from a cointegrating regression to test a null of no cointegration. A proof of this is available from the authors on request.

Assumption 1. *Let t_b and t_a be statistics with identical asymptotic distributions under H_0 and let them both be consistent under H_1, such that, for some appropriate power, $k > 0$,*

$$t_i/(n_i)^k \xrightarrow{\text{a.s.}} c_i < \infty \quad \textit{as} \quad n_i \to \infty.$$

Proposition 2. *Under the conditions of Assumption 1:*

$$\lim\{(\bar{n}_b/\bar{n}_a)^k\} = c_a/c_b\,.$$

This result is a straightforward extension of (Geweke, 1981, p. 1431) to which the reader is referred for a proof.

In Proposition 2 as t^* tends to infinity, the sample sizes n_b and n_a must also tend to infinity because the tests are consistent. Both tests have the same asymptotic distribution under H_0 so the ratio in the proposition may be interpreted as being evaluated at some constant power $1-\overline{\beta}$ for a common asymptotic test size of zero. The conditions for the proposition are weak and are satisfied by the commonly used tests.

To see how Proposition 2 may be used to assess the power-preserving span increase referred to above, take an example. Suppose that u_t in (1) is white noise. Let $\widehat{\rho}$ be the OLS estimate of ρ in (1) and $\widehat{\rho}^*$ the OLS estimate of ρ^* in (2) and consider the test statistics $\widehat{t}_b = n_b(\widehat{\rho} - 1)$ and $\widehat{t}_a = n_a(\widehat{\rho}^* - 1)$. For this case $k = 1$ and under the H_1 of $\rho < 1$, the almost sure limits of t_b/n_b and t_a/n_b are $\rho - 1$ and $\rho^m - 1$, respectively. The ratio of sample sizes (n_a to n_b) required to maintain power at the arbitrary level $1 - \overline{\beta}$ for a critical value t^* is, therefore, $(1-\rho)/(1-\rho^m)$ and the corresponding relative span is simply $m(1-\rho)/(1-\rho^m) = m/\sum i = 0^{m-1}\rho^i$. For example, if quarterly and annual data are being compared and if at quarterly frequency $\rho = 0.9$, then the proposition indicates that a 16.3% increase in span is required so that 40 quarterly observations should give the same power as 12 annual ones. If the Dickey-Fuller t-ratio is used, then the formula for the ratio of spans is $\{m(1-\rho)(1+\rho^m)/((1+\rho)(1-\rho^m))\}^{0.5}$ which gives a required increase in span for $\rho = 0.9$ of 0.7%.

More complicated cases are the Z_α and the ADF t-ratio tests when y_t is an $ARMA(1,1)$ flow variable given by

$$y_t = \rho y_{t-1} + \epsilon_t + \theta\epsilon_{t-1}.$$

In this case the process remains $ARMA(1,1)$ after temporal aggregation. For the Z_α test of Phillips (1987a), $k = 1$ and

$$\operatorname{plim}\left(\frac{t_b}{n_b}\right) \equiv \operatorname{plim}\left(\frac{Z_\alpha}{n_b}\right) = -\left\{1 - \rho - \frac{\theta(1-\rho^2)}{(1+\theta^2+2\theta\rho)}\right\}^2 \times \left\{\frac{1+\theta^2+2\theta\rho+\theta(1+\rho^2)}{(1+\theta^2+2\theta\rho)(1-\rho)}\right\}.$$

Replacing θ and ρ in this expression with θ^* and ρ^* from the aggregated model gives $\operatorname{plim}(t_a/n_a)$, and the power-preserving span ratio is simply $m\{\operatorname{plim}(t_b/n_b)/\operatorname{plim}(t_a/n_a)\}$. For $m = 2$, $\rho = 0.9$, and $\theta = 0.25$, the required span increase is 3.6%.

In the case of the ADF t-test, the MA error is approximated by the addition of lagged differences to the regression. For low values of θ, an ADF of order one may be sufficient to approximate the process. For this case $k = \frac{1}{2}$ and

$$\operatorname{plim}\left(\frac{t_b}{\sqrt{n_b}}\right) = \frac{-(1+\theta^2+\theta\rho)}{\sqrt{2(1+\theta+\theta\rho+\theta^2)/(1-\rho)}}.$$

Again, replacing θ and ρ with θ^* and ρ^* from the aggregated model gives a corresponding form for $\operatorname{plim}(t_a/\sqrt{n_a})$, and the power-preserving span ratio is $\sqrt{m}\{\operatorname{plim}(t_b/\sqrt{n_b})/\operatorname{plim}(t_a/\sqrt{n_a})\}$. For $m = 2$, $\rho = 0.9$, and $\theta = 0.25$, the formula predicts a required span increase of 3.3%. If $\theta = 0.6$, then this figure rises to 17.6%.

Proposition 2 provides a measure of 'efficiency' in terms of the relative power under H_1. This must be distinguished from the more usual notion of efficiency which is in terms of the relative variances of parameter estimates. The latter has been examined by Palm and Nijman (1994) who find that the increase in span required to achieve equal efficiency of parameter estimates for the AR(1) model with $m = 2$ is $2\rho^2/(1+\rho^2)$. This compares with our measure of power-preserving span for the $n(\widehat{\rho}-1)$ statistic of $2/(1+\rho)$ which is strictly larger. Our measure will of course vary according to the test statistic used.[6]

[6] In the case of estimation ML is asymptotically efficient, and so is a natural choice for an estimator. In the context of test against distinct alternatives there is no such natural candidate.

As well as having a measure of relative power of basic and aggregated samples under H_1, it is also interesting to gauge absolute power in each case. The power, P, of the unit root test of sample size n is given by

$$P(\gamma, n, t_n^*) = \int_{-\infty}^{t_n^*} f_n(x; \gamma, n)\, dx \tag{7}$$

where f_n is the pdf of the test statistic, γ is a vector of nuisance parameters and t_n^* is the critical value for sample size n. The pdf in (7) is in general unknown but a change of variables is usually possible such that

$$P(\gamma, n, t_n^*) = \int_{-\infty}^{g(t_n^*, \gamma, n)} f_n^*(x; \gamma, n)\, dx \tag{8}$$

where the sequence of functions $f_n^*(x; \gamma, n)$ converges to the asymptotic distribution $f(x)$ independent of γ and n as $n \to \infty$. An approximation to (8) is, therefore,

$$P(\gamma, n, t_n^*) \simeq \int_{-\infty}^{g(t_n^*, \gamma, n)} f(x)\, dx\,. \tag{9}$$

To illustrate, consider the Dickey-Fuller t-ratio test for the simplest case of (1) where u_t is white noise. In this case, $f(x)$ is the standard normal distribution function and

$$g(t_n^*, \gamma, n) = t_n^* - \frac{\sqrt{n}(\rho - 1)}{\sqrt{1-\rho^2}}. \tag{10}$$

In the same model, the test statistic based on $n(\hat{\rho} - 1)$ also has an asymptotic standard nortmal distribution, but here

$$g(t_n^*, \gamma, n) = \frac{t_n^*}{\sqrt{n(1-\rho^2)}} + \frac{\sqrt{n}(1-\rho)}{\sqrt{1-\rho^2}}. \tag{11}$$

Finally, consider the Augmented Dickey-Fuller t-ratio test for (1), where u_t is now ARMA(1,0) so that y_t is ARMA(2,0). Again $f(x)$ is standard normal with

$$g(t_n^*, \gamma, n) = t_n^* + \frac{\sqrt{n}(1-\rho_1-\rho_2)\sqrt{1+\rho_1-\rho_2}}{\sqrt{2\{(1-\rho_1^2-\rho_2^2)(1-\rho_2) - 2\rho_1^2\rho_2\}}}. \tag{12}$$

[The critical values t_n^* for all these tests are tabulated in (Fuller, 1976, Tables 8.5.1 and 8.5.2).] It should be clear that both the relative power measure in Proposition 2 and the approximation to the power function in (9) can be computed and in most practical circumstances given values for γ and ρ.

4 Monte Carlo simulation

Both Propositions 1 and 2 and the approximations in (9) of Section 3 are based on asymptotic results. To assess how useful these are in finite samples, Monte Carlo simulations were conducted for a number of simple models for y_t and illustrating the use of different unit root tests. Simulations are reported for three models of the y_t process:[7] the ARMA(1,0) stock variable case and two ARMA(1,1) flow models, one with a low value of the moving average parameter ($\theta = 0.25$) and one with a high value ($\theta = 0.6$).

For each model a different test statistic is illustrated: the Dickey-Fuller $n(\rho-1)$ test for model 1, the Phillips (1987a) Z_α test for model 2, and the ADF t-ratio test for model 3. the ADF t-ratio test was computed following the procedure recommended by Said and Dickey (1984), where the number of lagged differences included in the regression is increased with the number of observations, n, but at a slower rate [a rate of order $O(n^{1/4})$ was used]. The Z_α test was implemented using a Parzan lag window[8] for the weights on the correction factor, with the lag length increasing with n, again at rate $O(n^{1/4})$.

The results, based on 5000 replications, are presented in Figs. 1–3, where power curves are graphed for values of $m = 1, 4, 12$, where $\rho = 0.9$ in the 'basic ' ($m = 1$) model. These curves illustrate the gain in power of moving from quarterly ($m = 1$) to annual ($m = 4$) or from monthly ($m = 1$) to annual ($m = 12$) data. the dotted lines in the figures plot the power for $m = 2$ against the number of observations scaled by the power-preserving span ratio for the test statistic as derived in Section 3. If Proposition 2 were to

[7] These were selected as being representative of a large number of simulations that were run, all programmed using the GAUSS econometric programming language.

[8] This was the choice in the simulations reported by Phillips and Perron (1988).

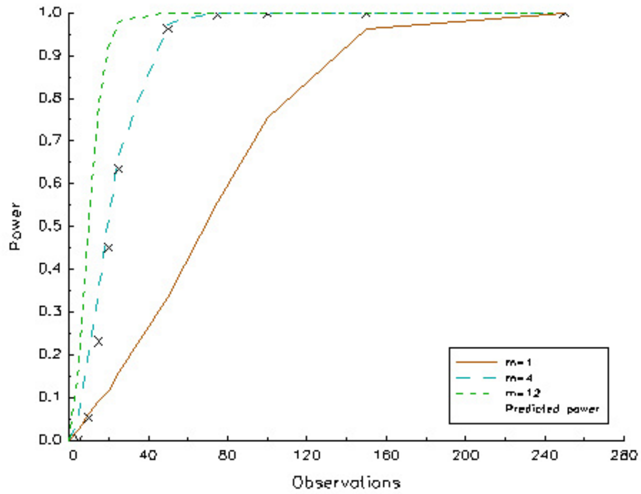

Figure 1: Power functions: ARMA(1,0) stock case, $n(\rho - 1)$ test, $\rho = 0.9$.

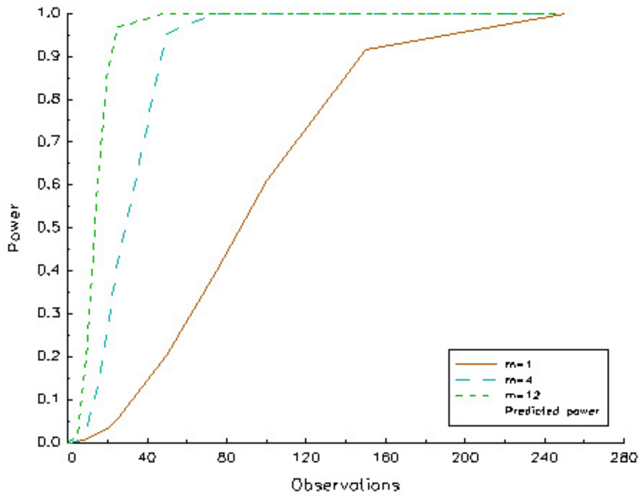

Figure 2: Power functions: ARMA(1,1) flow case, Z_α test, $\rho = 0.9$, $\theta = 0.25$.

hold *exactly* in finite samples, then this dotted line would lie exactly on top of the $m = 1$ curve. The closeness of the dotted lines to the $m = 1$ curves in all three figures thus shows how remarkably well Proposition 2 holds for all our models, even in very small samples.

Finally, the power approximation to the power function based on (9) was computed for the $m = 4$ curves for the Dickey-Fuller $n(\rho - 1)$ test and the ADF t-ratio test, and is plotted by the crosses in Figs. 1 and 3. For the DF test in Fig. 1 it can be seen that the approximation, using Eq. (11), while it slightly underestimates power at very small sample sizes, holds very well at all other sample sizes. For the ADF t-ratio test in Fig. 3 the approximations were computed as in Section 3, using (9) and the formula

$$g(t_n^*, \gamma, n) = t_n^* + \sqrt{n} z_n,$$

where z_n is the plim of the estimated AR($O(n^{1/4})$) process under the truer ARMA(1,1) model. It can be seen from the figure that the approximation overestimates power at small sample sizes and underestimates it at large sample sizes. However, taken overall the approximation, while obviously not as good as that for the simpler model, is still not bad, and can serve as a rough guide to finite-sample power.

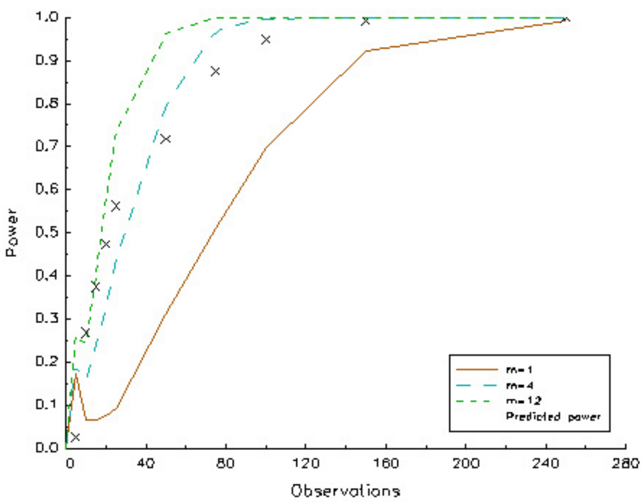

Figure 3: Power functions: ARMA(1,1) flow case, ADF test. $\rho = 0.9$, $\theta = 0.6$.

5 The cointegration of consumption and wealth

In a recent paper, Molana (1991) extends the intertemporal model of consumer behaviour of Hall (1978) to allow for credit constraints and derives a specification relating (nondurable) consumption to the stock of wealth. However, on the basis of an augmented Dickey-Fuller (ADF) test using quarterly observations from 1966(4) to 1981(4), he was unable to reject the null hypothesis that consumption and wealth are not cointegrated. Although quarterly data on wealth are only available from 1966(4), annual data go back to 1957. We repeated Molana's exercise[9] using the 25 annual data points between 1957 and 1981 and our results are compared to his in Table 1.

Table 1: Tests for the cointegration of non-durable consumption and wealth

Span	Frequency	$CRDW$	ADF	Z_t	$CRDW^c$	ADF^c	Z_t^c
1966-1981	Quarterly	0.83	-1.22	—	0.11	-3.18	-3.15
1966-1981	Annual	0.63	-2.53	-1.99	0.13	-3.47	-3.34
1957-1981	Annual	0.81	-3.56	-3.02	0.17	-3.46	-3.13

$CRDW$ is the cointegrating regression Durbin-Watson statistic, ADF the cointegrating regression augmented Dickey-Fuller statistic, Z_t the Phillips Z_t test statistic and superscript ^{c} denotes the appropriate 5% critical values (obtained by simulation).

The cointegrating regression was $\ln C_t = \alpha + \beta \ln W_t + u_t$ where C_t is real nondurable consumption and W_t is real net household wealth (seasonals were included in the quarterly model). $CRDW$, ADF, and Z_t are the cointegrating regression Durbin-Watson, augmented Dickey-Fuller and Phillips Z_t test statistics respectively, and the superscript ^{c} denotes the appropriate 5% critical value. Critical values were estimated by numerical simulation on the assumption that $\ln C$ and $\ln W$ follow independent random walks with drift given by the average quarterly or annual growth rate of the respective series.

The quarterly roots in the ADF auxiliary regression were not reported in Molana (1991). In the annual data however, where an AR(2) auxiliary regression proved adequate, the largest of the two roots was found to be 0.552. An estimate of the largest root in the quarterly cointegrating regression is therefore given by $0.552^{\frac{1}{4}} = 0.862$ (see for example Amemiya and Wu (1972)). Propositions 1 and 2 indicate that with a root so close to unity, we should suffer little loss in power in moving from quarterly to annual data with the same time span. The first two rows of Table 1 show that despite relatively high $CRDW$ statistics, the ADF and Z_t tests based on the 1966-1981 span and using either frequency of data observation, fail to reject noncointegration by a large margin. Using the formula for the power

[9] In Molana's work the use of quarterly data was necessary to explore dynamics etc. However, this does not preclude the use of annual data to establish cointegration as the first stage of a two-step procedure.

approximation derived in (12) in Section 3, (consistent) estimates of ρ_1 and ρ_2 from the annual data, and setting $T = 15$ gives a power estimate of 0.08.[10] In the light of this, Molana's own failure to reject is unsurprising. The third row shows the gain from increasing the span; the ADF now rejects noncointegration and although Z_t remains insignificant, it only does so by the smallest of margins. Finally, updating the span to 1987 gives highly significant $CRDW$, ADF and Z_t values of 0.78, −4.09 and −3.31, respectively. We may conclude therefore that consumption and wealth in the UK are indeed cointegrated and that Molana's failure to find a cointegrating relationship can be explained by the low power of unit root tests when applied to data of a short time span.

References

Amemiya, T. and R. Y. Wu (1972). The effect of aggregation on prediction in the autoregressive model. *Journal of the American Statistical Association 67*, 628–632.

Bergstrom, A. R. (1984). Continuous time stochastic models and issues of aggregation over time. In Z. Griliches and M. Intriligator (Eds.), *Handbook of Econometrics*, Chapter 20. Amsterdam: North-Holland.

Brewer, K. R. W. (1973). Some consequences of temporal aggregation and systematic sampling for ARMA and ARMAX models. *Journal of Econometrics 1*, 133–154.

Dickey, D. A. and W. A. Fuller (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association 74*, 427–431.

Dickey, D. A. and W. A. Fuller (1981). Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica 49*, 1057–1073.

Engle, R. F. and C. W. J. Granger (1987). Co-integration and error correction: representation, estimation and testing. *Econometrica 55*, 251–276.

Fuller, W. (1976). *lntroductlon to Statlstical Time Series*. New York, NY: Wiley.

Geweke, J. (1981). The approximate slopes of econometric tests. *Econometrica 49*, 1427–1442.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *Journal of Political Economy 86*, 971–987.

Harvey, A. C. and R. G. Pierse (1984). Estimating missing observations in economic time series. *Journal of the American Statistical Association 79*, 125–131.

Hotta, L. K., R. G. Pierse, and P. L. V. Pereira (1993). *The effect of overlapping aggregation in time series models*. London: London Business School. Unpublished paper.

Lippi, M. and L. Reichlin (1991). Trend-cycle decompositions and measures of persistence; does time aggregation matter? *Economic Journal 101*, 314–323.

Molana, H. (1991). The time series consumption function: error correction, random walk and the steady-state. *Economic Journal 101*, 382–403.

Nelson, C. R. and G. J. Plosser (1982). Trends and random walks in macroeconomic time series. *Journal of Monetary Economics 10*, 139–162.

Nijman, T. E. and F. C. Palm (1990). Predictive accuracy gain from disaggregate sampling in ARIMA models. *Journal of Business and Economic Statistics 8*, 405–415.

Palm, F. C. and T. E. Nijman (1994). Missing observations in the dynamic regression model. *Econometrica 52*, 1415–1435.

Perron, P. (1989). Testing for a random walk: a simulation experiment of power when the sampling interval is varied. In B. Raj (Ed.), *Advances in Econometrics and Modelling*. Dordrecht: Kluwer Academic Publishers.

[10] The corresponding figure for the 1957-1981 span is 0.26, a threefold increase.

Perron, P. (1991). Test consistency with varying sampling frequency. *Econometric Theory 7*, 341–368.

Phillips, P. C. B. (1987a). Time series regression with unit roots. *Econometrica 55*, 277–302.

Phillips, P. C. B. (1987b). Towards a unified asymptotic theory for autoregression. *Biometrika 74*, 535–547.

Phillips, P. C. B. and P. Perron (1988). Testing for a unit root in time series regression. *Biometrika 75*, 335–346.

Said, S. E. and D. A. Dickey (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika 71*, 599–607.

Shiller, R. J. and P. Perron (1985). Testing the random walk hypothesis: power versus frequency of observation. *Economics Letters 18*, 381–386.