CONSUMPTION AND FRACTIONAL DIFFERENCING: OLD AND NEW ANOMALIES

by

Joseph G. Haubrich

(20-89)

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

CONSUMPTION AND FRACTIONAL DIFFERENCING: OLD AND NEW ANOMALIES

BY

Joseph G. Haubrich

Department of Finance

Wharton School

University of Pennsylvania

April, 1989

Thanks to Andrew Abel, Andrew Lo, and Joseph Ritter for stimulating discussions.

13.21.3

I. Introduction

Theories of consumption have evolved hand in hand with theories of income, from the static Keynesian MPC to the unit root anomalies of Deaton and Campbell. Recently, "long memory" stochastic processes, compromises between a random walk and a finite order ARMA model, have become popular time series models for macroeconomic data (Haubrich and Lo, 1988, Diebold and Rudebusch 1988, Sowell 1987). It seems natural for consumption studies to consider this new class--the fractionally differenced processes.

This new approach has profound implications for the consumption function. On the theoretical level a fractional process for income predicts consumption exhibiting both Flavin's (1981) excess sensitivity and Campbell and Deaton's (1987) excess smoothness. On the empirical level, the existence of long memory in consumption--both in particular components and across different countries--provides a new and intriguing set of puzzles.

This paper first calculates the stochastic properties of consumption when income follows a fractional stochastic process, and shows how this may explain both the excess sensitivity results and the excess smoothness results. It then uses a recently developed improvement of the Rescaled Range Statistic and finds long term memory in consumption. The remaining sections undertake Monte Carlo simulations to assess the finite sample size and power of the test, conduct cross country comparisons, and discuss possible explanations.

II. Fractional Methods

Intuition suggests that differencing a time series roughens it up while summing it smooths it down. A fractional difference between 0 and 1 can be described as a filter that roughens a series less than a first difference: one that yields a series rougher than a random walk but smoother than white noise. This is apparent from the infinite-order moving average representation. Let X_t follow

$$(1 - L)^d X_t = \varepsilon_t \tag{1}$$

where ε_t is white noise, d is the degree of differencing, and L is the lag operator. If d = 0 then X_t is white noise, and if d = 1, X_t is a random walk. However, as Granger and Joyeux (1980) and Hosking (1981) show, d need not be an integer. The binomial theorem provides the relation:

$$(1 - L)^d = \sum_{k=0}^{\infty} (-1)^k \binom{d}{k} L^k \tag{2}$$

with the binomial coefficient $\binom{d}{k}$ defined as:

$$\binom{d}{k} \equiv \frac{d(d - 1)(d - 2)\cdots(d - k + 1)}{k!} \tag{3}$$

for and real d and non-negative integer k. Using this, the AR form of X_t follows:

$$A(L)X_t = \sum_{k=0}^{\infty} A_k L^k X_t = \sum_{k=0}^{\infty} A_k X_{t-k} = \varepsilon_t \tag{4}$$

with the AR coefficient expressed compactly in terms of the gamma function

$$A_k = (-1)^k \binom{d}{k} = \frac{\Gamma(k - d)}{\Gamma(-d)\Gamma(k + 1)} . \tag{5}$$

Manipulating (5) yields the corresponding MA representation of X_t:

$$X_t = (1 - L)^{-d} \varepsilon_t = B(L)\varepsilon_t \qquad B_k = \frac{\Gamma(k + d)}{\Gamma(d)\Gamma(k + 1)} . \tag{6}$$

The time series properties of X_t depend crucially on the difference parameter d. For example, when d is less than 1/2, X_t is stationary; when d is greater than -1/2, X_t is invertible (Granger and Joyeux 1980, Hosking 1981). Likewise the autocorrelation properties of X_t depend on the parameter d. The MA coefficients B_k tell the effect of a shock k periods ahead, and indicate the extent to which current levels depend on past values. Using Stirling's approximation, we know:

$$B_k \approx \frac{k^{d-1}}{\Gamma(d)} . \tag{7}$$

Comparing this with the decay of an AR(1) process highlights the central "long memory" feature of fractional processes: they decay hyperbolically, at rate k^{1-d}, rather than at the exponential rate ρ^k of an AR(1).

III. Fractional Differencing and the Theory of Consumption

Fractional processes can seriously change how we think about the consumption function. Long term dependence in Y_t, income, will affect the pattern of consumption in ways that mimic several of the anomalies discovered. Likewise, finding fractional differencing in consumption would have profound implications: by standard permanent income theory, it should not be there.

One well-known set of paradoxes concerns the variability of consumption relative to income. When income is represented as a trend plus ARMA process, consumption looks too volatile, it displays "excess sensitivity" (Flavin, 1981); when income is represented as a difference stationary process, consumption looks too smooth (Campbell and Deaton, 1987). Letting income follow a fractional process has a potential to resolve this paradox. With a fractional income process current shocks will persist much longer than with any finite ARMA process. The shock thus persists longer than the conventional trend plus stationary representation admits. Consumers, following the true fractional process, look irrational to the econometrician. Another econometrician, estimating an income process with a unit root, will make the opposite conclusion, assuming that a shock persists forever, whereas the true (fractional) shock dies out, albeit slowly.

To make the point more formally, consider the following example. Assume the standard certainty equivalence framework (e.g., quadratic utility, see Flavin 1981, Zeldes 1988) with a real interest rate of r = 0.01. We will compare the variance of consumption under three different income processes: AR(1), a fractional, and a random walk. Under any of these assumptions, we can describe the path of consumption using a formula from Flavin (1981), shown to hold for stationary and non-stationary process by Hansen and Sargent (1981).

$$\Delta C_t = \frac{r}{1+r} \frac{[1 + \Sigma(1+r)^{-k}\theta_k]}{[1 - \Sigma(1+r)^{-k}\phi_k]} \varepsilon_t . \tag{8}$$

Where C_t is consumption, θ_k are the MA coefficient of Y_t, ϕ are the AR coefficients, and Δ is the difference operator, $\Delta = (1 - L)$. From this, we can calculate the variance of ΔC_t

$$\text{var}(\Delta C_t) = \{\frac{r}{1+r} \frac{[1 + \Sigma(1+r)^{-k}\theta_k]}{[1 - \Sigma(1-r)^{-k}\phi_k]}\}^2 \sigma_\varepsilon^2 . \tag{9}$$

We now use equation 9 to compare the variance of consumption under three different specifications for income. Consider an AR(1) with $\rho = 0.9$, a fractional process with d = 1/2, and a random walk. Furthermore, normalize σ_ε^2 to 1, and choose a real interest rate of 0.01. For the AR(1),

$$\text{var}(\Delta C_t) = [\frac{r}{1+r-\rho}]^2 = [\frac{0.9}{1+0.01-0.9}]^2 = 0.00826 \tag{10}$$

The corresponding expression for the fractional process does not have a closed form solution, but since all terms are positive we can take a lower bound by summing up the first 500 terms. Using this and equation 7, we find the variance when income follows a fractional process:

$$\text{var}(\Delta C_t) = \left[\frac{r(9.186)}{1 + r}\right]^2 = 0.01017 \, . \tag{11}$$

Finally, expressing a random walk as an AR(1) with $\rho = 1.0$, we have

$$\left[\frac{r}{(1 + r)\left(1 - \frac{1}{1 + r}\right)}\right]^2 = \left[\frac{r}{(1 + r)} \frac{(1 + r)}{r}\right]^2 = 1 \, . \tag{12}$$

Equations (10)-(12) keep the variance of innovations constant, resulting in different variances for income. However, we observe the (sample) variance of Δy. A more realistic approach keeps it constant and requires differing innovation variances. Fortunately, this produces similar results.

Take as a base the random walk case, setting σ_ε^2, the variance of innovations, to 1, which also implies that the variance of income (Δy) is 1. For a (detrended) AR(1), $\text{var}(\Delta y) = \frac{2}{1 + \rho} \sigma_\varepsilon^2$. If $\rho = 0.9$, normalizing $\text{var}(\Delta y)$ to 1 requires $\sigma_\varepsilon^2 = \frac{1}{1.053}$. Likewise for a fractional process of $\Delta y = (1 - L)^{1-d}\varepsilon_t$, $\text{var}(\Delta y) = = \frac{\Gamma(1 - 2d)}{\{\Gamma(2 - d)\}^2} \sigma_\varepsilon^2$. For $d = 3/4$ (the variance is infinite for $d = 1/2$) this implies a σ_ε^2 of 0.5113814.

Plugging these values into the consumption equations (10)-(12) yields a variance of Δc of 1 for a random walk, 0.007847 for an AR(1) with $\rho = 0.9$, and 0.288196 for the fractional process with $d = 3/4$. Fractional income still produces consumption that looks too variable for an AR(1) and too smooth for a random walk. Note that in the fractional case, the variance of income innovations exceeds the variance of consumption even with non-stationary income. This clarifies the Campbell and Deaton claim that consumption variance should be higher; a unit root, not just non-stationarity, is required for consumption variance to exceed income innovation variance.

Thus, fractional processes hold promise of explaining both excess sensitivity and excess smoothness results by the intuitively appealing approach that income is in fact more persistent than ARMA suggest, but not as long lasting as unit root models claim. The caveat "holds promise" appears because current estimation methods cannot accurately estimate the fractional process in income. Furthermore, explicit tests (Haubrich and Lo, 1988) indicate that income does not follow a fractional process. Still, several recent innovations (Sowell 1987, Geweke and Porter-Hudak 1983) suggest that it is only a matter of time before practical estimation is possible (Diebold and Rudebusch, 1988, Sowell, 1987).[1]

Fractional differencing also introduces a cautionary note in comparing the variability of consumption and income. The most commonly used mechanism generates a fractional process via aggregation of many AR(1) processes (Bernamont 1937, Granger 1980, Haubrich and Lo 1988). Even if individuals conform to the permanent income hypothesis, the aggregate values of consumption and income may not: the representative agent approach breaks down. A simple example will show that consumption appears too variable for the observed aggregate income process.

If individuals face an AR(1) income process

$$Y_{t+1} = \rho Y_t + \varepsilon_t \tag{13}$$

---

[1]For an estimate of income with a view to explaining consumption anomalies in the spirit of this section, see the interesting (independent) work of Diebold and Rudebusch 1989.

then by (8) consumption is

$$\Delta C_t = \frac{r}{1 + r - \rho} \varepsilon_t . \qquad (14)$$

Now we invoke the aggregation theorems to generate a fractional difference in aggregate income. If ρ^2 is distributed across individuals according to Beta distribution B(p, q), and each individual's shock ε_t is independent of others, aggregate income Y_t will follow a fractional process of order 1 - (q/2). Choosing p = q = 1 picks out a particularly simple Beta distribution, the uniform. Aggregating the many individual consumptions c_t together produces the aggregate C_t:

$$\Delta C_t = \int_0^1 \frac{r}{1 + r - \rho} 2\rho d\rho . \qquad (15)$$

Integrating (15) and collecting terms results in the variance of ΔC_t when we aggregate up from individuals who face an AR(1) income stream:

$$\text{var}(\Delta C_t) = 2r + 2r^2 \ln r - 2r^2 \ln(1 + r) . \qquad (16)$$

The uniform distribution makes it easy to calculate the variance of consumption for an agent facing the aggregate income stream. Since d = 1/2, we use equation (9), again with an interest rate of 0.01, to compute $\text{var}(\Delta C_t)$. Comparing (16) with (10) shows that consumption will look excessively variable relative to income since 0.018 > 0.008. On an abstract level this result is well known: aggregation can introduce serious biases. The present context makes the problem particularly acute because the most plausible method of generating fractional processes is by aggregating. If income follows an ARMA (2, 2), one may perhaps pretend that aggregation problems do not arise. When income follows a fractional process, it signals that aggregation is playing a major role.

IV. Testing for Fractional Differencing

Permanent income theory predicts that consumption should be a random walk (Hall 1978). Other theories, involving liquidity constraints or habit formation, often imply otherwise. In this section we search for a particular alternative, for fractional differencing, using the methodology developed in Haubrich and Lo (1988) and in Lo (1988). Finding a fractional process in consumption means rejecting the standard random walk form of the permanent income hypothesis. It may be consistent with other theories, say a liquidity constrained consumer forced to consume his long term dependent income.

The test statistic we use is the modified rescaled range or R/S statistic, based on a statistic originally developed by Hurst (1951) and popularized by Mandelbrot (1975). The statistic tests whether a process X_t shows long term dependence or not. More formally, we express the null hypothesis for a process defined as

$$X_t = \mu + \varepsilon_t \tag{17}$$

where μ is an arbitrary but fixed constant. As the null hypothesis H, assume that the disturbances $\{\varepsilon_t\}$ satisfy the conditions

(C1) $$E[\varepsilon_t] = 0 \text{ for all } t \text{ .}$$

(C2) $$\sup_t E\left[|\varepsilon_t|^\beta\right] < \infty \text{ for some } \beta > 2 \text{ .}$$

(C3) $$\sigma^2 = \lim_{n\to\infty} E\left[\frac{1}{n}\left(\sum_{j=1}^{n} \varepsilon_j\right)^2\right] \text{ exists and } \sigma^2 > 0 \text{ .}$$

(C4) $\{\varepsilon_t\}$ is strong-mixing with mixing coefficients α_k that satisfy:

$$\sum_{k=1}^{\infty} \alpha_k^{1-\frac{2}{\beta}} < \infty \text{ .}$$

Conditions (C2)-(C4) allow dependence and heteroskedasticity, but prevent them from being too large. Thus, short term dependent processes, such as finite order ARMA models, are included in the null hypothesis, as are models with conditional heteroskedasticity. Unlike the statistic used by Mandelbrot, the modified R/S used here is robust to short term dependence. A deeper discussion of these conditions appears in Phillips (1987) and Haubrich and Lo (1988), which interested readers should consult.

To construct the modified R/S statistic, take a sample X_1, $X_2 \ldots X_n$, with sample mean $\bar{X}_n$, choose q lags, and calculate:

$$Q_n \equiv \frac{1}{\hat{\sigma}_n(q)} \left[\underset{1 \le k \le n}{\text{Max}} \sum_{j=1}^{k} (X_j - \bar{X}_n) - \underset{1 \le k \le n}{\text{Min}} \sum_{j=1}^{k} (X_j - \bar{X}_n) \right]$$

where

$$\hat{\sigma}_n^2(q) \equiv \frac{1}{n} \sum_{j=1}^{n} (X_j - \bar{X}_n)^2 + \frac{2}{n} \sum_{j=1}^{q} \omega_j(q) \left\{ \sum_{i=j+1}^{n} (X_i - \bar{X}_n)(X_{i-j} - \bar{X}_n) \right\} \qquad (18)$$

$$= \hat{\sigma}_x^2 + 2 \sum_{j=1}^{q} w_j(q) \hat{\gamma}_j \qquad w_j(q) \equiv 1 - \frac{j}{q+1} \qquad q < n \ .$$

Intuitively, the numerator in (18) measures the memory in the process via the partial sums. White noise does not stay long above the mean: positive values are soon offset by negative values. A random walk will stay above or below 0 for a long time, and the partial sums (positive or negative) will grow quickly, making the range large. Fractional processes fall in between. Mandelbrot (1972) refers to this as the "Joseph Effect" of seven fat and seven lean years. The denominator normalizes not only by the variance, but by a weighted average of autocovariances.[2] This innovation over Hurst's R/S provides the robustness to short term dependence.

The partial sums of white noise constitute a random walk, so $Q_n(q)$ grows without bound as n increases. A further normalization makes the statistic easier to work with and interpret:

$$V_n(q) = Q_n(q)/\sqrt{(n)} \ . \tag{19}$$

Haubrich and Lo derive the asymptomatic distribution of V, calculating the mean and standard deviation of approximately 1.25 and 0.27. Table 1 presents fractiles of the distribution of V and confidence intervals about the mean. Figure 1 plots the distribution and density. Note that the distribution is skewed, with most of its mass between 3/4 and 2.

We apply this test to several consumption series: Quarterly real total consumption, durables, non-durables, and services. The data range from 1959:1 to 1988:4 and are taken from the NIPA files on WEFA. We calculate $V_n(q)$ for the first difference of the log of these variables. The results are reported in table 2. The first column of numerical entries are the estimates of the classical V statistic which is not corrected for short term dependence. The next columns contain estimates of the modified $V_n(q)$ for lag values of 1, 2, 3 and 5.

Table 2 shows that the evidence for long term dependence occurs for service expenditures. Both the classical statistic and the modified statistic reject the null hypothesis of only short term correlation. Increasing the number of lags decreases both the power and the size of the test, so it is also not surprising that the test does not reject at three lags or above.

---

[2]These weights define the Bartlett window. Newey and West, 1987, enumerate the advantages of this specification.

Still, this finding of fractional differencing in consumption services raises several interesting anomalies. Why is there long term dependence in services and not in other consumption data? How can long term dependence occur in consumption without occurring in income? What mechanism causes the fractional exponent?

The distribution reported above is asymptotic; the small sample properties of the test may differ considerably. Simulation methods provide a standard assessment of the actual size and power in small samples. Table 3 reports the results of size and power simulations, each panel using 10,000 independently generated series. Against a fractional process with d = 1/3, the test shows considerable power, even though this drops off rapidly as lag length increases. The size calculations are for three processes included in the null. Panel B reports the results when the null is an IID process, which is the null from permanent income theory. Panels C and D report size for AR(1) nulls with autoregressive parameters of 0.5 and -0.25. (Components show both positive and negative first order autocorrelation.) These give an idea of how well the test distinguishes long term dependence from short term dependence. For two or fewer lags, actual size of the test appears larger than the theoretical 5%. Even so, services show long term dependence at a conventional level of 10%. Like power, size decreases quickly, and beyond 2 lags it is not surprising that the test does not reject the null.

V. Cross Country Comparisons

The U.S. results are suggestive but not fully persuasive. Imperfect size and power mean that any test sometimes accepts and sometimes rejects when it shouldn't. A standard solution is more data: in this case, other countries. I check for long term dependence in four data sets from three countries:

annual data on income and consumption in the U.K. from 1870 to 1965, and quarterly consumption and its components in the U.K., Canada and France.

The annual data, index numbers from Feinstein (1972), allow a truly long term look at consumption. Long term dependence seems more naturally a hypothesis about a near century of data than about post-war sub-samples. Most national accounts divided consumption into its components only recently. Only Canada, the U.K., and France have component series long enough for a test with reasonable size and power. Checking for long term dependence in less than twenty years of data seems perverse.

Tables 4, 5 and 6 show the results of applying the modified Rescaled Range Test to data from the United Kingdom, Canada, and France. Results for the historical British series appear in the first two rows of Table 4. The other series start in 1959 and 1963 and run through 1988. The Canadian data starts in 1960, and the French in 1967, though some components start later.

In only one case, France, do GDP or total consumption show long term dependence. The U.S., U.K., and Canada show no such dependence. Long term dependence appears in total consumption in the one place it also appears in GDP. This suggests a general lack of memory in GDP and total consumption, but it also points out strong differences across countries. Conventional theories do not predict different behavior for France.

One puzzle arises because long term memory frequently shows up in consumption components, but not in output. One plausible story, after all, is that consumption closely follows output, either because of liquidity constraints or because of general equilibrium considerations. This has no plausibility when services are fractionally differenced and GDP is not.

Two related puzzles arise in looking at the long term dependence properties of consumption components. Why services, and why not durables?

Service expenditures emerge as significant in three of the four countries, strongly so in two. Undoubtedly some services, such as divorces and appendectomies, have permanent features, but it seems surprising that services should consistently show more evidence of long term dependence when other components do not. The most likely component was probably durable goods, which both theoretically and empirically exhibits strong serial correlation. It shows up significant only once, using the classical Rescaled Range, which does not control for short term serial correlation. This could easily reflect the MA(1) component expected from Mankiw (1982).

No patterns emerge from the other components. Non-durables are significant only once, in Canada. Certainly some variation arises from the size and power of the test, but it also suggests that the time series properties of consumption and its components differ markedly across countries. Future work ought to exploit those differences, and compare stochastic structures across countries.

The results of the cross country comparison should give us additional confidence in the test. It could fail either by having no power, and never finding any fractionally differenced series, or by being too sensitive to short-term patterns, and rejecting everywhere. The test finds some long memory, in a rather interesting pattern; it neither always rejects nor always accepts the null. The strong pattern of no long memory in output, total consumption, or durables seems particularly persuasive. Mistaking short term for long term persistence is the greatest weakness of the R/S test, especially its classical variation. It is too sensitive to deviations from white noise that are officially part of the null. Yet it does not reject the null for output and most classes of consumption components.

The simulation results confirm this. With no lags, the test has both good power and a large size. It should find long term dependence, if anything, too often. With more lags, the size and power of the test decrease rapidly. When a weak test with a small size consistently rejects, it's a good sign something is there.

VI. Potential Explanations

Finding that services consumption displays long-term dependence presents a new anomaly. Services show such dependence while durables, non-durables and output do not. This does more than reject the random walk hypothesis; it presents a new stochastic structure needing explanation. It also rules out some explanations, such as postulating a fixed marginal propensity to consume out of a fractionally differenced income stream. The simulations and patterns across countries rule out small sample biases as an explanation.

Creating long term dependence via aggregation (mentioned above in section III) presents only a very partial answer. Thus, consider the standard Euler equation for the consumption choice problem

$$E[u'(c_{t+1})] = \frac{\beta}{R} u'(c_t) \tag{19}$$

where u is the utility function, and β and R denote the discount factor and the interest factor (one plus the rates). For quadratic utility, this reduces to an AR(1) consumption process,

$$c_{t+1} = \frac{\beta}{R} c_t + \left(\frac{1}{b}\right)\left(1 - \frac{\beta}{R}\right) - \frac{1}{b} \xi_{t+1} \ . \tag{20}$$

b is a constant from the utility function.

The variance of the error term is in general not independent of β, so we cannot immediately apply the Granger theorem. A simple trick will surmount the difficulty. Let a draw from nature give the individual not only a value of $(\frac{\beta}{R})$ but also the mean and variance of his income shock process. This can be chosen to cancel the connection between the shock and β arising from the maximization problem. Invoking the aggregation result (Granger, 1980) yields that if $(\frac{\beta}{R})^2$ is distributed as a $\beta(p, q)$ distribution, consumption will be a fractionally differenced process with order $d = 1 - \frac{q}{2}$. For differenced consumption, this produces a $d < 0$, which explains only the least interesting empirical finding, for only the U.K. durables show any significant evidence of $d < 0$. The more striking evidence of long term dependence indicates $d > 0$.

Different preferences hold some possibility of inducing long memory in consumption, perhaps because knowing so little we can rule little out. If not producing long memory directly, it may introduce a process that works via aggregation. Habit formation is another possibility. Probably a representative agent model would not produce a fractional series; the standard Ryder-Heal (1973) preferences give the past a geometrically declining weight, which does not fit the slower hyperbolic decay of a fractional process. Again, however, it may make aggregation possible. Providing one of the few closed form solutions, Sundaresan (1989) derives optimal consumption when utility depends on last period's consumption. Consumption still follows a random walk if the interest rate equals the rate of time preference.

VII. Conclusion

There are really two reasons to consider consumption from the standpoint of fractional differencing and its related long term dependence. First, if income follows a long memory process, consumption responds to an income stream unlike those previously considered, and so the relation between income and consumption should look strange. The examples in section 3 show that fractional differencing can produce many of the interesting anomalies observed. The impact must be somewhat muted, however, both because closed form predictions and accurate estimation techniques are hard to come by, and because serious tests suggest that income does not show long term dependence. The examples further show, however, that when dealing with fractional differences working only with aggregate quantities is particularly dangerous: aggregation problems cannot be assumed away, because aggregation causes fractional differencing.

Secondly, consumption of services shows long term dependence. Though consistent with previous work showing consumption is sensitive to past changes in income (Flavin 1981, Deaton and Campbell 1987), it adds a new anomaly. Output and other components do not show long term dependence. Services alone show up significant in different countries. Why services occupy this unique position remains a puzzle.

BIBLIOGRAPHY

Bernamont, J., "Fluctuations De Potential Aux Bornes D'un Conducteur Metallique De Faible Volume Parcouru Par un Courant," Annales de Physique (Leipzig), 1937, pp. 71-140, chapitre II.

Campbell, John, and Angus Deaton, "Why is Consumption so Smooth?" Princeton Univ., Working Paper, July, 1987.

Deaton, Angus, "Life Cycle Models of Consumption: is the evidence consistent with the Theory?" NBER Working Paper No. 1910, April, 1986.

Diebold, Francis X., and Glenn D. Rudebusch, "Long Memory and Persistence in Aggregate Output," Federal Reserve Board of Governors Finance and Economics Discussion Series, No. 7, Jan., 1988.

Diebold, Francis X., and Glenn D. Rudebusch, "Is Consumption Too Smooth? Long Memory and the Deaton Paradox," Board of Governors, March, 1989.

Feinstein, C. H., Statistical Tables of National Income, Expenditure and Output of the U.K. 1855-1965, Cambridge University Press, New York, 1972.

Flavin, Marjorie, "The Adjustment of Consumption to Changing Expectation about Future Income," Journal of Political Economy, 89, 1981, pp. 974-1009.

Geweke, John, and Susan Porter-Hudak, "The Estimation and Application of Long Memory Time Series Models," Journal of Time Series Analysis, 4, 1983, p. 221-238.

Granger, Clive, "Long Memory Relations and the Aggregation of Dynamic Models," Journal of Econometrics, 14, 1980, pp. 227-238

Granger, Clive, and Roselyne Joyeux, "An Introduction to Long-Memory Time Series Models and Fractional Differencing," Journal of Time Series Analysis, 1, 1980, pp. 14-29.

Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, 86, 1978, pp. 971-87.

Hansen, Lars P., and Thomas J. Sargent, "A Note on Wiener-Kolmogorov Prediction Formulas for Rational Expectations Models," Economics Letters, 8, 255-60.

Haubrich, Joseph G., and Andrew W. Lo, "The Sources and Nature of Long-Term Memory in the Business Cycle," Working Paper, July, 1988.

Hosking, J. R. M., "Fractional Differencing," Biometrica, 68, 1981, pp. 165-176.

Hurst, Harold E., "Long Term Storage Capacity of Reservoirs," Transactions of the American Society of Civil Engineers, 116, 1951, pp. 770-799.

Lo, Andrew W., "Long-Term Memory in Stock Market Prices," Working Paper, Sloan School of Management, 1988.

Mandelbrot, Benoit, "Statistical Methodology for Non-Periodic Cycles: From the Covariance to R/S Analysis," Analysis of Economic and Social Measurement, 1, 1972, pp. 259-290.

Mankiw, N. Gregory, "Hall's Consumption Hypothesis and Durable Goods," Journal of Monetary Economics 10, 3, Nov. 1982, pp. 417-425.

Newey, Whitney K., and Kenneth D. West, "A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," Econometrica, 55, 1987, pp. 703-708.

Phillips, Peter C. B., "Time Series Regression With a Unit Root," Econometrica, 55, 1987, pp. 277-301.

Ryder, H. E. and G. M. Heal, "Optimal Growth with Intertemporally Dependent Preferences," Review of Economic Studies, 40, 1973, pp. 1-31.

Sowell, Fallaw, "Maximum Likelihood Estimation of Fractionally Integrated Time Series Models," Duke University Working Paper, 1987.

Sundaresan, S. M., "Intertemporally Dependent Preferences in the Theories of Consumption, Portfolio Choice and Equilibrium Asset Pricing," Review of Financial Studies, forthcoming, 1989.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," Wharton Working Paper, 1985.

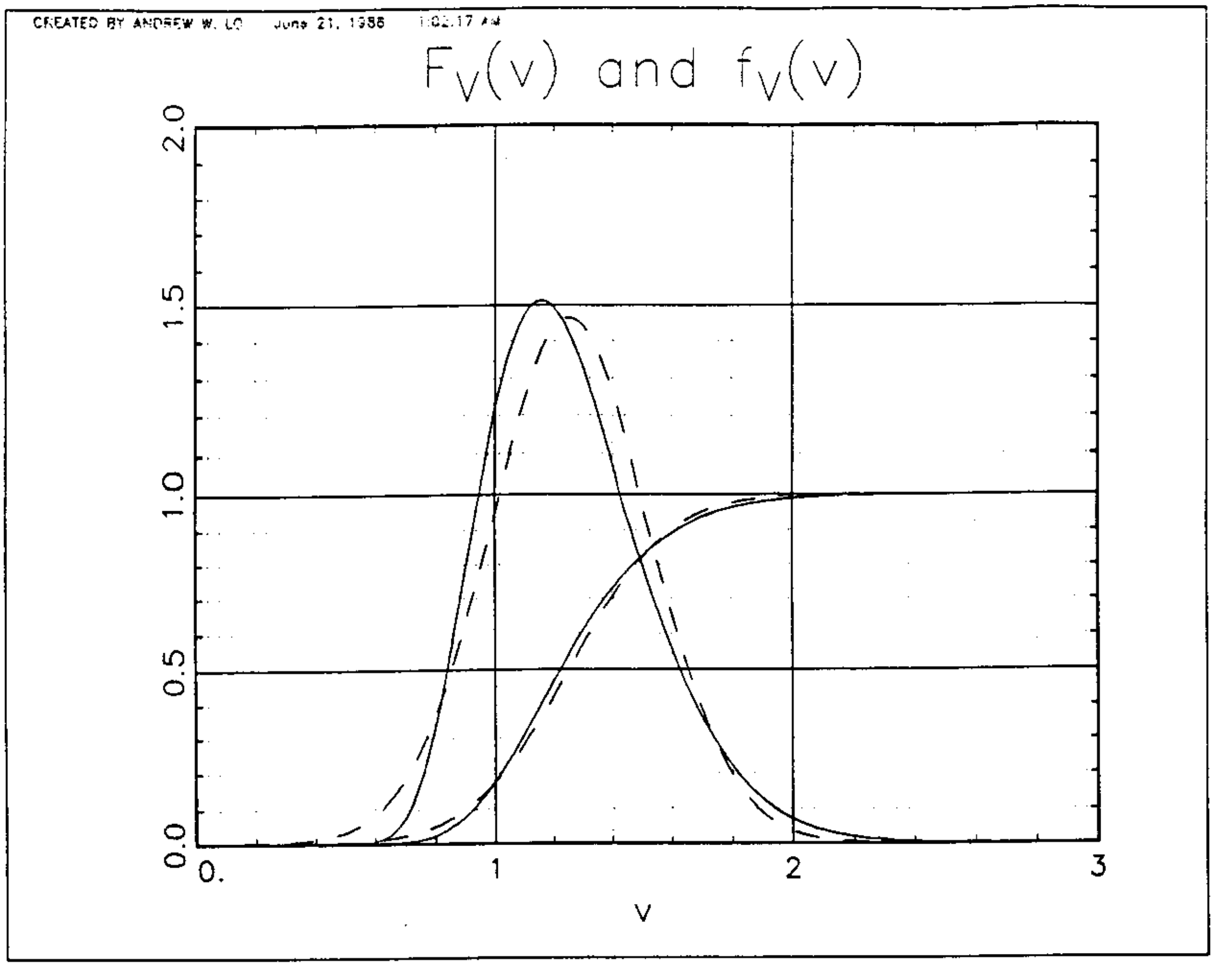

Figure .

Distribution and density function of the range V of a Brownian bridge. Dashed curves are the normal distribution and density functions with mean and variance equal to those of V.

Table 1a
Fractiles of the Distribution $F_V(v)$

$P(V < v)$	.005	.025	.050	.100	.200	.300	.400	.500
v	0.721	0.809	0.861	0.927	1.018	1.090	1.157	1.223

$P(V < v)$	.543	.600	.700	.800	.900	.950	.975	.995
v	$\sqrt{\frac{\pi}{2}}$	1.294	1.374	1.473	1.620	1.747	1.862	2.098

Table 1b

Symmetric Confidence Intervals about the Mean

$P\left(\sqrt{\frac{\pi}{2}} - \gamma < V < \sqrt{\frac{\pi}{2}} + \gamma\right)$	γ
.001	0.748
.050	0.519
.100	0.432
.500	0.185

Table 2
R/S Analysis of Consumption

Series	$V_n(0)$	$V_n(1)$	$V_n(2)$	$V_n(3)$	$V_n(5)$
Total C	1.607	1.478	1.379	1.287	1.193
Durables	0.989	1.031	1.105	1.016	0.983
Non-Durables	1.405	1.254	1.191	1.117	1.067
Services	1.936*	1.899*	1.829*	1.718	1.595

All Data: <u>Quarterly Real Expenditures</u>, Seasonally Adjusted, 1959:1-1988:4.
*Significant at 5% level.
$V_n(q)$ calculated on first differences of logs.

Table 3
Size and Power Simulations

A: POWER AGAINST $d = \frac{1}{3}$

N	Q	PWR 1%	PWR 5%	PWR 10%
100	0	0.430	0.603	0.684
100	1	0.192	0.373	0.478
100	2	0.065	0.214	0.326
100	3	0.011	0.114	0.209
100	5	0.001	0.012	0.062

B: IID NULL SIZE RESULTS

N	Q	SZE 1%	SZE 5%	SZE 10%
100	0	0.031	0.089	0.153
100	1	0.020	0.071	0.126
100	2	0.013	0.054	0.106
100	3	0.008	0.041	0.085
100	5	0.003	0.021	0.053

C: SIZE RESULTS AR(1) e = 0.5

N	Q	SZE1%	SZE 5%	SZE 10%
100	0	0.201	0.380	0.490
100	1	0.029	0.114	0.191
100	2	0.004	0.045	1.101
100	3	0.003	0.027	0.070
100	5	0.003	0.022	0.052

D: SIZE RESULTS AR(1) e = -0.25

N	q	SZE 1%	SZE 5%	SZE 10%
100	0	0.153	0.308	0.414
100	1	0.034	0.104	0.176
100	2	0.015	1.067	0.123
100	3	0.008	0.039	0.085
100	5	0.001	0.015	0.041

Each panel uses 10,000 independently generated series of 100 done with Fortran on a VAX 8700. The programs are a minor modification of the ones written by Lo for Haubrich and Lo (1988).

Table 4

United Kingdom

Series	$V_n(0)$	$V_n(1)$	$V_n(2)$	$V_n(3)$	$V_n(5)$
(i) 1870-1965					
Income	1.140	1.021	0.954	0.927	1.149
Consumption	1.506	1.337	1.349	1.375	1.505
(ii) 1959:I-1988:III					
GDP	1.077	1.213	1.274	1.310	1.375
Total C	1.315	1.440	1.429	1.353	1.334
Durables	0.660*	0.780	0.882	0.860	0.962
Non-Durables	1.209	1.302	1.305	1.268	1.260
Services	1.809*	1.708	1.593	1.470	1.305

(i) Annual Data 1870-1965: Feinstein 1972.

(ii) Quarterly Data Central Statistical Office. Real, Seasonally Adjusted, 1959:I-1988:III. Services and Non-Durables, 1963:I-1988:III.

*Significant at 5% level.
$V_n(q)$ calculated on first differences of logs.

Table 5

Canada

Series	$V_n(0)$	$V_n(1)$	$V_n(2)$	$V_n(3)$	$V_n(5)$
GDP	1.628	1.492	1.439	1.381	1.366
Total C	1.634	1.650	1.610	1.523	1.450
Durables	1.184	1.298	1.300	1.250	1.290
Non-Durables	1.705	1.955*	2.112*	2.016*	1.909*
Services	1.389	1.334	1.336	1.315	1.225

All Data: Real, Quarterly, Seasonally Adjusted CANSIM, Statistics Canada, 1960:I-1988:III.

*Significant at 5% level.
$V_n(q)$ calculated on first differences of logs.

Table 6

France

Series	$V_n(0)$	$V_n(1)$	$V_n(2)$	$V_n(3)$	$V_n(5)$
GDP	1.699	2.023*	2.073*	1.975*	1.889*
Total C	1.828*	1.856*	1.894*	1.797*	1.760
Durables	1.039	1.043	1.078	1.091	1.129
Non-Durables	1.019	1.110	1.440	1.398	1.430
Services	2.046*	2.039*	1.877*	1.796*	1.630

All Data: Real, Quarterly, Seasonally Adjusted. Data Stream Inc., from OECD Quarterly National Accounts.

*Significant at 5% level.

$V_n(q)$ calculated on first differences of logs.

GDP, Total C, Durables 1967:I-1988:II

Non-Durables 1970:I-1988:II

Services 1969:I-1988:II