An ARFIMA Model for Volatility Does Not Imply Long Memory

Richard A. Ashley
Department of Economics
Virginia Tech
(540) 231-6220
ashleyr@vt.edu

and

Douglas M. Patterson
Department of Finance
Virginia Tech
(540) 231-5737
amex@vt.edu

May 4, 2011

An ARFIMA Model for Volatility Does Not Imply Long Memory

ABSTRACT

Jiang and and Tian (2010) have estimated an ARFIMA model for stock return volatility. We argue that this result does not imply actual 'long memory' in such time series -- as any kind of instability in the population mean yields apparent fractional integration as a statistical artifact. Alternative high-pass filters for studying stock market volatility data are suggested.

Keywords: Long Memory, Fractional Integration, Volatility, Stock Returns, Time Series

I. Introduction.

This paper is motivated by Jiang and Tian (2010). They estimate ARFIMA models for stock return volatility. These models are used to forecast return volatility, with the goal of improving option valuation under the Statement of Financial Accounting Standards (SFAS) 123R. They show that when a fractional difference is incorporated into a vector-autoregressive (VAR) specification, superior volatility forecasts are obtained. Their fitted model includes, in addition to the ARFIMA fractional difference operator $-\ (1-B)^d,\ 0 < d < 0.5\ -$ some common market factors such as the volatility of the S&P 500 index. They describe this as an "LM-VAR' model, in which the 'LM' term refers to 'long memory.'

Our purpose is not to criticize their identification or estimation of these LM-VAR models as volatility forecast models. However, we point out that the fractional difference operator, $(1-B)^d$ at the heart of any $ARFIMA(p,d,q)$ model, is only one of many high-pass filters suitable for modeling stock market volatility.

The unfortunate aspect of Jiang and Tian's paper is its assertion that their estimated $ARFIMA$ model necessarily implies the existence of 'long memory' in stock return volatility data, where 'long memory' is intended to assert the existence of actual serial correlations in the data which die out only very slowly at long lags. That assertion is incorrect, basically because any estimated $ARFIMA(p,d,q)$ (or test of $H_0: d=0$) in essence relies on sample autocovariances as consistent estimators of population autocovariances. But these sample autocovariance estimators are inconsistent in the presence of any time variation in the population mean of the time series. Thus, the apparent presence of fractional integration in a time series is more likely to be signaling the presence such time variation – structural shifts or weak trends – than to be an authentic indicator of 'long memory.'

II. Fractional integration and 'long memory.'

It is well known – e.g., Granger (1980), Granger and Joyeaux (1980), Beran (1994), and Baillie (1996) – that fractional integration of order d in a time series implies that its population autocorrelation at lag k decay slowly as k increases. In particular, if $0 < d < 1/2$, for the $ARFIMA(p,d,q)$ model, then $\rho_k \propto k^{2d-1}$ $as\, k \to \infty$. Concomitantly, fractional integration of order d is equivalent, under certain conditions, to an exploding spectral density at zero frequency – i.e., to $s(\omega) \propto \omega^{-2d}$ $as\ \omega \to 0^+$, where $s(\omega)$ is the power spectrum of $\{x_t\}$.[1]

For this reason, observed slow decay in the sample autocorrelation function r_k, is taken to be evidence that $\{x_t\}$ is generated by an $ARFIMA(p,d,q)$ model. Similarly, an observation that $\log[\hat{s}(\omega)]$ -- the logarithm of the sample spectrum $\hat{s}(\omega)$ -- goes to zero linearly as $\log(\omega) \to 0$ is also taken to be evidence that $\{x_t\}$ is generated by an $ARFIMA(p,d,q)$ model. Indeed, the most popular ways to estimate d and to test the null hypothesis H_0: $d = 0$ are based on regressing

[1]The power spectrum is defined as the Fourier transform of the autocovariance function ; the corresponding population autocorrelation function $\rho_k \equiv \dfrac{\mathrm{cov}(x_t, x_{t-k})}{\sqrt{\mathrm{var}(x_t)\,\mathrm{var}\,(x_{t-k})}}$ is a simple transform of the autocovariance function.

$\log[\hat{s}(\omega)]$ of $\{x_t\}$ against $\log(\omega)$ for low frequencies – see Geweke and Porter-Hudak (1983) and Robinson (1995).

This, in fact, is precisely how Jiang and Tain (2010) estimate the value of d in their $ARFIMA(p,d,q)$ models for stock return volatility. And, in principle there is nothing wrong with estimating d is this way and using these estimates to identify and estimate short-term forecasting models for stock return volatility.

The problem arises when one interprets this modeling effort as implying the existence of actual long-range serial correlation in the data, as this interpretation is valid only if the d estimates are based on consistent estimates of $s(\omega)$ or, equivalently, on consistent estimators of the autocovariances underlying the ρ_k .

But these autocovariance estimates cannot be consistent if $E\{x_t\}$ varies over time. This proposition is proven in Ashley and Patterson (2010, Section 3.1); but this demonstration is so brief and straightforward that we repeat it here. The *population* autocovariance of $\{x_t\}$ with itself lagged k periods is defined as:

$$\text{cov}(x_t x_{t-k}) = \frac{1}{T}\sum_{t=1}^{T}[x_t - E\{x_t\}][x_{t-k} - E\{x_{t-k}\}] \quad ,$$

whereas the *sample* autocovariance of $\{x_t\}$ with itself lagged k periods is defined as:

$$\widehat{\text{cov}}(x_t, x_{t-1}) = \frac{1}{T}\sum_{t=k+1}^{T}[x_t - \overline{x}][x_{t-k} - \overline{x}].$$

But if $E\{x_t\}$ varies over time – due to either structural shifts or due to any kind of smooth variation or trend-like behavior – then the sample mean $\overline{x}$ cannot possibly be a consistent estimator of $E\{x_t\}$.[2] Thus, in that case, the sample autocovariances cannot possibly be consistent estimators of the population autocovariances. Hence, neither the autocovariances nor the power spectrum based on them is consistently estimated in that case.

This result rationalizes what otherwise might appear to be a contradiction between our acceptance of Jian and Tian's *ARFIMA* forecasting model and our rejection of their interpretation of this model as implying slowly decaying long-range linear dependence in these data. Their model might in this situation be a useful approximation for short-run forecasting, but their interpretation of their estimated model as implying that ρ_k is non-negligible for large k is not credible. As noted above, it is much more likely that the sample autocovariances underlying their estimated short-run model are inconsistent estimators of the population autocovariances because of time variation in $E\{x_t\}$ than it is that there is any noticeable relationship between current stock return volatility and its distant past.

[2] See Ashley and Patterson (2010, Section 3.1).

III. Dealing with a slowly-evolving $E\{x_t\}$.

Like a fractionally integrated process, a process with a slowly evolving mean is also characterized by an autocorrelation function that slowly decays as the lag becomes large[3]. Given the similarity between these two models in the sample statistics, it becomes an empirical question as to whether or not the process is an example of fractional integration or a slowly evolving mean; a serial correlation coefficient that slowly decays is not a sufficient condition for fractional integration. In Ashley and Patterson (2010) we studied the weekly volatility of the CRSP value weighted stock index. There we present evidence based on the sample power spectrum that stock market volatility is not generated by a fractional difference process. In particular, the behavior of the spectrum at low frequencies is not indicative of a fractional difference process.

In Figure 1 we display the sample power spectrum at low frequencies from a generated approximation to a fractionally integrated $(ARFIMA)$ model, using an

[3] This is proven in Theorem 2 of Ashley and Patterson (2010).

$MA(1,000,000)$ approximation to the *ARFIMA* process.[4] The *MA* model was truncated after 1,000,000 terms. For reference purposes in Figure 1, the theoretical spectrum for the fractionally integrated process is plotted with the "+" symbol. Contrast the plot in Figure 1 with Figure 2, where we plot the low frequency spectral behavior of the observed weekly volatility of returns to the value weighted CRSP index[5]. The sample period is the 2,705 weeks spanning March 1956 through the last week in December of 2007.

Figure 2 shows that the observed power spectrum of the CRSP return weekly volatility actually dips at the lowest frequencies, rather than exploding – per the fractional integration hypothesis and Figure 1 – as the frequency approaches zero from the right. Thus, a closer look at related sample evidence is actually not very supportive of the conclusion that weekly stock market index volatility is generated by an *ARFIMA* process. From this evidence the conclusion that stock market volatility is generated by an *ARFIMA* process is not exactly compelling.

---

[4] An *ARFIMA* process can be very compactly expressed using the $(1-B)^d$ operator, but one of its awkward features is that the notion of actually differencing a time series a fractional number of times is intuitively opaque. Relatedly, another awkward feature is that *ARFIMA* processes can only be approximately simulated, using expansions of $(1-B)^d$ or $(1-B)^{-d}$ to obtain $AR(p)$ or $MA(q)$ approximations to a fractionally integrated process. See Hamilton (1994, pp. 447-9) where formulas for these expansions are derived. These expansions converge very slowly, thus extremely large values of p or q are necessary in order for the expansion to yield an adequate approximation to the correlogram, spectrum, etc. of a fractionally integrated process.

[5] Weekly volatility is measured as the root mean square of the daily returns during the calendar week.

Nevertheless, as Jiang and Tian (2010) observe, one can fit an *ARFIMA* model to the data and obtain useful short-term forecasting models. And this is by no means wrong, so long as one recognizes that one is not adducing evidence for 'long memory' in the actual data generating process.

Alternatively, one could estimate less elegant – but more easily interpretable and (perhaps) more appropriate – models for these data by observing that the fractional difference operator is in this context simply serving as a high-pass filter for the data.

Many alternative high-pass filters are well known in the time series analysis and Electrical Engineering literatures. These include:

1) Exponential Smoothing[6]
2) Moving Average[7]
3) Nonlinear Bandpass filtering[8]
4) Butterworth filter[9]

---

[6] See Granger and Newbold (1977) pages 162-165, and SAS/ETS User's Guide, Version 6, Second Edition (1995) Chapter 9, page 443.
[7] See Ashley and Patterson (2007) for a simple example applied to weekly stock index volatility.
[8] These are common in the macroeconomic time-series literature – see Baxter and King (1999).

5) Non parametric time regression[10]

High-pass filtering of these types can eliminate a slow, smooth time variation in the mean of a financial return volatility series just as easily as does a fractional difference. Indeed, Ashley and Patterson (2007, 2010) used some of these filters to eliminate the sample evidence of fractional integration in several such time series.

Of course, some high pass filters are more intuitively appealing than others; some readers might even prefer the compactness of the fractional difference filter, despite its intuitive opaqueness. One would expect to obtain similar short-term forecasts from models based on any of these choices, so our main complaint with the *ARFIMA* approach is its concomitant implication of 'long memory' in the time series. We also note that the *ARFIMA* model eliminates any trends in the time series at the outset, whereas an analyst applying some other high-pass filter is more likely to also examine the trend which, albeit weak, might be of economic interest.

---

[9] See A. Antoniou, *Digital Filters: Analysis, Design, and Applications*, New York, NY: McGraw-Hill, 1993, and, S.K. Mitra, *Digital Signal Processing: A Computer-Based Approach*, New York, NY: McGraw-Hill, 1998..
[10] See Ashley and Patterson (2010).

REFERENCES

Antoniou, A., *Digital Filters: Analysis, Design, and Applications*, New York, NY: McGraw-Hill, 1993.

Ashley, R. A., and D. M. Patterson, "Apparent Long Memory in Time Series as an Artifact of A Time-Varying Mean: Considering Alternatives to the Fractionally Integrated Model," *Macroeconomic Dynamics*, 14 (Supplement 1) (2010), 59-87.

______________________________, "Apparent Long Memory in a Time Series and the Role of Trends: a 'Moving Mean' Filtering Alternative to the Fractionally Integrated Model," unpublished working paper, Department of Economics, Virginia Polytechnic Institute, (2007).

Baxter, M. and R. King, "Measuring Business Cycles: Approximate Band-Pass Filters foe Economic Time Series," *Review of Economics and Statistics*, 81 (1999), 575-593.

Granger,C. W. J., and P. Newbold, *Forecasting Economic Time Series*, New York: Academic Press (1997).

Hamilton, J. D., *Time Series Analysis*, Princeton NJ: Princeton University Press (1994).

Jiang, G. J., and Y. S. Tian, "Forecasting Volatility Using Long Memory And Comovements: An Application to Option Valuation Under SFAS 123R," Journal Of Financial And Quantitative Analysis, 45(2) (2010), 503-533.

Mitra, S. K., *Digital Signal Processing: A Computer-Based Approach*, New York, NY: McGraw-Hill, 1998.

SAS/ETS User's Guide, Version 6, Second Edition, 2nd printing, Cary NC: SAS Institute Inc. (1995).

Schwert, G. W., "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance*, 44 (1989), 1115-1153.

Figure 1. The sample power spectrum at very low frequencies from a generated approximation to a fractionally integrated model is plotted as the solid line. The theoretical spectrum for the fractionally integrated process is plotted with the “+” symbol.

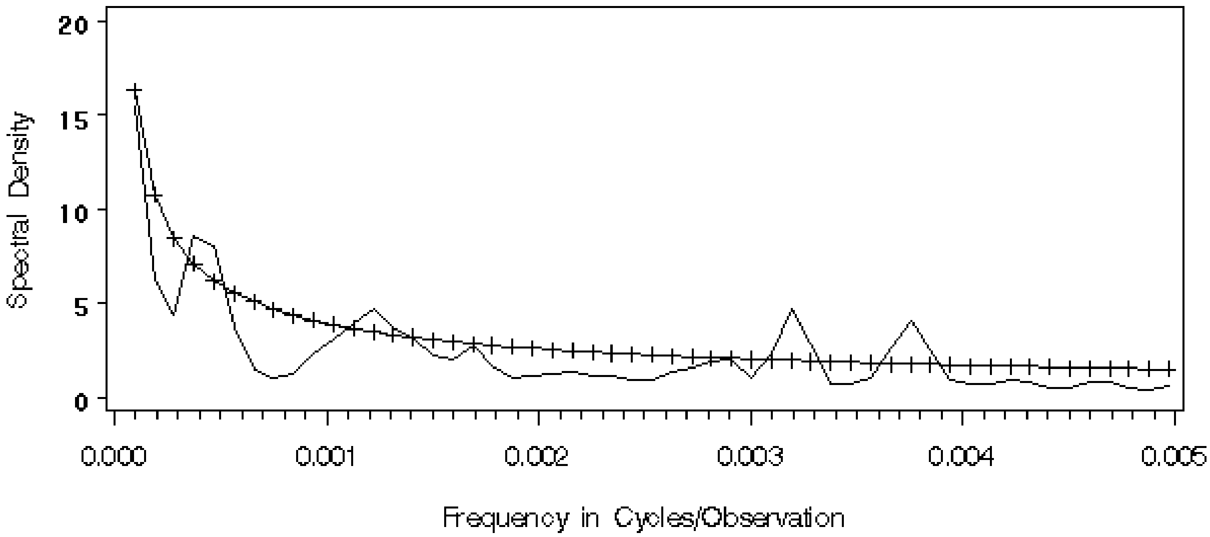

Figure 2. The sample power spectrum of the weekly volatility of the CRSP value weighted stock market index. The period shown are the 2,705 weeks from March 1956 through December 2007.

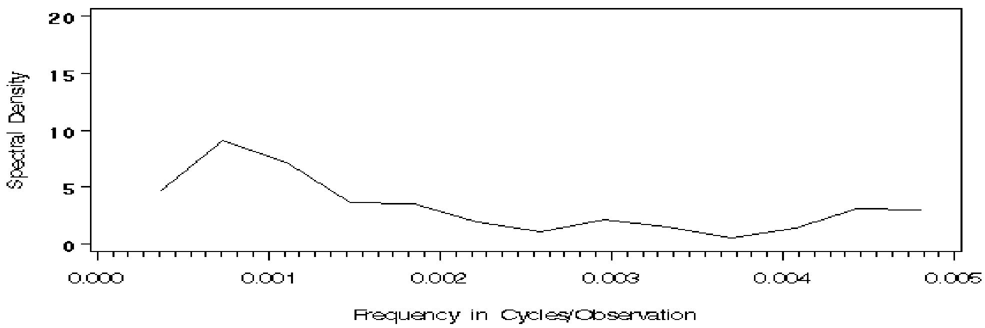

20
15
10
5
0
Spectral Density
0.000
0.001
0.002
0.003
0.004
0.005
Frequency in Cycles/Observation