# Median-Unbiased Estimation of Higher Order Autoregressive/Unit Root Processes and Autocorrelation Consistent Covariance Estimation in a Money Demand Model

J. Huston McCulloch[1]
Ohio State University

June 2008
Revised Oct. 2008

Econometric Society 2008 North American Summer Meetings
Pittsburgh, June 19-22, 2008

14th Annual Conference on Computing in Economics and Finance
Paris, June 26-28, 2008

**ABSTRACT**

It is shown that the Newey-West (1987) Heteroskedasticity and Autocorrelation Consistent (HAC) covariance matrix estimator can greatly understate the standard errors of OLS regression coefficient estimates in finite samples, and therefore comparably overstate t-statistics. Although the bias vanishes in infinite samples and is tolerable in samples as small as $10^6$, it can lead to t-statistics that are too high by a factor of 1.7-2.2 with a sample size in the range 65-1200 and first order autoregressive serial correlation with AR coefficient 0.9.

The Exactly Median Unbiased estimator of Andrews (1993) for a *directly* observed AR(1) process is extended to the case of an AR(p) process that is only *indirectly* observed via OLS regression residuals. By allowing the maximum permitted order to increase without limit with the sample size, the estimator consistently estimates a stationary process with any autocovariance function. It also provides a unit root test (and therefore a test for cointegration of the regressors) that is exact up to median unbiased estimates of the higher order persistences.

These Median Unbiased Autoregressive (MUAR) estimates of the autocovariance function are then used to construct an Autocorrelation Consistent (MUAR-AC) covariance matrix for the OLS coefficient estimates. Applied to a simple model of US demand for narrow money M1-S (official M1 + estimated Retail Sweep Accounts), it is found that a unit root in the errors and therefore absence of cointegration can be at least weakly rejected (at the 10% test size). The MUAR-AC standard errors are 128-153% higher than HAC standard errors, or equivalently, HAC t-statistics for any hypothesis concerning the coefficients are 128-153% too large.

Despite the greatly increased MUAR-AC standard errors, the income elasticity and interest semielasticity of demand for M1-S remain highly significant, suggesting that narrow money (currency plus all checking accounts) may still be a useful indicator of monetary policy, and that the Fed should resume collecting direct data on it.

## 1. Introduction

Serial Correlation is a pervasive problem in time series models in econometrics, as well as in statistics in general. When, as is ordinarily the case, the serial correlation is positive, Ordinary Least Squares (OLS) standard errors are generally too small, and the derived t-statistics are too large.

The truncated-kernel Heteroskedasticity and Autocorrelation Consistent (HAC) covariance matrix, introduced by Newey and West (1987), is now widely used by economists to "correct" the standard errors of OLS time series coefficients for serial correlation. Greene (2003: 201) reports that its use is now "standard in the econometrics literature." Hayashi (2000: 409-12) mentions only it, the similar Quadratic Spectral HAC of Andrews and Monahan (1992), and the Vector Autoregression HAC (VAR-HAC) method of den Haan and Levin (1996) as appropriate methods for correcting OLS standard errors for serial correlation. Stock and Watson (2007) present HAC as the *only* method worthy of consideration.

This study shows that despite their consistency, HAC standard errors can greatly overstate the precision of OLS coefficient estimates with sample sizes and serial correlation commonly found in economic studies when, as has become standard, "automatic bandwidth selection" is employed. Practitioners are always in search of big t-statistics and therefore small standard errors, so it is entirely understandable that the HAC under-correction for serial correlation has become so popular. However, it is bad econometric practice to systematically overstate the significance of one's results by deliberately choosing a deficient estimator.

The vintage Cochrane-Orcutt estimator (Greene 2003: 273) uses OLS estimates of an AR(1) model of the residuals to transform the regression equation in an attempt to remove the serial correlation. Higher order AR disturbances may also be accommodated in a similar manner (Beach and MacKinnon 1973; Greene 2003: 274). However, it is well known (e.g. Greene 2003: 636) that OLS and even Maximum Likelihood (ML) estimates of AR models are biased so as to understate their persistence. The bias does go away in large samples, so that consistency can be claimed as long as the order of the autoregression is allowed to increase without bound with sample size as in den Haan and Levin (1996). With typical econometric sample sizes, however, the bias can be substantial. The standard errors from the transformed regression will therefore also have a downward bias.

Andrews (1993) shows how exactly median-unbiased estimates of the autoregressive coefficient in a Gaussian AR(1) model can be obtained by means of a simple simulation procedure. Andrews and Chen (1994) extend this to obtain only approximately median-unbiased estimates of the coefficients of an AR(p) model. The present paper shows how to extend the method of Andrews (1993) so as to obtain *exactly* median unbiased estimates of the coefficients of an AR($p$) model with Gaussian errors.

This paper also goes beyond both Andrews (1993) and Andrews and Chen (1994), which consider only the residuals of a univariate model with a constant and/or time trend, by applying the method to the *residuals of a general OLS mode*l with unobserved AR(*p*) errors. It obtains exact confidence intervals for each of the AR(*p*) coefficients (conditional on the median-unbiased estimates of each of the other coefficients), and shows how to conduct an exact unit root test on the residuals, which in the case of OLS residuals, amounts to an exact, *finite sample* test for cointegration, conditional again on the median-unbiased estimates of each of the other coefficients.

The paper devises an algorithm to obtain the exactly Median Unbiased AR(*p*) (MUAR) coefficients from the residuals of a general OLS regression with Gaussian errors. These MUAR coefficients are then used to obtain Autoregressive Consistent (MUAR-AC) estimates of the standard errors of OLS regression coefficients.

The MUAR-AC algorithm is then illustrated by using it to consistently adjust the standard errors of an updated OLS money demand equation for serial correlation, and to test it for cointegration. The results are compared to HAC and uncorrected AR(p)-based standard errors.

Regressor-Conditional Heteroskedasticity (RCH), discussed at greater length in the final section, can also distort OLS standard errors, if and when it is present. The paper discusses an autocorrelation-corrected test for RCH and possible ways to modify the MUAR-AC estimator to deal with this problem when it has been detected. At present this aspect of the paper is still "under construction."

## 2. Finite Sample Bias of Truncated-Kernel HAC

Consider a time-series linear regression of the form
$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$
where $\mathbf{X} = (x_{tj})$ is a $T \times k$ matrix of exogenous regressors whose first column is a vector of units. The OLS estimator of $\boldsymbol{\beta}$,
$$\hat{\boldsymbol{\beta}} = (\mathbf{X'X})^{-1}\mathbf{X'y} = \boldsymbol{\beta} + (\mathbf{X'X})^{-1}\mathbf{X'\varepsilon}$$
then has covariance matrix
$$\mathbf{C} = \mathbf{Cov}(\hat{\boldsymbol{\beta}}) = E((\mathbf{X'X})^{-1}\mathbf{X'\varepsilon\varepsilon'X}(\mathbf{X'X})^{-1}). \tag{1}$$
If the $T \times 1$ error vector $\boldsymbol{\varepsilon}$ is independent of $\mathbf{X}$ and has a time-invariant or Toeplitz autocovariation structure,
$$\boldsymbol{\Gamma} = E(\boldsymbol{\varepsilon\varepsilon'}) = \left(\gamma_{|t-t'|}\right),$$
this becomes
$$\mathbf{C} = (\mathbf{X'X})^{-1}\mathbf{X'\Gamma X}(\mathbf{X'X})^{-1}. \tag{2}$$
Under the classic OLS assumption $\boldsymbol{\Gamma} = \gamma_0\mathbf{I}$, (2) becomes
$$\mathbf{C}^{OLS} = \gamma_0(\mathbf{X'X})^{-1}. \tag{3}$$
The OLS residuals
$$\mathbf{e} = (\mathbf{I} - \mathbf{X}(\mathbf{X'X})^{-1}\mathbf{X'})\boldsymbol{\varepsilon} = \mathbf{M}\boldsymbol{\varepsilon} \tag{4}$$
may be used to compute the (strong) estimates of the autovariances:

$$\hat{\gamma}_j = \frac{1}{T-j} \sum_{t=j+1}^{T} e_t e_{t-j}, \quad j = 0, \dots T-1.^2 \tag{5}$$

The Newey-West truncated-kernel HAC estimator is

$$\hat{\mathbf{C}}^{HAC} = (\mathbf{X'X})^{-1} \mathbf{X'FX}(\mathbf{X'X})^{-1},$$

where

$$\mathbf{F} = \left( \varepsilon_t \varepsilon_{t'} k(|t-t'|) \right)$$

and

$$k(j) = \max((m-j)/m, \, 0) \tag{6}$$

is the *truncated Bartlett Kernel function* for some *bandwidth m*. Most econometric packages provide "automatic bandwidth selection" for HAC, using a formula similar to the following "benchmark rule" recommended by Stock and Watson (2007: 607):

$$m = \text{round}(0.75 \, T^{1/3}).^3 \tag{7}$$

Since this formula forces $m$ to rise without bound as $T$ rises to infinity, it ensures that all autocorrelations will eventually be used by $\hat{\mathbf{C}}^{HAC}$. At the same time, $T/m$ rises without bound, so that the $m$-1 included autocorrelations will be estimated with increasingly high precision. Furthermore, for any $j$ the Bartlett factor $(m-j)/m$ will eventually become arbitrarily close to 1. As a result, the HAC covariance matrix consistently estimates the true covariance matrix, provided a finite eighth moment condition for the errors is met.

Stock and Watson do suggest that a higher or lower value of $m$ be tried as well, depending on the degree of serial correlation, but provide no procedure to implement this suggestion. The vast majority of practitioners who use HAC simply use whatever automatic rule is provided by their software and assume that this provides a state-of-the-art cure for any serial correlation problem.

In the benchmark case of homoskedasticity, the expectation of the HAC estimator is

$$\mathrm{E}\,\hat{\mathbf{C}}^{HAC} = (\mathbf{X'X})^{-1} \mathbf{X'\Gamma}^{HAC}\mathbf{X}(\mathbf{X'X})^{-1}, \tag{8}$$

where

$$\mathbf{\Gamma}^{HAC} = \left( \gamma_{|t-t'|} k(|t-t'|) \right)$$

HAC thus effectively employs only the first $m$-1 autocovariances, and replaces the others with zeros. At the same time it down-weights or damps the autocovariances it does use by the Bartlett Kernel factor $(m-j)/m$. For both these reasons, it tends to underestimate the coefficient variances. However, the amount by which it does this depends on both the $\gamma_j$ and the degree of serial correlation of the regressors themselves.

---

[2] Many authors, including Hayashi (2000: 408), divide by $T$ rather than $T$-$j$ in (5), giving what may be called the *weak* estimates of the autocovariances. This effectively applies the untruncated Bartlett filter weights of Kiefer and Vogelsang (2002) etc. to the strong autocovariance estimates. (See below.) Applying an additional factor of T/(T-k) will debias the estimate of $\gamma_0$ in the absence of serial correlation, but this adjustment is small in even moderate samples.

[3] EViews, following a suggestion of Newey and West, uses $m = \text{floor}(4\,(T/100)^{2/9})+1$, which has very similar effect for $T$ in the range 50 – 2000, and generally identical effect in the range 400 – 1000.

To illustrate this effect, we may, with no loss of generality, choose an orthogonal basis for the regressors, retaining a vector of units as the first regressor, so that $\mathbf{X'X}$ and its inverse are diagonal, and each regressor after the first has zero mean. Let

$$r_{ij} = \sum_{t=i+1}^{T} x_{tj} x_{t-i,j} / \sum_{t=1}^{T} x_{tj}^2$$

be a measure of the serial correlation of the $j$-th orthogonalized regressor at lag $i$. It can then be shown that the true variance of $\hat{\beta}_j$ given by (2) exceeds its OLS variance as given by (3) by a factor of

$$f_j^{OLS} = c_{jj} / c_{jj}^{OLS} = 1 + 2\sum_{i=1}^{T-1} \gamma_j r_{ij} / \gamma_0 , \tag{9}$$

whereas the expected NW variance estimator (8) exceeds the OLS variance by a factor of only

$$f_j^{HAC} = c_{jj}^{HAC} / c_{jj}^{OLS} = 1 + 2\sum_{i=1}^{T-1} \gamma_i \, \text{k}(i) r_{ij} / \gamma_0$$

The serial correlation in econometric regressions is often to a first approximation AR(1) in structure,

$$\varepsilon_t = \varphi \varepsilon_{t-1} + u_t , \tag{10}$$

where $0 << \ < 1$ and the innovations $u_t$ are iid. Under (10), the autocovariances are

$$\gamma_i = \gamma_0 \varphi^i . \tag{11}$$

In a simple money demand regression, for example, with log real money balances as the dependent variable and a constant, log real income, and an interest rate variable as the explanatory variables, it is not unusual for the estimated first order serial correlation to be on the order of 0.9 using quarterly data.[4]

The unitary first regressor is perfectly correlated with itself at all lags, so that

$$r_{i1} = (T-i)/T . \tag{12}$$

The other $r_{ij}$ are necessarily less than this in absolute value. If regressor $j$ has a pronounced upward trend, as does real income, for example, $r_{ij}$ will be near this upper bound for the smaller values of $i$ for which $\varphi^i$ is still perceptible. For $i > T/2$, $r_{ij}$ will actually tend to be negative, but by then $\varphi^i$ will be essentially zero. Even if the regressor is stationary without drift, such as the unemployment rate, interest rates, or inflation, it may still have a value very near this upper bound for the values of $i$ for which $\varphi^i$ still matters. In many cases, therefore, the $f_j$ will be near the maximal value as determined by (12).

Table 1 shows illustrative values of $f_j^{OLS}$, $f_j^{HAC}$, and their ratio $f_j / f_j^{HAC}$, the factor by which HAC underestimates the variance, for various sample sizes $T$, using

---

[4] For example, the median-unbiased estimate of the first order persistence of the residuals of the money demand equation reported in Table 3 below is 0.920.

bandwidths $m$ determined by the Stock and Watson "benchmark" formula (7).[5] It is assumed for this illustration that the errors are AR(1) with $\varphi = 0.9$, and that the $r_{ij}$ take on their "worst case" value (12).

**Table 1**
Factor by which HAC can underestimate coefficient variance.
AR(1) errors with $\varphi = 0.9$, maximal $r_{ij}$ as in (12).

| T | m | $f_j^{OLS}$ | $f_j^{HAC}$ | $f_j^{OLS} / f_j^{HAC}$ |
|---|---|---|---|---|
| 65 | 3 | 14.27 | 2.67 | 5.34 |
| 150 | 4 | 16.75 | 3.47 | 4.83 |
| 300 | 5 | 17.84 | 4.22 | 4.23 |
| 500 | 6 | 18.29 | 4.91 | 3.73 |
| 800 | 7 | 18.56 | 5.56 | 3.34 |
| 1200 | 8 | 18.70 | 6.12 | 3.04 |
| $10^4$ | 16 | 18.96 | 9.83 | 1.93 |
| $10^5$ | 35 | 19.00 | 13.98 | 1.36 |
| $10^6$ | 75 | 19.00 | 16.60 | 1.14 |
| $\infty$ | $\infty$ | 19.00 | 19.00 | 1.00 |

Note: $f_j^{OLS}$ = True variance/OLS variance

$f_j^{HAC}$ = HAC variance/OLS variance

$f_j^{OLS} / f_j^{HAC}$ = True variance/HAC variance

It may be seen from the last column of Table 1 that with as few as $10^6$ observations, the potential bias in the HAC covariance matrix is down to the tolerable level of 14%. A million is admittedly far short of infinity. However, with more typical macroeconometric sample sizes, on the order of a few hundred if that, HAC can easily underestimate the covariance matrix by an unacceptable factor of 3 – 5, and therefore overstate t-statistics by an equally unacceptable factor of approximately 1.7 – 2.2. Since econometric practitioners are always in search of "good" (i.e. big) t-statistics, it is small wonder that the HAC "correction" for serial correlation has become so popular.[6]

The situation is even worse in a regression whose dependent variable consists of overlapping multi-period averages. Multi-period averaging does greatly reduce the

---

[5] The selected values of $T$ up to 1200 are round numbers that generate the indicated integer values of $m$, with only minimal rounding.

[6] The findings of Table I are corroborated already by the simulations of Andrews and Monahan (1992). They develop a Quadratic-Spectral HAC (QS-HAC) estimator that has better asymptotics than truncated-kernel HAC, but which has similar effect with moderate sample sizes. Their simulations comparing the coverage of parametric AR(1) confidence intervals to that of their QS-HAC estimator find that the parametric estimator always does better in the absence of Regressor Conditional Heteroskedasticity (RCH), even without the AR(1) debiasing of Andrews (1993), and even when the true model is MA(m) rather than AR(1). See their Table IV top and Table V bottom. The AR($p$) correction proposed here would fit MA($m$) even better than AR(1). Their other tables do demonstrate that in the presence of acute RCH, it can be quite beneficial to take this into account as well.

residual variance and therefore OLS estimates of standard errors. However, the averaging also generates serial correlation even if none was present to begin with, and therefore makes these standard errors invalid. Using non-overlapping observations would eliminate the induced serial correlation, but would then reduce the sample size and push the OLS standard errors back up, perhaps even higher than with a non-averaged dependent variable. The HAC undercorrection for serial correlation has understandably become virtually *de rigueur* in such studies.

Stock and Watson (2007: 647) do recognize this overlapping average problem, and suggest increasing *m* in this circumstance, but again provide no rule for how to do this. A value of *m* equal to a large multiple of the averaging horizon would be required to overcome the Bartlett factors employed by the standard HAC estimators. And simply replacing the Bartlett factors with unity would not address the problem that any autocorrelations of order *m* and higher are completely ignored by HAC.

## 3. Median-Unbiased Estimation of Higher Order Autoregressive/Unit Root Models

Andrews (1993) corrects the OLS estimate of an AR(1) regression such as (10) by simulating the process with Monte Carlo simulations and then finding the function $m(\varphi)$ that gives the median of the OLS estimate of the AR(1) coefficient as a function of the true coefficient. He then inverts this function at the actual OLS estimate $\hat{\varphi}^{OLS}$ to find the median-unbiased (MU) estimate $\hat{\varphi}^{MU}$:

$$\hat{\varphi}^{MU} = m^{-1}(\hat{\varphi}^{OLS}).$$

This method is exact, to within Monte Carlo simulation error.[7]

Since the median-unbiased estimate of $\varphi$ will lead to median-unbiased estimates of the higher order autocorrelations $\varphi^i$ (for $\varphi \geq 0$) and of the Cumulative Impulse Response, CIR = $1/(1-\varphi)$, whereas a mean-unbiased estimate of $\varphi$ would lead to mean-biased estimates of these and other important nonlinear functions of $\varphi$ because of Jensen's Inequality, he argues that the median-unbiased criterion is more useful for the AR(1) coefficient than the mean-unbiased criterion.

Although Andrews (1993) dealt only with a univariate process that incorporates a constant and/or time-trend, the method can easily be applied to the residuals of an OLS regression, as follows: Zero-mean, trendless AR(1) errors **ε** may be simulated directly from (10) for any value of $\varphi$. The errors may then be converted into simulated residuals **e** using the **X** matrix in question, by means of (4). Equation (10) may then be estimated by OLS using these residuals, and the median function computed from these OLS estimates. Fortunately, this procedure does not depend on the true regression coefficients, or on the unknown variance of the innovations.

---

[7] The method of Roy and Fuller (2001) is similar in spirit to that of Andrews (1993). However, Roy and Fuller estimate the AR coefficients indirectly from the t-statistic for the hypothesis of a unit root, rather than directly from the estimated AR(1) coefficient or sum of AR coefficients as in Andrews and Chen (1994). Since the expected AR(p) standard errors are themselves functions of the AR coefficients, little if anything is gained by taking them into account in addition to the estimated coefficients.

Andrews (1993) notes that an exact confidence interval for $\varphi$ can easily be found by also computing quantile functions for quantiles other than the median. Since the distribution of the OLS estimator is monotonic in the true parameter, the .025 quantile function, for example, can be inverted at the sample OLS estimate to obtain the upper bound of a one-sided .975 confidence interval.[8] Similarly, the .975 quantile function can be inverted to obtain the (necessarily lower) lower bound of a one-sided .975 confidence interval. The intersection of the two one-sided intervals is then a two-sided 95% confidence interval. If the inverted $p$-quantile function of $\hat{\varphi}^{OLS}$ is less than unity, a unit root may be rejected at level $p$.

Although Andrews (1993) does not mention it, if the method is applied as described above to the residuals of an OLS regression, his unit root test becomes an exact, finite sample test for cointegration that does not rely on any asymptotic properties of the estimator. It is, however, contingent on an AR(1) error structure.
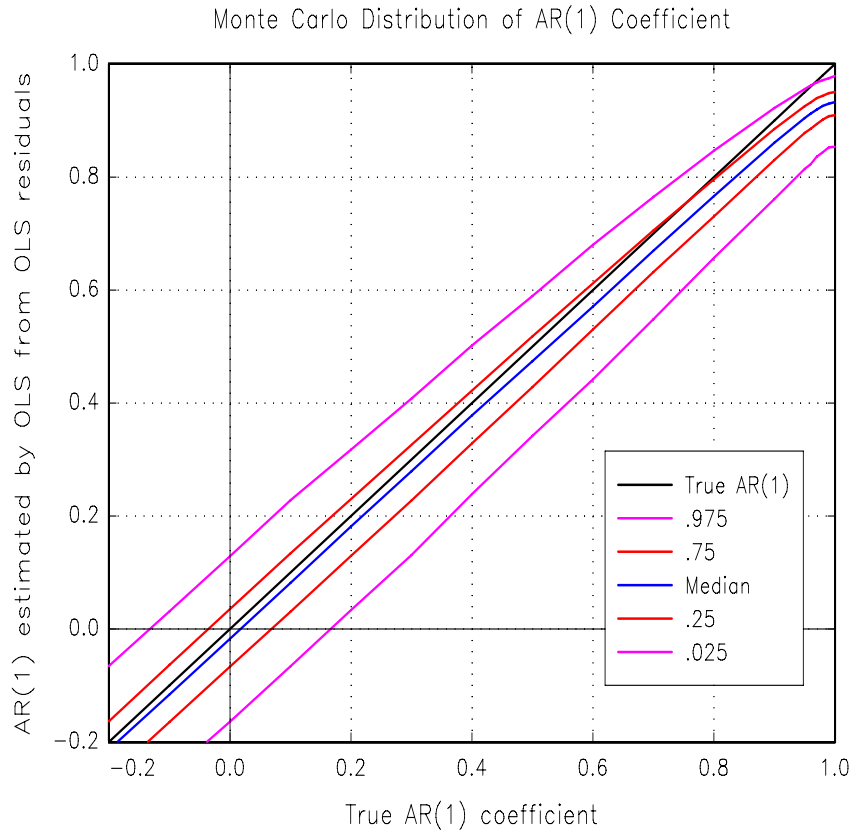


**Figure 1**

Figure 1 illustrates these extensions of Andrews' method, using the residuals of the OLS money demand regression discussed in Section 5 below. For each value of the

---

[8] When applied to regression residuals, this upper bound is essentially equivalent to the exact critical value of the Durbin-Watson (1950) one-sided test for zero first order serial correlation.

true AR(1) coefficient $\varphi$, 999 simulated sets of errors obeying (10) are generated.[9] These are converted to simulated OLS residuals using (4). These residuals are then fit to (10) by OLS (with no constant, since the column of units in **X** forces the residuals to sum to 0) to obtain 999 simulated values of $\hat{\varphi}^{OLS}$. These are sorted and used to compute the median and any other desired quantiles of the distribution. Since there is no unconditional distribution from which to initialize a simulated AR(1) process with $\varphi > 1$, it is assumed that $\varphi \leq 1$.

The monotone increasing blue line representing the median is consistently below the black 45 degree line representing the true AR(1) coefficient. The gap between the two grows as $\varphi$ increases, reflecting the increasing bias in the OLS estimator as the unit root is approached. The actual residuals gave $\hat{\varphi}^{OLS} = 0.904$. This matches the simulated median for the (AR(1)) Median Unbiased estimate of $\hat{\varphi}^{MU} = 0.950$. That is to say, for any $\varphi \leq 0.950$, the probability of obtaining an OLS estimate of 0.904 or higher is no more than 0.50, while for any $\varphi \geq 0.950$ the probability of obtaining 0.904 or lower is at least 0.50.

A 95% confidence interval for $\varphi$ can similarly be obtained by matching the .975 and .025 quantiles of the simulated distribution (in magenta) to the actual OLS point estimate. $\hat{\varphi}^{OLS} = 0.904$ is the .975 simulated quantile for $\varphi = 0.873$. It follows from the monotonicity of each of the quantile functions that for any $\varphi \leq 0.873$, the probability of obtaining .904 or higher is 0.025 or less. Furthermore, 0.873 is the smallest value for which this is true, so that 0.873 is the appropriate lower bound for a 95% confidence interval.

In the present illustration, matching the .025 quantile of the simulated distribution to the OLS estimate only occurs for $\varphi = 1$. The 95% confidence interval therefore is (.873, 1).

The simulated probability of $\hat{\varphi}^{OLS} \leq 0.904$ with $\varphi = 1$ is .207, so that a unit root in the errors cannot be rejected at even the .20 level, assuming for the moment an AR(1) structure. Note that this is in fact an exact finite sample test for cointegration, conditional on an AR(1) structure. Since a unit-root test is appropriately single-sided, a (positive) unit root can be rejected at say the 5% level if and only if unity lies outside the 90% CI for $\varphi$.

Although an AR(1) model is often a good first approximation to the autocovariation function, there is frequently evidence of higher order serial correlation. This may not be truly autoregressive, but an AR($p$) model can approximate any stationary $\Gamma$ to any desired precision, given a high enough value of $p$.

---

[9] Throughout this paper, GAUSS proc rndKMu was used to generate iid normal pseudo-random variables, together with the hexadecimal seed E41E9415, obtained by radioactive decay courtesy of Walker (2001). Using the same seed for each value of the parameter(s) ensures that each quantile is a smooth function of $\varphi$.

Andrews and Chen (1994) approximately median unbias a univariate AR($p$) model similar to[10]

$$\varepsilon_t = \sum_{j=1}^{p} \varphi_j \varepsilon_{t-j} + u_t ,\qquad(13)$$

by first rewriting it in the Augmented Dickey-Fuller (ADF) form

$$\varepsilon_t = \alpha \varepsilon_{t-1} + \sum_{j=1}^{p-1} \psi_j \left( \varepsilon_{t-j} - \varepsilon_{t-j-1} \right) + u_t ,\qquad(14)$$

where

$$\alpha = \sum_{j=1}^{p} \varphi_j , \quad \psi_j = -\sum_{h=j+1}^{p} \varphi_h .$$

They argue persuasively that the sum of the AR coefficients $\alpha$ is a more natural measure of persistence than the largest inverted root of the AR polynomial, since it determines the Cumulative Impulse Response of the model, CIR = $1/(1-\alpha)$. Furthermore, although $\alpha < 1$ is not the only necessary condition for stationarity, it is the one that is most likely to be an issue in econometric data. Since linear combinations of median unbiased estimators are not necessarily median unbiased, it does make some difference whether (13) or (14) is considered.

Andrews and Chen argue that the ADF-form AR($p$) model (14) cannot be exactly median unbiased by generalizing the method of Andrews (1993), because the distribution of the OLS estimator of each parameter depends not only on the parameter itself, but also on the true (and therefore unknown) values of the other $p$-1 parameters as well, and a single equation cannot be solved for $p$ unknowns. Consequently, they only univariately median-unbias their estimate of $\alpha$, conditional on OLS estimates of the $\psi_j$. Since the OLS estimates of the $\psi_j$ are biased as well (though ordinarily not to the same degree), their estimates are only approximately median unbiased.

However, there is in fact no obstacle to median-unbiasing all $p$ parameters of this model. For this purpose we first further rearrange (13) into what might be called Recursive ADF, or Persistence form, as follows:

$$\varepsilon_t = \sum_{j=1}^{p} \alpha_j \Delta^{j-1} \varepsilon_{t-1} + u_t ,\qquad(15)$$

where

---

[10] Andrews and Chen in fact include a constant and time trend term in a univariate model. The present proposal is concerned with the less restrictive case of a general OLS regression. If the regression includes nothing but a constant and time trend, the results of the present two-step procedure will be very similar, if not quite identical, to those of their one-step procedure.

$$\boldsymbol{\alpha} = \mathbf{A}\boldsymbol{\varphi},$$
$$\boldsymbol{\alpha} = (\alpha_1, \ldots \alpha_p)',$$
$$\boldsymbol{\varphi} = (\varphi_1 \ldots \varphi_p)', \tag{16}$$
$$\mathbf{A} = \left( \binom{j-1}{i-1}(-1)^{i-1} \right),$$

and $\Delta^j$ indicates the j-th difference operator $(1\text{-}L)^j$. Note that $\mathbf{A} = \mathbf{A}^{-1}$, so that $\boldsymbol{\varphi} = \mathbf{A}\boldsymbol{\alpha}$ as well. We will refer to $\alpha_j$ as the j-*th order persistence*.

The first order persistence $\alpha_1$ of (15) is identically equal to the Andrews-Chen persistence $\alpha$ of (14). When this equals unity, (14) becomes a *p*-1 order autoregression in first differences. The persistence of this new autoregression (and therefore its greatest potential source of bias) is then the sum of the $\psi_j$'s, which is simply our second order persistence $\alpha_2$, and so forth. The *p*-th order persistence $\alpha_p$ identically equals the *p*-th order autoregressive coefficient $\varphi_p$ times $(-1)^{p+1}$, so that $\varphi_p = 0$ if and only if $\alpha_p = 0$.

Let $m_i(\boldsymbol{\alpha})$ be the Monte-Carlo median of the OLS estimate of $\alpha_i$ obtained from regression residuals generated by (4), with errors simulated using $\boldsymbol{\alpha}$. Equating these *p* functions to the *p* actual OLS estimates of $\alpha_j$, we in fact have *p* equations in *p* unknowns that can readily be solved for $\hat{\boldsymbol{\alpha}}^{MU}$, using a simple iterative procedure. Since $\mathbf{A} = \mathbf{A}^{-1}$, we then have $\hat{\boldsymbol{\varphi}}^{MU} = \mathbf{A}\hat{\boldsymbol{\alpha}}^{MU}$ as "median-unbiased" estimates of $\boldsymbol{\varphi}$, in the sense that they correspond to truly median-unbiased estimates of $\boldsymbol{\alpha}$.

An exact 100(1-*p*)% confidence interval for each parameter (conditional on median-unbiased estimates of the other parameters) can be found by solving its *p*/2 and 1-*p*/2 quantile functions simultaneously with the median functions of the other *p*-1 parameters.[11] Likewise, an exact size *p* test for cointegration (again conditional on median-unbiased estimates of $\alpha_2, \ldots \alpha_p$) can be obtained by solving the 1-*p* quantile function for $\alpha_1$ together with the median functions of the other parameters.

The autocovariation function may be estimated consistently by considering values of the autoregressive order *p* up to and including a value such as
$$p_{\max} = \text{round}(.75\, T^{1/3}), \tag{17}$$
the "benchmark" formula (7) recommended by Stock and Watson (2007) for HAC.[12] Parsimony may be enforced with a general-to-specific model selection procedure that starts with $p = p_{max}$ and tests the hypothesis $\alpha_p = \varphi_p (-1)^{p+1} = 0$ as described above at

---

[11] The median-unbiased values taken on by the other parameters will ordinarily be different at the upper and lower confidence bounds.

[12] As noted in Footnote 3 above, the formula $m = \text{floor}(4\,(T/100)^{2/9}) + 1$ has very similar effect for *T* in the range 50 – 2000.

some appropriate test size, say .05, sequentially reducing *p* by 1 if the hypothesis cannot be rejected.[13]

## 4.  MUAR-AC Covariance Matrix Estimation

Once the autoregressive order has been determined and $\hat{\alpha}^{MU}$ is found, the sum of squared residuals of (15) may be used to estimate var($u_t$), and the Yule-Walker equations solved for $\hat{\Gamma}^{MU}$.  The Median-Unbiased Autoregressive, Autocorrelation Consistent (MUAR-AC) covariance matrix is then

$$\hat{\mathbf{C}}^{MUAR-AC} = (\mathbf{X'X})^{-1}\mathbf{X'}\hat{\Gamma}^{MU}\mathbf{X}(\mathbf{X'X})^{-1}.$$

The MUAR-AC estimates could in principle be used to iteratively re-estimate the regression by Feasible Generalized Least Squares (FGLS).  However, Hayashi (2000: 59) warns that "Very little is known about the finite-sample properties of the FGLS estimator."  It therefore may be safer, when the correlation structure must be estimated, to stick with the OLS regression coefficients, and to only use the estimated correlation structure to adjust the standard errors, as in the Newey-West procedure.

## 5.  Money Demand

The demand for money has long been of concern to macroeconomists, but the search for it has been plagued by the problems of serial correlation and near-unit-root dynamics in the residuals.  Meltzer (1963), for example, early-on estimated a plausible U.S. money demand equation by OLS, but Courchene and Shapiro (1964) quickly demonstrated that Meltzer's equation suffered from severe positive serial correlation and hence that his standard errors were greatly understated and his t-statistics comparably overstated.  Goldfeld (1973) and Hallman, Porter and Small (1991) add one or more lags of their respective dependent variables in an effort to reduce or eliminate serial correlation, but these lag coefficients also suffer from AR bias.  Point estimates of the sum of the lag coefficients that appear to be different from unity by the standard OLS t-test may in fact be masking a spurious regression.

The present paper illustrates the use of the MUAR-AC covariance matrix by applying it to OLS estimates of the postwar US money demand function.  For this purpose, money is measured as "M1-S", the sum of official M1 plus estimated retail sweep accounts.[14]

Prior to 1995, official M1 consisted of currency in circulation plus all checking accounts (plus a negligible quantity of non-bank travelers' checks), and thus represented

---

[13]  This procedure does not consistently estimate the autoregressive *order* of a pure AR(*p*) process, since there is always a 5% chance that it will be too high.  Nevertheless, MUAR will consistently estimate each coefficient, even if the order is unnecessarily high.

[14]  Dutkowsky and Cynamon (2003) apply the term "M1-S" to a similar aggregate, which is a non-additive combination of official M1 and estimated sweep accounts.  However, as Anderson (2003) points out, sweep accounts and M1 checking accounts are equivalent from the depositors' point of view, and therefore should be aggregated one-for-one.

"narrow money." Since 1995, however, official M1 has omitted the rapidly growing portion of checking accounts known as "Retail Sweep Accounts" (Anderson 2002). Official M1 is therefore a meaningful measure of narrow money only up to 1995. Since 1995, data is available from the Federal Reserve Bank of St. Louis for estimated cumulative conversions to sweeps. This gives at least a rough idea of the magnitude of total sweep accounts.[15] Figure 2 shows, on a dollar scale, monthly official M1 (SA), estimated sweep accounts S, and their sum, M1-S. This series was used to compute m = Real M1-S, averaged quarterly and deflated by the GDP deflator (SA).[16]
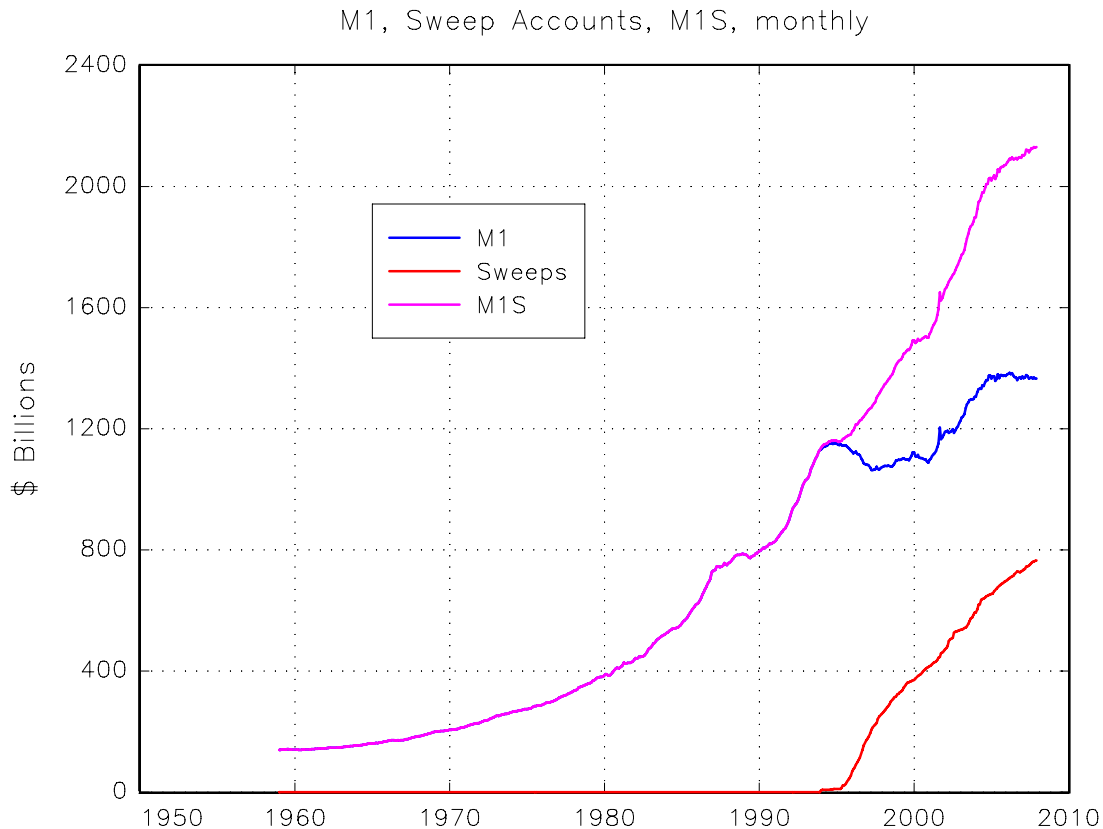


**Figure 2**
Estimated Sweep Accounts (red), Official M1 (SA, blue),
and their sum, M1S (magenta), monthly.

Figure 3 depicts M1-S velocity, $V = y/m$, in units of $yr^{-1}$, computed from m = Real M1-S and y = real GDP (SA, quarterly, not shown). It may be seen that velocity

---

[15] Estimated sweeps data was acquired from St. Louis Fed Research Dept., <http://research.stlouisfed.org/aggreg/swdata.html>. It is unfortunately to be expected that the quality of the measure may deteriorate toward the end of the period, since the sweep account proxy neglects any growth that has occurred in sweep accounts since their conversion.

[16] Data other than estimated Sweeps were obtained from FRED (<http://research.stlouisfed.org/fred2/>): series GDPDEF (GDP Implicit price deflator, SA), GDPC96 (real GDP in billions of chained 2000 dollars, SA annual rate), M1SL (BOG monthly M1, SA), and TB3MS (3-month Treasury Bill Rate, secondary market, NSA, monthly, averaged here to quarterly).

rose fairly steadily to 1981. Since then it has generally declined but at an uneven rate, with local peaks around 1984, 1990, 2000, and 2007. Despite the general decline, it has never fallen below its pre-1975 levels.
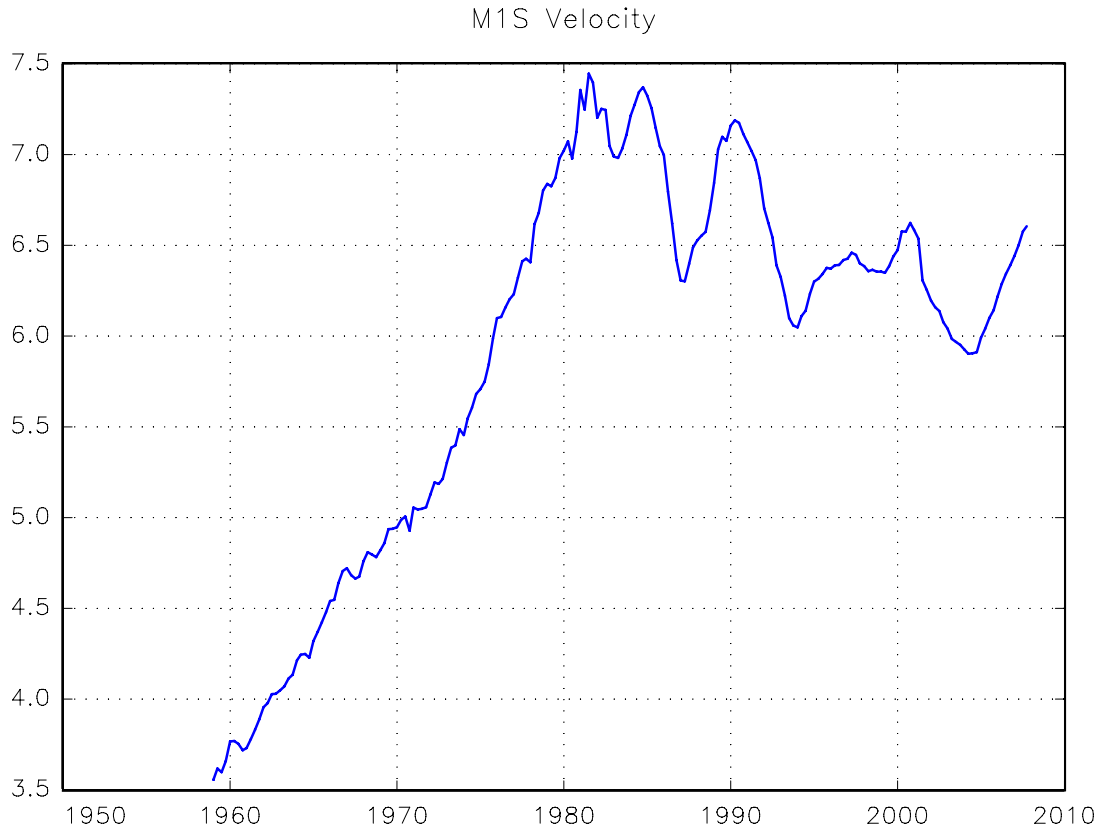


**Figure 5**
M1S Velocity, $V = y/m$ (yr$^{-1}$), quarterly

The 3-month Treasury Bill rate R (not shown) has an uptrend to 1981 and subsequent downtrend with numerous vacillations, including local peaks near 1984, 1990, 2000, and 2007. Its behavior may help account for the fluctuations in M1-S velocity, though there are apparently economies of scale in money holding that account for the residual uptrend in this velocity measure.

It is assumed that observed real M1-S money balances $m_t$ equal desired real money balances plus (in logs) an error $\varepsilon_t$ that is normal but may be serially correlated. Desired real money balances are assumed to have a constant real income elasticity $a$, and a constant interest semielasticity $b$:

$$\log m_t = c + a \log y_t + b\, R_t + \varepsilon_t. \tag{18}$$

The regressors are assumed for the purposes of the present study to be exogenous and thus independent of the errors, though it may be appropriate to revisit this assumption in future work.

The regressors are first normalized by subtracting their terminal (2007 Q4) values, so that the intercept $c$ represents estimated log real money demand for this terminal quarter. With this normalization, the t-statistic on the intercept is not interesting, but its standard error then directly indicates the precision with which equilibrium 2007Q4 M1-S demand is determined.

Equation (18) was fit by OLS with data for 1959Q1 – 2007Q4 (n = 196 observations). Columns 2-4 of Table 2 shows the OLS estimates, standard errors, and t-statistics. The OLS t-statistics for the regressors are offscale for significance, while the intercept results appear to indicate that equilibrium 2007Q4 M1-S demand is known to within a s.e. of only 1%.

**Table 2**
OLS coefficient estimates of (18),
with OLS and Newey-West standard errors and t-statistics.
(196 observations, 1959Q1 – 2007Q4)

| coef. | OLS | | | Newey-West | |
| --- | --- | --- | --- | --- | --- |
| | est. | s.e. | t-stat | s.e. | t-stat |
| $c$ | 7.4426 | 0.0104 | ------- | 0.0185 | ------- |
| $a$ | 0.6621 | 0.0114 | 58.24 | 0.0195 | 33.95 |
| $b$ | -0.0403 | 0.0019 | -21.25 | 0.0032 | -12.52 |
| DW | 0.1835 | | | | |

Unfortunately, the Durbin Watson statistic DW is also offscale for significance, resoundingly indicating at least first order serial correlation. Figure 4 depicts the OLS residuals. These obviously indicate serially correlated errors. The residuals do appear to be mean-reverting, but this may simply be an artifact of the circumstance that the errors themselves have been demeaned by the constant term in the regression, virtually detrended by the inclusion of log(y), and then further manipulated by the inclusion of R in the regression, as governed by (4). The appearance of stationarity may therefore be illusory, and the regression therefore spurious.
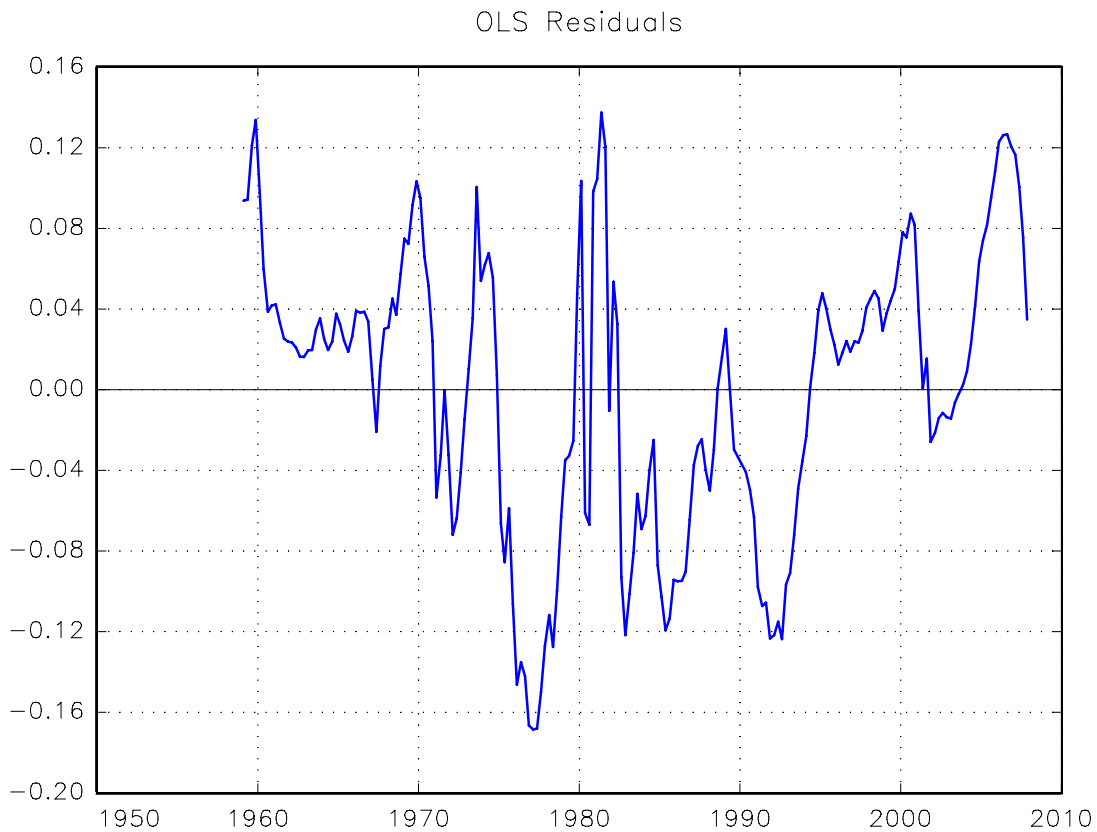
**Figure 4**
OLS residuals of equation (18).

The last two columns of Table 2 give the Newey-West HAC standard errors and t-statistics for the three coefficients, using a bandwidth of m = 4 as computed from the Stock and Watson "benchmark" formula (7). The standard errors are 68-78% larger than the OLS standard errors, yielding t-statistics that are proportionately smaller, but still in the double digits.

The blue line in Figure 5 shows the autocorrelations of the residuals in Figure 4, for lags up to 16 quarters. These die virtually to zero by lag 12. However, these may be biased downwards, for much the same reason autoregressive coefficients are biased in small samples.
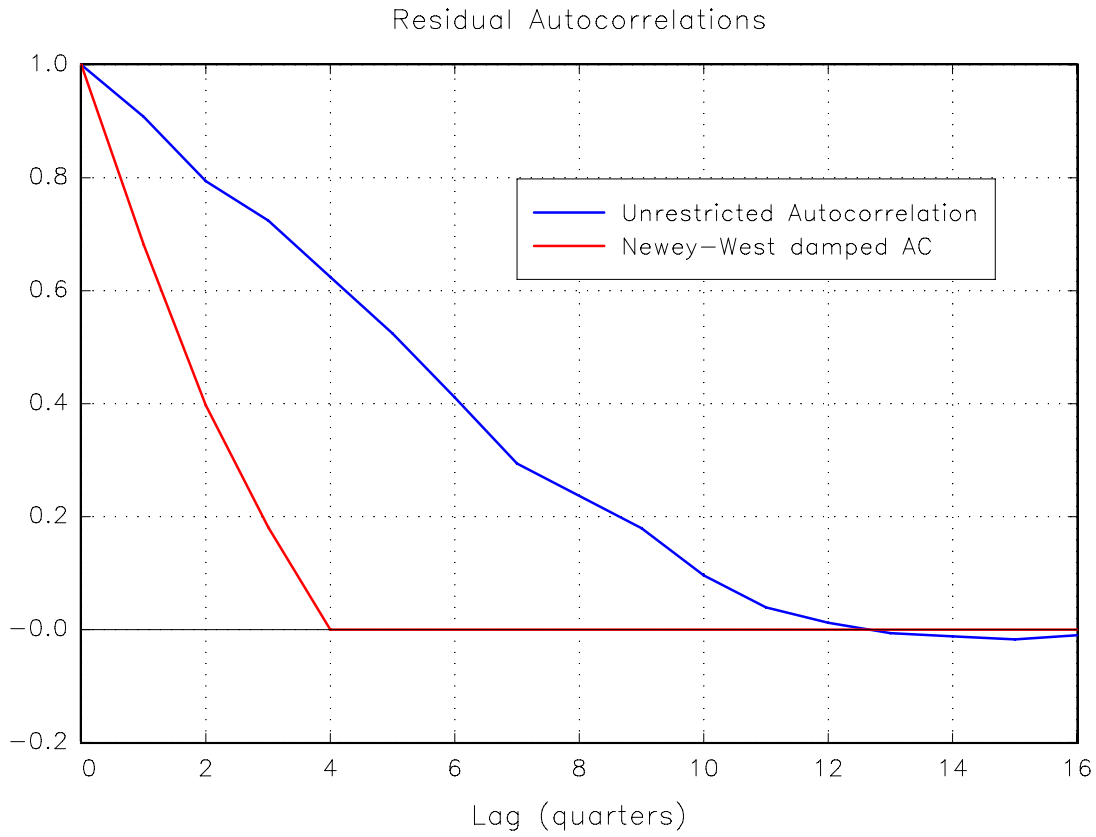
Residual Autocorrelations



**Figure 8**
Autocorrelations of the residuals of equation (18) (blue),
with Newey-West attenuated autocorrelations (red).

      The red line in Figure 5 indicates the estimated autocorrelations, after attenuation by the Bartlett kernel factor (6) employed in the Newey-West estimator. With bandwidth $m = 4$, the autocorrelations of order 4 and higher are all assigned weight 0 and therefore ignored altogether, while the lower order autocorrelations are greatly attenuated. With an infinite sample, the bandwidth would become infinite and the attenuation factors would all be unity, so that the HAC standard errors do consistently estimate the true standard errors. However, with this or almost any other finite sample, they will greatly understate the true uncertainty of the parameters.

      Our MUAR-AC estimator requires that we first compute $p_{max} = 4$ using (17), and then consider AR(p) models beginning with this order and working downwards. Column 2 of Table 3 below gives the OLS estimates of the standard-form AR(4) coefficients $\hat{\varphi}_j^{OLS}$ computed from the OLS residuals of (18). Column 3 shows the corresponding persistence- or recursive ADF-form coefficients $\hat{\alpha}_j^{OLS}$.

**Table 3**
OLS and Median-Unbiased estimates of AR(4) coefficients.

| lag $j$ | $\hat{\varphi}_j^{OLS}$ | $\hat{\alpha}_j^{OLS}$ | $\hat{\alpha}_j^{MU}$ | 95% CI | | $\hat{\varphi}_j^{MU}$ |
|---|---|---|---|---|---|---|
| | | | | lower | upper | |
| (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| 1 | 1.140 | 0.884 | 0.920 | 0.848 | 1.000 | 1.152 |
| 2 | -0.503 | 0.292 | 0.242 | -0.085 | 0.467 | -0.499 |
| 3 | 0.528 | -0.317 | -0.288 | -0.622 | -0.008 | 0.545 |
| 4 | -0.282 | 0.282 | 0.278 | 0.130 | 0.414 | -0.278 |

Column 4 of Table 3 gives the median-unbiased persistence-form estimates $\hat{\alpha}_j^{MU}$, found by simultaneously matching the median functions for each coefficient to the OLS estimates in column 3, using 1000 Monte Carlo replications for each trial value.[17] Although the MU estimate of $\alpha_1$ is only 0.036 above the OLS estimate, this reflects a 45% increase in persistence, as the Cumulative Impulse Response increases from 1/(1-0.884) = 8.62 to 1/(1-0.920) = 12.5.

Column 7 of Table 4 gives the median-unbiased standard-form estimates $\hat{\varphi}_j^{MU}$ derived from the $\hat{\alpha}_j^{MU}$ using (16). Although their sum has increased by the same substantial 0.036, this increase is spread almost unnoticeably over the four coefficients.

Columns 5 and 6 of Table 4 provide the lower and upper bounds of a 95% confidence interval for each $\alpha_j$. The lower bound of 0.130 for $\alpha_4$, for example, was found by simultaneously matching the simulated median functions for $\alpha_1$, $\alpha_2$ and $\alpha_3$ together with the .975 quantile function for $\alpha_4$ to the OLS estimates in column 3. Assuming that each quantile for $\hat{\alpha}_j^{OLS}$ is monotone increasing in the true value of $\alpha_j$, holding the other coefficients constant at their true value as proxied by their median values, the probability is thus .025 or less that $\hat{\alpha}_4^{OLS}$ could be as high as 0.282 for any $\alpha_4 \leq 0.130$. Likewise the upper bound of 0.414 for $\alpha_4$ means that the probability is .025 or less that $\hat{\alpha}_4^{OLS}$ could be as *low* as 0.282 for any $\alpha_4 \geq 0.414$.

Note that the lower and upper bounds of the 95% CI are based on two separate one-tailed tests at the .025 level, using estimates of the first 3 parameters that are potentially different from one another as well as from the MU estimates in column 4. Thus for the lower bound, the four estimates are (0.920, 0.243, -0.276, -.130), while for the upper bound they are (0.920, 0.243, -0.290, 0.414).[18]

---

[17] The entire estimation, with 1000 replications for each trial set of parameters, a tolerance of .0001, and including confidence intervals for 4 different confidence levels (not all tabulated) takes about 90 seconds in GAUSS on a recent PC. It is important that the same seed be used for each trial value, so that the simulated median function will be a smooth function of the trial parameters.

[18] Since the first two coefficients are essentially unchanged and the third varies only slightly in this example, it may in general be adequate simply to use the MU point estimates of the nuisance parameters when calculating the confidence intervals.

Since the 95% CI for $\alpha_4$ well excludes 0, an AR(3) model may easily be rejected at the .05 level. A 2-tailed test is appropriate here, since if $\alpha_4$ is not 0, it could equally well be on either side of 0. Our general-to-specific model selection procedure therefore retains the AR(4) model.

As it happens, in an AR(2) model (not tabulated), $\alpha_2$ is insignificantly different from 0 at the 5% level. A specific-to-general model selection procedure starting with the AR(1) model of Figure 1 would therefore stop with AR(1). However, since AR(2) is a special case of the rejected and therefore presumably false AR(3) model, this test is invalid.

Since unity is included in the 95% CI for $\alpha_1$ in Table 3, a unit root cannot be rejected at the .025 test size. In the case of a unit root, a one-tailed test is appropriate, since we are not considering values of the persistence greater than unity.[19] As it happens, a unitary value for $\alpha_1$ cannot be rejected at the more conventional .05 test size (1-tail), either. However, it *can* be rejected at the .10 test size, so that we may say that we have at least weakly rejected a unit root.[20]

So long as a unit root can be rejected at even the 50% level (i.e. so long as $\hat{\alpha}_1^{MU} < 1$), $\hat{\Gamma}^{MU}$ and thence $\hat{C}^{MUAR-AC}$ may be computed and used to calculate standard errors for the money demand coefficients themselves.

Table 4 gives standard errors for the three money demand coefficients, using both the OLS AR(4) coefficients of column 2 of Table 3, and the MUAR-AC coefficients of column 7 of Table 3, along with t-statistics computed from the MUAR-AC standard errors. For comparison, the OLS and Newey-West HAC standard errors are repeated from Table 1.

**Table 4**
OLS coefficient estimates of (18), with OLS, Newey-West,
AR(4) and MUAR-AC standard errors, and MUAR-AC t-statistics.
(196 observations on M1-S, 1959Q1 – 2007Q4)

| coef. | OLS | | NW | AR(4) | MUAR-AC | |
|---|---|---|---|---|---|---|
| | est. | s.e. | s.e. | s.e. | s.e. | t-stat. |
| $c$ | 7.4426 | 0.0104 | 0.0185 | 0.0334 | 0.0468 | ------ |
| $a$ | 0.6621 | 0.0114 | 0.0195 | 0.0364 | 0.0509 | 13.00 |
| $b$ | -0.0403 | 0.0019 | 0.0032 | 0.0055 | 0.0073 | -5.52 |

It may be seen that untruncated and unattenuated AR(4) standard errors are already 72-87% higher than the HAC standard errors. However, these are based on

---

[19] In fact, there is no way to initialize the simulation of such a process, since it must start infinitely far from its fixed point. An exactly unit root process, on the other hand, has no fixed point, and so one starting point is as good as another when a constant is included in the regression.

[20] In fact, $\alpha_1 = 1$ is only one of several ways a unit root could occur in an AR(4) model. Nevertheless, $\alpha_1 < 1$ is the necessary condition for stationarity most likely to be violated in econometric data.

median-biased AR coefficients. The median-unbiased MUAR(4) standard errors are another 33-40% higher than the AR(4) standard errors, and thus 128-153% higher than HAC, not to mention 284-350% higher than OLS. Equivalently, the HAC t-statistics of Table 1 are 128-153% too high, while the OLS t-statistics are 284-350% too high.

Figure 6 shows the unrestricted autocorrelations of Figure 5, alongside the autocorrelations implied by the OLS AR(4) and MUAR(4) coefficients of Table 3. Since the OLS AR(4) coefficients simply match the first four autocorrelations as computed from almost the same set of points as the unrestricted autocorrelations, these match almost exactly out to lag 4. Although higher order than $p_{max} = 4$ was not considered, AR(4) does happen to match the other unrestricted autocorrelations out to lag 16 fairly well. Beyond lag 16, the AR(4) autocorrelations continue to decay to zero, whereas the unrestricted autocorrelations (not plotted) flop persistently but aimlessly on either side of zero.[21] The MUAR(4) autocorrelations are naturally much more persistent than the AR(4) autocorrelations, because $\hat{\alpha}_1^{MU} > \hat{\alpha}_1^{OLS}$.
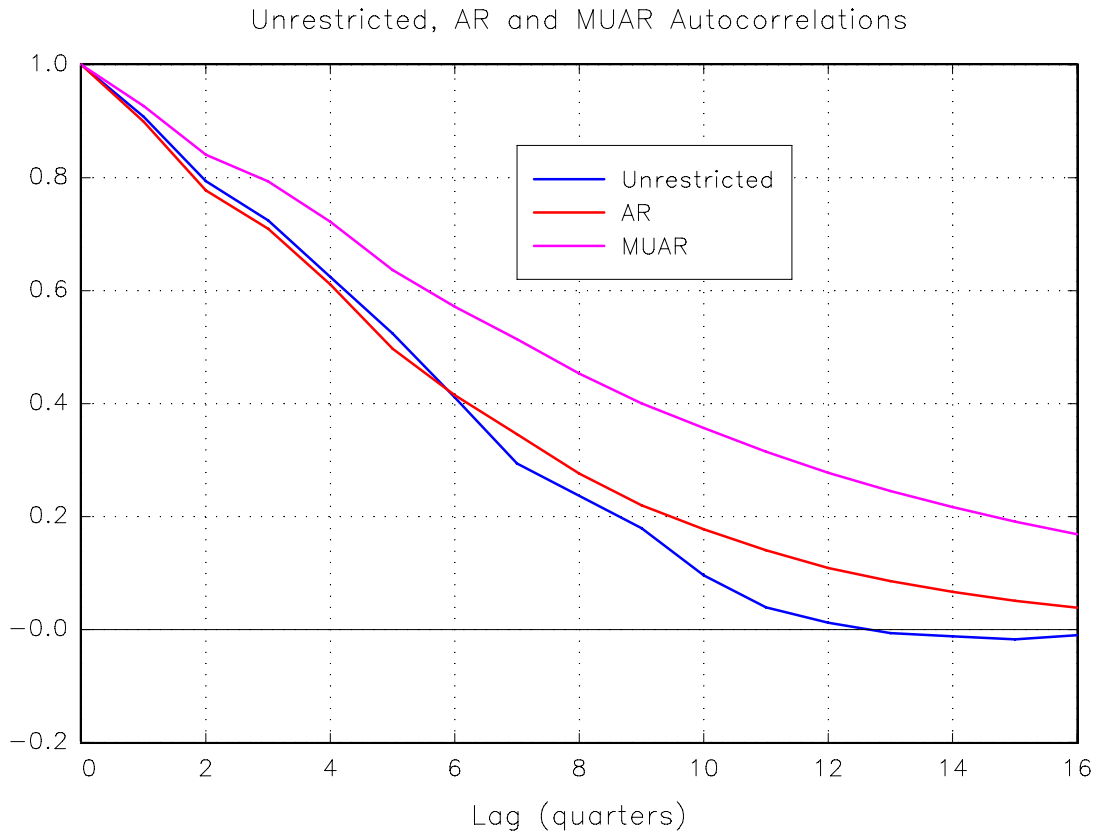


**Figure 9**
Unrestricted, AR(4) and MUAR(4) autocorrelations.

---

[21] Since the lower order autocorrelations are non-zero and large, it is to be expected that the estimated higher order autocorrelations will be persistent, and frequently lie outside the bounds for autocorrelations based on the null that all the autocorrelations are zero.

The greater persistence of the MUAR(4) autocorrelations illustrated in Figure 6 is only one of two reasons why the MUAR standard errors in Table 4 are higher than the AR(4) standard errors. The second reason is illustrated in Figure 7, which shows the actual unrestricted, AR(4) and MUAR(4) autocovariance estimates, i.e. the autocorrelations of Figure 9 times their respective estimates of $\gamma_0$. The unrestricted and AR(4) estimates of $\gamma_0$ are somewhat different, because the one is computed directly from the sum of squared residuals $\hat{\varepsilon}_t$, while the other is reconstructed from the sum of squared OLS residuals of the autoregression $\hat{u}_t^{OLS}$, together with the autoregressive coefficients. The MUAR(4) estimate of $\gamma_0$ is much higher than either of these, both because it is based on the MUAR autoregressive residuals $\hat{u}_t^{MU}$, whose sum of squares is necessarily greater than that of the Least Squares (LS) AR(4) residuals, and because the MUAR autocorrelations are more persistent than the LS AR residuals.
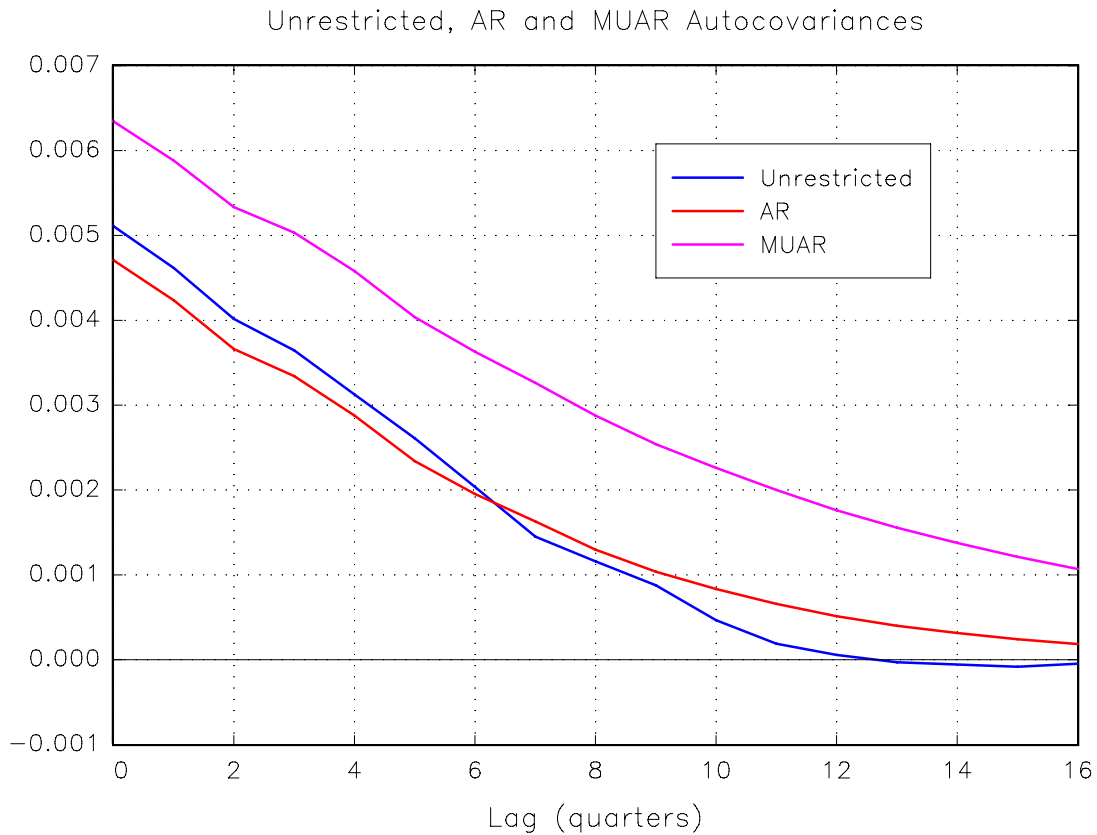


**Figure 10**
Unrestricted, AR(4), and MUAR-AC autocovariances.

Even though the MUAR t-statistics of Table 4 are greatly reduced in comparison to OLS or even HAC, they are still offscale for significance. The M1-S money demand income elasticity and interest semi-elasticity are therefore very well defined by the data, with the appropriate signs. The intercept, and therefore the equilibrium level of 2007Q4 M1-S demand, is accurate to approximately 4.7%. This suggests that M1-S may be a

useful, if inexact, indicator of monetary policy, and that the Fed should make every effort to collect true data on sweep accounts.[22]

## 6. Possible Extension to Regressor-Conditional Heteroskedasticity

Heteroskedasticity (see McCulloch 1985b) is a distinct problem from autocorrelation, but one that also frequently arises in econometric models. The present section discusses possible extensions of the MUAR-AC to incorporate *Regressor-Conditional Heteroskedasticity*.

Heteroskedasticity may be of (at least) four types: A Priori, External, Sequential, and Regressor-Conditional. *A Priori Heteroskedasticity* arises when the variance is known or may be computed from observable variables, up to an unknown constant of proportionality. Weighted Least Squares (WLS) provides efficient coefficient estimates with correct standard errors.

In *External Heteroskedasticity*, the variances are some unknown function of variables not necessarily included in the regression. Packages such as E-Views estimate this function from a user-provided list of variables, and then estimate the regression by WLS.[23]

In the very common problem of *Sequential Heteroskedasticity*, the absolute values of sequential errors are correlated, without regard to the regressors or external variables. Models that incorporate such "volatility clustering" include Autoregressive Conditional Heteroskedasticity or ARCH (Engle 1982), Generalized ARCH or GARCH (McCulloch 1985a; Bollerslev 1986), and the Local Scale Model (Shephard 1994; McCulloch 2006).

In *Regressor Conditional Heteroskedasticity* (RCH), the variances are some unknown function of the regressors themselves. This is not necessarily the most important or most prevalent form of heteroskedasticity, but is the form addressed by the famous White (1980) Heteroskedasticity Consistent Covariance (HCC) matrix that has become intimately intertwined with the Autocorrelation Consistent covariance matrix literature.[24]

White (1980) provides a test for RCH in the absence of autocorrelation based on the $R^2$ of a regression of the squared errors on a generalized quadratic function of the regressors that provides a second order approximation to the unknown true function.

---

[22] Because a unit root can only weakly be rejected (at the 10% level), these results should be interpreted with caution, and a model with explicitly time-varying coefficients such as Adaptive Least Squares (McCulloch 2005) might be more appropriate. This possibility goes beyond the scope of the present paper.

[23] It would be even easier to simply estimate the regression consistently by OLS, and then to adjust the standard errors of the OLS coefficient estimates for any external heteroskedasticity detected in the OLS residuals, in the spirit of NW-HAC and the present paper. Unfortunately, EViews does not offer this option.

[24] Andrews and Monahan (1992) do find that neglecting RCH can lead to substantial distortion of coverage, but only in simulations in which the RCH is of an artificially extreme form.

This test could readily be generalized to account for the presence of autocorrelation, as quantified by $\hat{\Gamma}^{MU}$. However, this could be done either by specifying that the variance of the regression errors $\varepsilon_t$ is a quadratic function of the regressors, or by specifying that the variance of the innovations $u_t$ is a quadratic function of the regressors. Future research should consider both approaches.

The ingenious Heteroskedasticity Consistent Covariance (HCC) matrix of White (1980)

$$\hat{\mathbf{C}}^{HCC} = (\mathbf{X'X})^{-1}\mathbf{X'DX}(\mathbf{X'X})^{-1}$$

where

$$\mathbf{D} = (diag(\mathbf{e}))^2$$

corrects asymptotically for RCH in the absence of serial correlation. Although RCH is a completely different, and less common, issue than serial correlation, a similar correction for RCH is incorporated into the HAC estimator of Newey and West (1987), so that the two literatures have become closely intertwined. MUAR-AC could be modified with a White-type adjustment to take RCH into account (MUAR-HAC), if and when it is detected, as follows:

$$\hat{\mathbf{C}}^{MUAR-HAC} = (\mathbf{X'X})^{-1}\mathbf{X'GX}(\mathbf{X'X})^{-1}$$

where

$$\mathbf{G} = \left(e_t e_{t'} \hat{\gamma}^{MUAR}_{|t-t'|} / \hat{\gamma}_{|t-t'|}\right)$$

This is not implemented in the present paper, but is left for future research.

The VAR-HAC estimator of den Haan and Levin (1997) incorporates a consistent autoregressive adjustment for serial correlation in the spirit of the present proposal, along with a VAR-based adjustment for RCH, whether or not RCH is present. However, they make no attempt to remove the bias in their AR coefficients, and their method can easily be parametrically extravagant. By focusing on the autocorrelation problem first, the MUAR-AC estimator eliminates the AR bias while at the same time being far more parsimonious. If RCH is found to be present, MUAR-HAC would remove the AR bias and still be far more parsimonious than VAR-HAC.

Yet another approach is the very interesting Heteroskedasticity and Autocorrelation Inconsistent (HAI) covariance estimator of Kiefer, Vogelsang and Bunzel (2000), as modified by Kiefer and Vogelsang (2002):

$$\hat{\mathbf{C}}^{HAI} = c(\mathbf{X'X})^{-1}\mathbf{X'}B\mathbf{X}(\mathbf{X'X})^{-1},$$

where

$$B = \left(e_t e_{t'}(T - |t - t'|)/T\right)$$

and $c$ is a correction for bias. The computations of Abadir and Paruolo (1996, 2002) demonstrate that in terms of the above 2002 version of the statistic (which differs by a factor of 2 from that in the 2000 paper), c = 5.588756592. While HAI makes no parametric assumptions about the form of the serial correlation, it is inconsistent in that the limiting covariance estimate is merely an unbiased random variable, rather than a limit in probability. When $t$ statistics are computed using these unbiased variances, Abadir and Paruolo's calculations imply that the 5% critical value for $|t|$ is 2.02,

approximately as expected.  However, the 1% critical value is 3.05, which is much higher than the conventional 2.58.  The HAI standard errors for the three coefficients of Table 4 are .0706, .0862, and .0041, respectively, which are sometimes higher and sometimes lower than the MUAR-AC estimates.

## 7.  Conclusion

It has been shown that the Newey-West (1987) Heteroskedasticity and Autocorrelation Consistent (HAC) covariance matrix estimator can greatly understate the standard errors of OLS regression coefficient estimates in finite samples, and therefore comparably overstate t-statistics.  Although the bias vanishes in infinite samples and is tolerable in samples as small as $10^6$, it can lead to t-statistics that are too high by a factor of 1.7-2.2 with a sample size of 65-1200 and first order autoregressive serial correlation with AR coefficient 0.9.

The present paper extends the Exactly Median Unbiased estimator of Andrews (1993) for a *directly* observed AR(1) process to the case of an AR(p) process that is only *indirectly* observed via OLS regression residuals.  By allowing the maximum permitted order to increase without limit with the sample size, the estimator consistently estimates a stationary process with any autocovariance function.  It also provides a unit root test (and therefore a test for cointegration of the regressors) that is exact up to median unbiased estimates of the higher order persistences.

These Median Unbiased Autoregressive (MUAR) estimates of the autocovariance function are then used to construct an Autocorrelation Consistent (MUAR-AC) covariance matrix for the OLS coefficient estimates.  Applied to a simple model of US demand for narrow money M1-S (official M1 + estimated Sweep Accounts), it is found that a unit root in the errors and therefore absence of cointegration can be at least weakly rejected (at the 10% test size).  The MUAR-AC standard errors are 128-153% higher than HAC standard errors, or equivalently, HAC t-statistics for any hypothesis concerning the coefficients are 128-153% too large.

Despite the greatly increased MUAR-AC standard errors, the income elasticity and interest semielasticity of demand for M1-S remain highly significant, suggesting that narrow money (including all checking accounts) may still be a useful indicator of monetary policy, despite the Fed's reluctance to collect direct data on it.

# REFERENCES CITED

Abadir, Karim M., and Paolo Paruolo. 1996. Two Mixed Normal Densities from Cointegration Analysis. *Econometrica* 65:671-80.

_____. 2002. Simple Robust Testing of Regression Hypotheses: A Comment. *Econometrica* 70:2097-9.

Anderson, Richard. 2002. Retail Sweep Programs and Money Demand. *Monetary Trends* (St. Louis Fed), Nov.. Online at http://research.stlouisfed.org/publications/mt/20021101/cover.pdf.

_____. 2003. Retal Deposit Sweep Programs: Issues for Measurement, Modeling and Analysis. Federal Reserve Bank of St. Louis Working Paper 2003-026A. Online at http://research.stlouisfed.org/wp/2003/2003-026.pdf.

Andrews, Donald W.K.. 1993. Exactly Median-Unbiased Estimation of First Order Autoregressive / Unit Root Models. *Econometrica* 61:139-65.

Andrews, Donald W.K., and Hong-Yuan Chen. 1994. Approximately Median-Unbiased Estimation of Autoregressive Models. *Journal of Business & Economic Statistics* 12:187-204.

Andrews, Donald W.K., and J. Christopher Monahan. 1992. An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator. *Econometrica* 60:953-66.

Beach, Charles, and James G. MacKinnon. 1978. Full Maximum Likelihood Estimation of Second Order Autoregressive Error Models. *Journal of Econometrics* 7:187-98.

Bollerslev, Tim. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *J. Econometrics* 31:1-50.

Courchene , Thomas J., and Harold T. Shapiro. 1964. The Demand for Money: A Note from the Time Series. *Journal of Political Economy* 72:498-503.

Dutkowsky, Donald H., and Barry Z. Cynamon, 2003. Sweep Programs: The Fall of M1 and Rebirth of the Monetary Aggregates. *Journal of Money, Credit and Banking* 35: 263-79.

Den Haan, Wouter J., and Andrew Levin. 1997. A Practitioner's Guide to Robust Covariance Matrix Estimation. In G.S. Maddala and C.R. Rao, eds., *Robust Inference* (*Handbook of Statistics* 15). Elsevier.

Durbin, James, and G.S. Watson. 1950. Testing for Serial Correlation in Least Squares Regression I. *Biometrica* 37:409-178.

Engle, Robert F. 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflations. *Econometrica* 50:987-1007.

Goldfeld, Stephen. 1973. The Demand for Money Revisited. *Brookings Papers on Economic Activity* 4:577-646.

Greene, William H. 2003. *Econometric Analysis*, 5th ed. Prentice Hall.

Hallman, Jeffrey J., Richard D. Porter, and David H. Small. 1991. Is the Price Level Tied to the M2 Monetary Aggregate in the Long Run? *American Economic Review* 81:841-58.

Hayashi, Fumio. 2000. *Econometrics*. Princeton Univ. Press.

Intergovernmental Panel on Climate Change. 2007. *IPCC Fourth Assessment Repor*t. <www.ipcc.ch>

Kiefer, Nicholas M., and Timothy J. Vogelsang. 2002. Heteroskedasticity-Autocorrelation Robust Standard Errors using the Bartlett Kernel without Truncation. *Econometrica* 70: 2093-5.

Kiefer, Nicholas M., Timothy J. Vogelsang, and Helle Bunzel. 2000. Simple Robust Testing of Regression Hypotheses. *Econometrica* 68: 695-714.

McCulloch, J. Huston. 1985a. Interest-Risk Sensitive Deposit Insurance Premia: Stable ACH Estimates. *J. Banking and Finance* 9: 137-156.

_____. 1985b. On Heteros*edasticity. *Econometrica* 53: 483.

_____. 2005. The Kalman Foundations of Adaptive Least Squares. Ohio State University Economics Dept. Working Paper 05-01. Presented at Econometric Society North American Summer Meetings, Duke University, June 2007. Online at <http://econ.ohio-state.edu/jhm/papers/KalmanAL.pdf>.

_____. 2006. Learning about Volatility: The Local Scale Model with Homoskedastic Innovations, with Application to Stock Returns. Presented at 12th International Conference on Computing in Economics and Finance, Limassol, Cyprus, June 2006. Online at < http://www.econ.ohio-state.edu/jhm/papers/ACH.pdf>.

_____. 2007. Adaptive Least Squares Estimation of the Time-Varying Taylor Rule. Presented at 13th International conference on Computing in Economics and Finance, HEC Montreal, June 2007. Online at <http://www.econ.ohio-state.edu/jhm/papers/TaylorALS.pdf>.

Meltzer, Allen. 1963. The Demand for Money: The Evidence from the Time Series. *Journal of Political Economy* 71:219-46.

Newey, Whitney, and Kenneth West. 1987. A Simple Positive Semi-Definite Heteroskedastic and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55:703-8.

Roy, Anindya, and Wayne A. Fuller. 2001. Estimation for Autoregressive Time Series with a Root Near 1. *Journal of Business & Economic Statistics* 19:482-93.

Shephard, Neil. 1994. Local Scale Models: State Space Alternative to Integrated GARCH Processes. *J. Econometrics* 60:181-202.

Stock, James H., and Mark W. Watson. 2007. *Introduction to Econometrics,* 2nd ed. Pearson/Addison-Wesley.

Walker, John. 2001. Hotbits: Genuine Random Numbers, Generated by Radioactive Decay. www.fourmilab.ch/hotbits. Downloaded June 12, 2001.

White, Halbert. 1980. A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica* 48:817-38.