2. Oil Shocks in the NAFTA Countries

Apostolos Serletis and Christopher Jacyk

1. Introduction

The relationship between oil prices and the economy is important and has been the subject of a large number of studies, exemplified by Hamilton (1983). These studies have, almost without exception, concentrated on the apparently adverse business-cycle effects of oil price shocks. For example, Hamilton (1983) working on pre-1972 data and based on vector autoregression (VAR) analysis, concluded that energy prices are countercyclical and lead the cycle. However, as Mork (1989, p. 74) put it "... his study pertained to a period in which all the large oil price movements were upward, and thus it left unanswered the question of whether the correlation persists in periods of price decline." In fact, as shown by Mork (1989), there is an asymmetry in the responses in that the correlation between oil price decreases and gross national product (GNP) growth is significantly different than the correlation between oil price increases and GNP growth, with the former being perhaps zero.

More recently, Serletis and Kemp (1997) investigate the cyclical behavior of energy prices using the methodology suggested by Kydland and Prescott (1990) and monthly data for crude oil, heating oil, unleaded gasoline, and natural gas for the period that each of these commodities has been traded on organized exchanges. Based on stationary Hodrick-Prescott (HP) cyclical deviations, their results are robust to alternative measures of the cycle and indicate that crude oil and heating oil prices are synchronous and procyclical whereas unleaded gasoline and natural gas prices are lagging procyclically. Moreover, energy prices are positively contemporaneously correlated with consumer prices and their cycles lead the cycle of consumer prices, suggesting a possible role for energy prices in the conduct of monetary policy.

However, as Serletis and Kemp (1997) conclude, "... the apparent phase-shift between energy prices and consumer prices should not be interpreted as supporting an effect from energy prices to consumer prices since using lead-lag relationships to justify causality is tenuous. Clearly, the investigation of the empirical relationship between energy prices and consumer prices, by looking at the performance of energy prices as indicators of inflation, is an area for potentially productive future research. Such an examination could utilize current state-of-the-art econometric methodology such as, for example, integration and cointegration theory as well as error-correction modeling (if applicable), using either the single-equation approach or a multi-equation (VAR) framework."

This is exactly what we do in the present chapter. We establish the empirical relationship between oil prices, consumer prices, and industrial production in the NAFTA countries -- United States, Canada, and Mexico -- using both the single-equation approach as well as the multi-equation (VAR) framework. The single-equation approach (although it can be interpreted as a VAR in which a specific subset of coefficients is restricted to equal zero) has clear implications about the response of consumer prices and industrial production to oil price changes. The multi-equation VAR allows us to investigate the dynamic relationships among the variables and to evaluate the effects of unanticipated oil price shocks by tracing out the implied impulse response functions.

Another objective of the chapter is to test the hypothesis that the volatility of oil prices causes (in the Granger sense) consumer prices and industrial production. In doing so, we use recent advances in the financial econometrics literature and model the changing volatility of oil prices by specifying parametric models for volatility and using these models to extract oil price volatility estimates from the data on oil prices. In particular, we use a generalized autoregressive conditional heteroskedasticity (GARCH) model to capture the conditional variation of oil price changes and to construct an oil price shock variable reflecting both the unanticipated component of oil price changes and the time-varying conditional variance of oil price changes.

The chapter is organized as follows. Section 2 briefly discusses the data. In section 3, we investigate the integration and cointegration properties of the variables (since the single-equation approach that this section takes critically depends on these properties) and present Granger causality test results. Section 4 investigates the robustness of the results and section 5 provides a description of the GARCH model that we use to construct the oil price shock variable and reports results regarding the role of oil price volatility. The last section summarizes and concludes the chapter.

2. Data

We study monthly data on crude oil, consumer prices, and industrial production for each of the United States, Canada, and Mexico. We use the West Texas Intermediate (WTI) spot price for crude oil, obtained from the Federal Reserve Economic Data (FRED) bulletin board of the Federal Reserve Bank of St. Louis. This price is generally viewed as the benchmark price for oil in North America. From 1947 to 1980 the WTI crude oil price was adjusted on a quarterly basis by the Texas Railroad Commission (TRC), between 1973 and 1980 it was based largely on the price set by OPEC, and it has been a market-based price since 1981.

The analysis is performed for three different sample periods for the United States and Canada; one using the full sample, beginning in 1947 for the U.S. and in 1961 for Canada; one using a sub-sample beginning in 1974 and one using a sub-sample beginning in 1981. For Mexico, we only use the full sample beginning in April 1972 and the 1981 sub-sample -- we don't have a 1974 sub-sample for Mexico, since the full sample begins in April 1972. The U.S. data end in May 1997, the Canadian in April 1997, and the Mexican in December 1996.

3. The Single - Equation Approach

Unit Root Tests

Since in the single-equation approach estimation and hypothesis testing critically depend on the univariate time series properties of the series, in what follows we test for unit roots by estimating by ordinary least-squares (OLS) the following augmented Dickey-Fuller (ADF) type regression [see Dickey and Fuller (1981)]

$$\Delta \log y_t = \alpha_0 + \alpha_1 t + \gamma \log y_{t-1} + \sum_{j=1}^{k} b_j \Delta \log y_{t-j} + \varepsilon_t \tag{1}$$

where Δ is the difference operator. The null hypothesis of a single unit root is rejected if γ is negative and significantly different from zero.

The k extra regressors in (1) are added to eliminate possible nuisance parameter dependencies in the limit distributions of the test statistics caused by temporal dependencies in the disturbances. In practice, the appropriate value of k is rarely known. One approach would be to use a model selection procedure based on some information criterion. However, Said and Dickey (1984) showed that the ADF test is valid asymptotically if the order of the autoregression is increased with sample size T at a controlled rate $T^{1/3}$. For the sample used, this translates into an order of 8 for the U.S. and Canadian data and an order of 7 for the Mexican data. It is to be noted that for an order of zero, the ADF reduces to the simple DF test. Also, the distribution of the t-test for γ in equation (1) is not standard; rather it is that given by Fuller (1976).

The results reported in Table 1 indicate that the unit root null hypothesis cannot be rejected at the 5% level. We also tested for a second unit root -- that is, for a unit root in the first (logged) differences of the series. The results (not reported here) indicate that the null hypothesis of a second unit root is rejected. Hence, we conclude that these series are having a stochastic trend. This is consistent with the Nelson and Plosser (1982) argument that most macroeconomic and financial time series have a stochastic trend.

Cointegration Tests

As mentioned in the introduction, causality tests critically depend on the integration and cointegration properties of the data. In particular, if the variables are integrated but not cointegrated, OLS yields misleading results. In fact, Phillips (1987) formally proves that a regression involving integrated variables is spurious in the absence of cointegration. In this case, the only valid relationship that can exist between the variables is in terms of their first differences. If, however, the variables are integrated and cointegrate, then the short-run dynamics can be described by an error correction model, in which the short-term dynamics of the variables in the system are influenced by the deviation from long-run equilibrium.

To present some empirical evidence on this issue, we test for cointegration using Johansen's (1988) maximum likelihood approach. In general, this approach to the estimation of the number of linearly independent cointegrating vectors for a vector autoregressive process, X_t, of order p involves (i) regressing ΔX_t on $\Delta X_{t-1}, ..., \Delta X_{t-p+1}$, (ii) regressing ΔX_{t-p} on the same set of regressors and (iii) performing a canonical correlation analysis on the residuals of these two regressions -- see Johansen (1988) for more details or Serletis (1994) for an application.

Table 2 shows the results for the different samples of the Johansen cointegration tests between the oil price and the CPI and between the oil price and the IPI for the United States (in panel A), Canada (in panel B), and Mexico (in panel C). Clearly, cointegration is found between the CPI and the price of oil in the full and 1974 samples of the U.S. data, in all three samples of the Canadian data, and in the 1981 sample of the Mexican data. Cointegration between the IPI and the price of oil is found only for the 1974 samples of the Canadian and U.S. data.

Granger Causality Tests

Finally, to test for causality, in the sense of Granger (1969), it must be assumed that the relevant information is entirely contained in the present and past values of the variables. A specification that suggests itself is

$$\text{(2)} \qquad \Delta z_t = \alpha_0 + \sum_{j=1}^{r} \alpha_j \Delta z_{t-j} + \sum_{j=1}^{s} \beta_j \Delta x_{t-j} + u_t$$

where Δz_t is the inflation rate or the growth rate of industrial production and Δx_t is the growth rate of the oil price. To test if x_t causes z_t in the Granger (1969) sense, equation (2) is first estimated by OLS and the unrestricted sum of squared residuals SSR_u is obtained. Then, by running another regression equation under the restriction that all β_j's are zero, the restricted sum of squared residuals (SSR_r) is obtained. If u_t is white noise, then the statistic computed as the ratio of $(SSR_r\text{-}SSR_u)/s$ to $SSR_u/(T\text{-}r\text{-}s\text{-}1)$ has an asymptotic F distribution with numerator degrees of freedom s and denominator degrees of freedom (T-r-s-1), where T is the number of observations and 1 is subtracted out to account for the constant term in equation (2).

In the case that z_t and x_t cointegrate we use the following specification, instead of (1),

$$\text{(3)} \qquad \Delta z_t = \alpha_0 + \alpha_z \hat{\varepsilon}_{t-1} + \sum_{j=1}^{r} \alpha_j \Delta z_{t-j} + \sum_{j=1}^{s} \beta_j \Delta x_{t-j} + u_t$$

where $\hat{\varepsilon}_{t-1}$ is the error correction term -- that is the OLS residual in the following regression $z_t = \alpha + \beta x_t + \varepsilon_t$. Hence, equation (3) is a bivariate autoregression in first differences, augmented by the error-correction term $\hat{\varepsilon}_{t-1}$. This error-correction model clearly shows how z_t changes in response to stochastic shocks (represented by $\tilde{u}_t$) and to the previous period's deviation from long-run equilibrium (represented by $g_{t\&1}$). In the context of (3), testing for causality from x_t to z_t involves testing the null that $\acute{a}_z = \hat{a}_j = 0$, for $j = 1, ..., s$.

We present p-values for Granger causality F-tests in Table 3 for each of the three samples and for each of the three countries. We find that the hypothesis that the price of oil does not Granger cause consumer prices is rejected at conventional significance levels for the U.S. and Canada, while the hypothesis that the price of oil does not Granger cause industrial production cannot be rejected, for each of the three samples and for each of the three countries. These results suggest that knowledge of past oil price changes improves the prediction of future consumer price changes beyond predictions that are based on past consumer price changes alone.

4. Robustness

The Granger-causality results just reported evaluate the proposition that both anticipated and unanticipated oil price changes influence the inflation rate and the real output growth rate. In this section, we investigate the robustness of these results. In particular, we use the multi-equation (VAR) framework in which the variables are treated as jointly determined, and investigate the dynamic relationships among oil prices, consumer prices, and industrial production. Although the single-equation approach has clear implications about the response of output and prices to changes in oil prices, under the VAR approach we have to deal with the identification and ordering problems, before the results are evaluated.

We start by using the oil price to identify oil-price shocks and investigate the dynamic effects of such shocks in a simple unrestricted three-variable VAR, consisting of the (logged) oil price (Oil), the consumer price index (CPI), and the industrial production index (IPI), in that order. We fit the VAR to the monthly data, we ignore low frequency variables (such as linear trends), and we set the lag length equal to 13. p-values of Granger-causality F-tests (not reported here) for the CPI equation show that the hypothesis of no causality from oil prices to consumer prices can be rejected (at conventional significance levels) for all three countries and for each sample, except for the 1981 Canadian sub-sample. The hypothesis, however, of no causality from oil prices to industrial production cannot be rejected. These results are consistent with those based on thse single-equation approach and point to the usefulness of oil prices as stand-alone leading indicators of inflation, but not of real output growth.

Solid lines in Figures 1-3 show the impulse response functions (over an expanse of five years) for each of the three countries and for each sample period considered earlier. Dashed lines denote ± standard-deviation bands computed using the Monte Carlo method described in RATS with 500 draws from the posterior distribution of the VAR coefficients and the covariance matrix of the innovations. In general, the qualitative pattern exhibited by consumer prices and industrial production, following a positive innovation in the oil price, is consistent with our priors. There is a strong and statistically significant increase in consumer prices and a decline in industrial production, which is not significantly different from zero.

5. The Role of Oil Price Volatility

In this section we explore the variability of oil prices as an influence on consumer prices and industrial production. In doing so, we use recent advances in the financial econometrics literature. In particular, we use a generalized autoregressive conditional heteroskedasticity (ARCH) model to capture the time-varying conditional variance as a parameter generated from a time-series model of the conditional mean and variance of the oil price change.

Let $\Delta\log(\text{Oil})_t$ be the oil price change with conditional forecast $E[\Delta\log(\text{Oil})_t]$ as in the following equation

$$\Delta\log(\text{Oil})_t = E[\Delta\log(\text{Oil})_t \mid I_{t-1}] + \varepsilon_t$$

where I_{t-1} is the conditioning information set on which forecasts are based and the forecast error ε_t has zero mean and conditional variance

$$E\left(\varepsilon_t^2 \mid I_{t-1}\right) = \sigma_t^2$$

Our objective is to use conditional volatility models to capture the time-dependent heteroskedastic distribution of ε_t. By capturing this feature of the data, we can produce a forecasted variance $\hat{\sigma}_t^2$, along with an oil price growth forecast error $\hat{\varepsilon}_t$, such that the standardized residuals, $\hat{\varepsilon}/\hat{\sigma}_t$, are homoskedastic and independent.

A great many conditional volatility models have been proposed in an effort to capture time-dependent heteroskedastic distributions, the most popular of which are members of Engle's (1982) autoregressive conditional heteroscedastic (ARCH) family. One such model, widely used in the literature, is Bollerslev's (1988) generalized ARCH (1,1), or GARCH (1,1) model

$$\sigma_t^2 = w + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

with $w > 0$, $\alpha \geq 0$, and $\beta \geq 0$. This model allows the conditional variance of ε_t to be an ARMA(1,1) process, and nests as special cases a variety of other models, including the Engle (1982) ARCH ($\beta_1 = 0$) and the Engle and Bollerslev (1986) integrated GARCH ($\alpha_1 + \beta_1 = 1$) models.

In the context of this model, the estimated σ_t^2 is the conditional variance of oil price changes -- that is, the variability of oil price changes expected to prevail (next period) given currently available information. The unexpected component of oil price changes is given by $\hat{\varepsilon}_t = \Delta\log(\text{Oil})_t - E[\Delta\log(\text{Oil})_t \mid I_{t-1}]$, where $E[\Delta\log(\text{Oil})_t \mid I_{t-1}]$ is the conditional expectation of $\Delta\log(\text{Oil})_t$. $\hat{\varepsilon}_t$, however, does not reflect changes in conditional variability over time. A measure of unanticipated volatility that reflects both the unanticipated component of oil price changes and the (time-varying) conditional variance of oil price change forecasts is given by $\hat{\varepsilon}_t / \hat{\sigma}_t$. This variable can be thought of as being a measure of how different a given oil price change is from the historical pattern.

To establish the relationship between anticipated and unanticipated oil price volatility and consumer prices and industrial production, we use the single equation approach, discussed in section III. In this approach, oil price volatility (irrespective of whether it is anticipated or unanticipated) is treated as predetermined and the time series properties of the data are imposed in estimation. Hence, to test if oil price volatility causes consumer prices or industrial production in the Granger sense, we employ the following regression equation

$$\Delta z_t = \alpha + \sum_{j=1}^{r} \alpha_j \Delta z_{t-j} + \sum_{j=1}^{s} \beta_j \left(\frac{\hat{\varepsilon}_{t-j}}{\hat{\sigma}_{t-j}} \right) + \sum_{j=1}^{k} \gamma_j \hat{\sigma}_{t-j}^2 + u_t$$

where $\hat{\sigma}_t^2$ is the GARCH (1,1) anticipated volatility of oil price changes and $\hat{\varepsilon}_t / \hat{\sigma}_t$ the unanticipated volatility of oil price changes. We choose common lags of four and have Δz_t be the inflation rate or the growth rate of industrial production.

The causality results are reported in Table 4. F_1 is the test of the null hypothesis that in a regression of Δz_t on lagged values of itself and anticipated and unanticipated oil volatility, the coefficients of unanticipated oil volatility are zero. F_2 is the test of the null hypothesis that in a regression of Δz_t on lagged values of itself and unanticipated and anticipated oil volatility, the coefficients of anticipated oil volatility are zero. Finally, F_3 is a test of the null hypothesis that in a regression of Δz_t on lagged values of itself and unanticipated and anticipated oil price volatility, the coefficients of both unanticipated and anticipated oil price volatility are zero.

Turning to the results, it is clear that oil price volatility (irrespective of whether it is anticipated or unanticipated) does not Granger cause either industrial production or consumer prices, except for unanticipated oil price volatility that causes U.S. consumer prices with the 1974 and 1981 sub-samples and anticipated oil price volatility that causes Mexican prices with the 1981 sub-sample. Results (not reported here) also indicate that this conclusion does not in general change when the unanticipated oil volatility is separated into positive and negative components (in order to test for asymmetric effects).

6. Conclusion

We have considered the empirical relationship between oil prices, consumer prices, and industrial production in the NAFTA countries. The analysis has been conducted using integration and cointegration tests, the single-equation approach (with the time series properties of the data imposed in estimation and hypothesis testing), as well as the multi-equation VAR framework, which represents one of the most promising tools for evaluating the dynamic effects of unanticipated shocks. There is strong evidence that oil price changes are significant stand-alone leading indicators of inflation, but not of real output growth. This result seems to be robust across the single equation and multi-equation approaches.

We also studied the relationship between consumer prices, industrial production, and the variability of oil price changes, using recent advances in the financial econometrics literature. In particular, we tested whether the GARCH (1,1) anticipated and unanticipated volatility of oil price changes has a systematic effect on consumer prices or industrial production. In general, we rejected this hypothesis of Granger causality from oil price volatility to consumer prices and industrial production although there is some indication that unanticipated volatility causes the CPI in the U.S.

Our results, however, must be interpreted carefully. The use of the WTI oil price as a representative for the oil price in Canada and Mexico was a matter of convenience. The actual oil price in Canada, however, was regulated between 1973 and 1984, and the Mexican oil price is still subject to government regulation. While regulators in both countries used the WTI as one benchmark in determining the regulated price, the actual domestic price was rarely, if ever, accurately reflected by the WTI. As a result, the macroeconomic responses to changes in the WTI price may not accurately reflect the response to changes in the domestic oil price in Canada and Mexico. A further consideration in this analysis is the level of economic development in Mexico. While the United States and Canada are both developed industrialized countries, Mexico is a developing economy and has, over the past three decades experienced, among other things, several currency crises and policy reforms. These factors tend to mask the true macroeconomic effects of oil price shocks for the Mexican economy, especially the effects on the CPI.

Table 1. ADF Unit Root Tests in Loggeg evels

Regression: $\Delta \log z_t = \alpha_0 + \alpha_1 t + \alpha_2 \log z_{t-1} + \sum_{j=1}^{m} \beta_j \Delta \log z_{t-j} + \varepsilon_t$

	No drift, No trend	With drift, No trend	With drift, With trend
		A. United States	
CPI	3.09	-1.25	-2.52
IPI	2.94	0.85	-2.42
Oil	0.88	-1.09	-1.60
		B. Canada	
CPI	2.08	-0.73	-1.35
IPI	2.21	-1.99	-2.89
Oil	0.87	-1.16	-1.19
		C. Mexico	
CPI	0.32	-0.36	-1.98
IPI	2.76	-1.84	-2.52
Oil	3.47	-0.86	-1.69

NOTE: The 95% critical value for the ‘no drift, no trend’ case is -1.94; for the ‘with drift, no trend’ case -2.87, and for the ‘with drift, with trend’ case -3.42. An asterisk indicates significance at the 5% level.

Table 2. Johansen Maximum Likelihood Cointegration Tests

	Full sample		1974 Sub-sample		1981 Sub-sample	
System	T	Trace	T	Trace	T	Trace
			A. United States			
CPI, Oil	600	19.009*	281	32.649*	197	15.069
IPI, Oil	600	8.590	281	17.914*	197	9.093
			B. Canada			
CPI, Oil	431	18.528*	280	38.818*	196	21.043*
IPI, Oil	431	10.232	280	17.114*	196	12.123
			C. Mexico			
CPI, Oil	297	11.724	N/A		192	19.025*
IPI, Oil	297	9.656	N/A		192	7.66

NOTE: All tests use a constant and trend variable. The number of lags is set equal to 4. T = number of observations. The 95% critical value is 15.41. An asterisk indicates significance at the 5% level (rejection of the null of no cointegration).

Table 3. Tail Areas of Tests of Granger causality on the Single - Equation Approach from Oil Prices to Consumer Prices and Industrial Produaction

Dependent Variable	Full Sample	1974 Sub-sample	1981 Sub-sample
		A. United States	
CPI	.008	.000	.000
IPI	.809	.844	.497
		B. Canada	
CPI	.000	.000	.025
IPI	.145	.134	.799
		C. Mexico	
CPI	.072	N/A	.165
IPI	.331		.330

NOTE: Low *p*-values imply strong marginal predictive power.

Table 4. Tail Areas of tests of Granger causality from Anticipated and Unanticipated Oil Price Volatility to Consumer Prices and Industrial Production

$$\Delta z_t = \alpha + \sum_{j=1}^{r} \alpha_j \Delta z_{t-j} + \sum_{j=1}^{s} \beta_j \left(\frac{\hat{g}_{t-j}}{\hat{\sigma}_{t-j}} \right) + \sum_{j=1}^{k} \gamma_j \hat{\sigma}^2_{t-j} + u_t$$

	p-values					
	F_1: $\beta_j = 0$, $j = 1, ..., s$		F_2: $\gamma_j = 0$, $j = 1, ..., k$		F_3: $\beta_j = \gamma_j = 0$, $j = 1, ..., 4$	
Sample	IPI	CPI	IPI	CPI	IPI	CPI
			A. United States			
Full	.272	.655	.108	.530	.114	.552
1974 Sub-sample	.280	.000	.332	.307	.327	.290
1981 Sub-sample	.551	.000	.841	.613	.843	.606
			B. Canada			
Full	.656	.194	.062	.282	.065	.322
1974 Sub-sample	.233	.125	.317	.072	.312	.087
1981 Sub-sample	.556	.365	.315	.129	.311	.128
			C. Mexico			
Full	.197	.374	.818	.421	.825	.414
1981 Sub-sample	.425	.214	.860	.005	.871	.005

NOTE: Low p-values imply strong marginal predictive power.

References

Bollerslev, T. (1988), 'Generalized Autoregressive Conditional Heteroskedasticity', *Journal of Econometrics* **31**, 307-327.

Dickey, David A. and Wayne A. Fuller (1981), 'Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root', *Econometrica* **49**, 151-174.

Engle, R.F. (1982), 'Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation', *Econometrica* **50**, 987-1008.

Engle, R.F. and T. Bollerslev (1986), 'Modeling the Persistence of Conditional Variances', *Econometric Reviews* **5**, 1-50.

Fuller, W.A. (1976), *Introduction to Statistical Time Series*. New York: John Wiley & Sons.

Granger, C.W.J. (1969), 'Investigating Causal Relations by Econometric Models and Cross-Spectral Methods', *Econometrica* **37**, 424-438.

Hamilton, James D. (1983), 'Oil and the Macroeconomy since World War II', *Journal of Political Economy* **91**, 228-248.

Johansen, Soren (1988), 'Statistical Analysis of Cointegration Vectors', *Journal of Economic Dynamics and Control* **12**, 231-254.

Kydland, Finn E. and Edward C. Prescott, Edward C. (Spring 1990), 'Business Cycles: Real Facts and a Monetary Myth', Federal Reserve Bank of Minneapolis *Quarterly Review*, 3-18.

Mork, Knut Anton (1988), 'Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results', *Journal of Political Economy* **97**, 740-744.

Nelson, Charles and Charles Plosser (1982), 'Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications', *Journal of Monetary Economics* **10**, 130-162.

Phillips, Peter (1987) 'Understanding Spurious Regressions in Econometrics', *Journal of Econometrics* **33**, 311-340.

Prescott, E.C. (1986), 'Theory Ahead of Business Cycle Measurement', Federal Reserve Bank of Minneapolis *Quarterly Review* **10**, 9-22.

Said, S. and David Dickey (1984) 'Testing for Unit Roots in Autoregressive-Moving Average Models with Unknown Order', *Biometrica* **71**, 599-607.

Serletis, Apostolos (1994), 'A Cointegration Analysis of Petroleum Futures Prices', *Energy Economics* **16**, 93-97.

Serletis, Apostolos and Todd Kemp (1998), 'The Cyclical Behavior of Monthly NYMEX Energy Prices', *Energy Economics*.