Quantifying the Impact of Oil Prices on Inflation

By Colin Bermingham*

Abstract

The substantial increase in the volatility of oil prices over the past six or seven years has provoked considerable comment within the international media. While recent spikes in oil prices have not had quite the same impact as that experienced in the 1970's, there are still significant implications from a macroeconomic perspective. This is particularly the case in terms of inflation. The re-emergence of the oil price issue necessitates a re-examination of econometric estimates of the influence of oil prices on inflation. We examine this issue in the case of a small open economy — that of Ireland.

* The author is an economist in the Bank's Economic Analysis and Research Department. The views expressed in the paper are the personal responsibility of the author and not necessarily those of the CBFSAI. The author woould like to thank Antonello D'Agostino, Kieran McQuinn, Derry O'Brien, Gerard O'Reilly and participants at an CBFSAI seminar for their comments.

1. Introduction

In recent years, oil prices have become increasingly volatile. The spot price of oil exhibited a general upward trend from the beginning of 2003 until the middle of 2008. However, in recent months, there has been a reversal of this trend on the back of a global economic slowdown. The unpredictable behaviour and wild swings in the price of oil have made inflation forecasting more difficult. Oil is a key determinant of the price of many goods in the consumer basket so any changes in oil prices will impact on inflation directly.

The recent spikes in oil prices had a significant effect on inflation internationally. However, the effects of these oil price shocks have been less damaging to economic growth and less inflationary than those experienced during the oil price crises of the 1970's. This issue has already received some academic interest as it suggests that our estimates of the pass-through from oil prices to inflation need to be re-visited. This paper investigates a related problem, which is whether it is possible to forecast the energy component of inflation using oil prices in the Irish case. The Irish situation provides an interesting case study because it typifies the case of a small open economy. Once the impact of oil prices on the energy component of inflation is known, its effect on overall inflation is also know because the energy component has a fixed weight when calculating overall inflation.

Quantifying the effect of oil prices on the energy component of inflation is open to the criticism that second round effects are ignored but there is limited evidence of second round effects to date. Second round effects refer to the situation where oil price increases lead to general inflation in other sectors of the economy through increased production costs. Some second round effects have been seen in the transport sector, operating in both the upward and downward direction. There has also been changes in some regulated prices. Outside the energy and food sectors, Harmonised Index of Consumer Price (HICP) inflation in other sectors is not significantly different from historical levels, suggesting limited evidence of second round effects.

The focus of this article is on short-term forecasts. This represents a realistic target. Accurate medium and long-term forecasts of energy inflation are very difficult to generate given the aforementioned volatility of oil prices. Using a simple regression based forecasting technique, it is possible to construct forecasts of the energy component that are more accurate than those generated by standard

Figure 1: Oil Prices Denominated in Dollars and in Euro

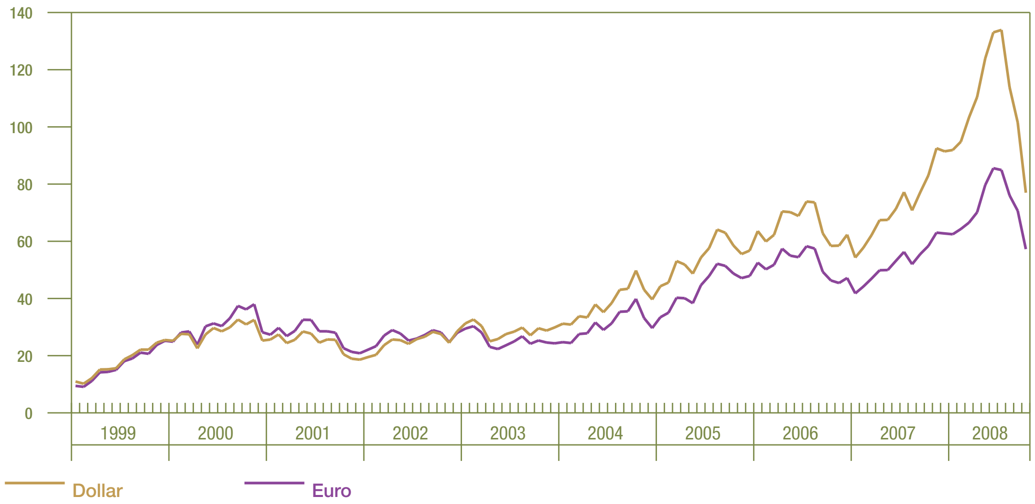

benchmark models. In some empirical applications, sophisticated econometric forecasts have failed to outperform these simple benchmark forecasts. Despite the focus on short-term forecasts, the models can also generate long-term forecasts but these forecasts are very sensitive to the assumptions made about future path of oil prices and exchange rates.[1]

2. Recent Developments

Energy price inflation has rarely been out of the business pages over the last three years. To put recent oil price volatility in context, Figure 1 graphs the price of a barrel of oil, which is priced internationally in US dollars, over the last ten years. The more relevant measure of oil prices in Ireland is the euro price of oil so the graph also includes this series.

It can be seen that the dollar price of oil, denoted in blue, increased more rapidly than the euro price over the last five years as the strength of the euro has insulated those within the euro system from the full impact of the oil price increase. The scale of the recent fall in oil prices is also evident from the graph and the smaller gap between the two series over the last couple of months is indicative of the recent gains of the dollar against the euro.

These oil price swings have had a considerable impact on inflation. Figure 2 graphs the contribution of the energy component to overall inflation in percentage terms over the past ten years. The inflation rate referred to in the graph is the HICP. This measure of inflation is very similar to the Consumer Price Index (CPI) but it is calculated for all European countries to allow for international comparisons. Over this period, year-on-year energy price inflation has generally been significantly higher and more volatile than overall inflation and this is reflected in the large contributions it has made to the inflation rate. There are periods when the contribution of the energy component has exceed 50%, meaning that if overall inflation was 2%, over 1% of this would be driven by the energy component. Over 2005, the influence of energy prices was particularly strong with a contribution of nearly 40% on average over the year. Given the importance of the energy component of inflation, it is critical to have a clear understanding of the impact of oil price changes on energy inflation.

[2] This article provides a non-technical summary of a recent study of the effects of oil prices on inflation but please refer to the technical paper available on the central bank website for the full technical details: Bermingham, C., 2008, "Quantifying the Impact of Oil Prices on Inflation", *CBFSAI Research Technical Paper*, 08/RT/08.

Figure 2: Contribution of Energy Component

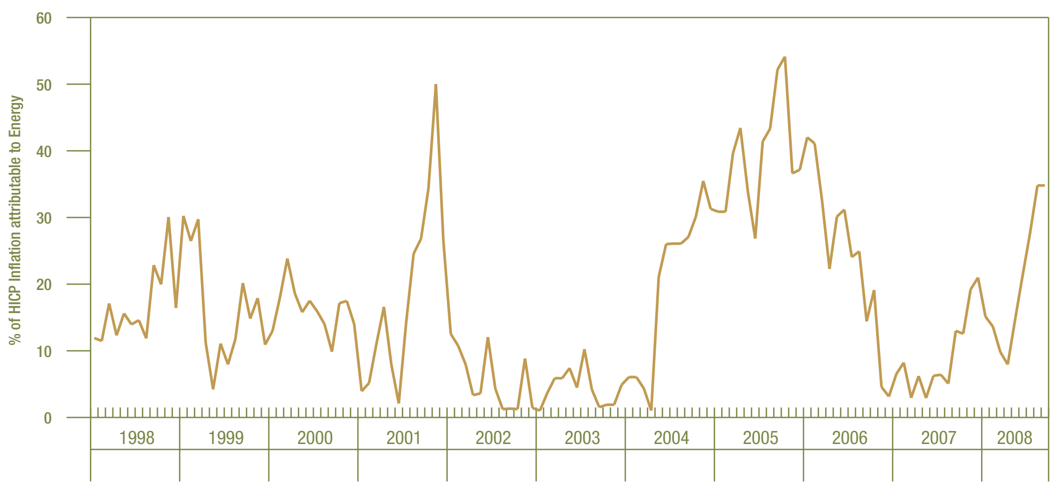

3. The Energy Component of Inflation

The standard approach to estimating the impact of oil prices on inflation is to specify the model in terms of the national consumer price index of the country, which in this case would be either the HICP or the CPI. There is one potential problem with this approach, in that oil prices are not the sole determinant of overall inflation. During the recent increase in oil prices, food prices were also increasing on international markets. Using standard approaches, it is difficult to disentangle which part of an increase in overall inflation is due to higher oil prices and which part is due to higher food prices. To address this problem, we focus only on the energy component of overall inflation in the estimates. Clearly, food price increases are not going to impact on the energy component of inflation so any estimates of the impact of higher oil prices on the energy component are not capturing higher food prices.

The weight of the energy component in the HICP is approximately 8.7% for Ireland. This is slightly higher than its weight of 7.8% in the CPI as the HICP is a smaller basket of goods and services. The composition of the energy components in the HICP and CPI are identical however. Thus, although the focus of this paper is the HICP, the results are equally valid with respect to the CPI energy component. In this paper, the energy component is split into its constituent parts. The first column in Table 1 shows the current weights of the various elements that constitute the energy component. The exact weight of each item and the mix of items in the component changes every five years when the CPI is rebased but changes to the make-up of the energy component over the last fifteen years have been fairly minor.

Unleaded petrol and diesel together account for approximately 48% of the energy component, but with petrol having a weight almost four times that of diesel. Home heating oil is referred to as fuel oil in this paper — this is how it is referred to in the HICP basket by our statistical agency. It accounts for 11% of the index. These three items are all heavily influenced by oil price developments. Other items generally used for home heating such as coal, turf, briquettes and fire-lighters together account for under 10% of the energy component. Electricity is an important item, with a weight of 20%. Of the two gas components, piped gas, at about 10% of the index, is much more important than bottled gas, which only has a weight of 1.5%. However, piped gas and electricity prices are both regulated and change price only occasionally and in discreet jumps. As they do not respond in a predictable way to international energy prices, they were excluded from the current energy component in order to form a new market driven component. This removed approximately 30% of the weight of the official energy component. The second column of Table 1 shows the weights of the newly constructed Non-Administered Energy (NAE) series once the remaining items are rescaled following the removal of the administered price series. This is a market driven energy series that responds more predictably to oil price developments and this series is used as the measure of the energy component in the analysis.

Table 1: Breakdown of the Energy Series

Item	Energy Weights	AE Weights
Petrol Unleaded	38.33%	54.46%
ELECTRICITY	19.89%	n/a
FUEL OIL	11.00%	15.63%
Diesel	9.79%	13.92%
PIPED GAS	9.74%	n/a
COAL	4.42%	6.28%
TURF & BRIQUETTES	4.39%	6.24%
Bottled Gas	1.52%	2.15%
Firelighters	0.63%	0.89%
Motor Oil	0.18%	0.26%
Fire Handy Packs	0.12%	0.17%
Total	100.00%	100.00%

4. Empirical Approach

Two econometric models are used to forecast the energy component in the paper. The first model is called an Autoregressive Distributed Lag Model (ARDL) model. In this type of model, the current inflation rate of the energy component depends on its lags and lags of oil price inflation. This means that current energy price inflation, π_t, is forecast based on previous values of itself and previous values of oil price inflation using this type of equation:

$$\pi_t = \gamma_0 + \sum_{i=1}^{n} \beta_i \pi_{t-i} + \sum_{i=1}^{n} \theta_i \Delta oil_{t-i} + \varepsilon_t$$

The second econometric model, known as a cointegration model, is predicated on the idea that oil prices and energy prices have a stable long-term relationship. Thus, for a given level of oil prices, we would have an expected level for energy prices in the economy. If actual energy prices turn out to be above or below this expected level, we would forecast them to revert back towards this expected level in the future. This approach also takes explicit account of recent developments in energy price inflation and oil price inflation.

There are also two different forecast strategies considered in the paper. The first strategy is to take the two econometric models just described and apply them to the NAE series. However, some of the components of the NAE series also exhibit a strong relationship with oil prices. Given that we know the importance of each item in the NAE series from Table 1, it is possible to forecast the individual items of the NAE series and then combine the forecasts to arrive at a second set of forecasts. In practice, only petrol, diesel and fuel oil are forecast individually but these items account for 85% of the series. It is assumed that the inflation rates of the other items remain unchanged.

5. Forecasts

5.1 Forecast Evaluation

Forecasts are performed on a recursive basis. The first sample period for recursive estimates is December 1996 — March 2005. Models are estimated over this time frame and forecasts for April 2005 — June 2005 are computed. One month is added to the sample and the process is repeated. In this way, 36 sets of forecasts at the three-month horizon were computed for each method. Once the forecasts have been constructed, we need some way to evaluate them. To measure the accuracy of the forecasts, we used a statistic called the Root Mean Squared Error (RMSE). It is calculated as follows

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum e_i^2}$$

where *e* is the forecast error. This equation says that we take the forecast error, square it, find the average and then take the square root. The reason to use a statistic like this is that if we only calculate the average forecast error, positive and negative forecast errors would offset each other so it would not provide a good indication of the accuracy of the forecast. The RMSE must be calculated separately for each forecast horizon (1-month, 2-month and 3-month) and for each econometric approach.

In addition to calculating the RMSE for the econometric models, we also calculate the RMSE for some benchmark models. These benchmarks provide very simple forecasts of the NAE series. The first benchmark forecast is for the inflation rate of the series to remain unchanged. For example, if yearly energy inflation is 4.5% for June 2005, the benchmark forecast for each month for July 2005 — September 2005 is also 4.5%. This will be referred to as the AO benchmark and this type of forecast is often called a naïve forecast. The second forecast is a simple econometric forecast that only uses past values of energy price inflation. This will be referred to as the AR benchmark. Comparing the accuracy of the model forecasts to these benchmarks gives an indication of the worth of the models, as we would not be interested in using models for forecasting purposes unless they prove more accurate than simplistic benchmarks. In some practical applications, these simple benchmarks have proven difficult to outperform.

Table 2 presents the results of the different forecasting methods applied to the individual

Table 2: Direct Monthly Forecasts

	Month 1	Month 2	Month 3
NAE			
AO	4.44	6.52	7.57
AR	4.33	6.29	7.12
ARDL	3.58	5.40	6.14
COINT	3.53	5.32	5.95
PETROL			
AO	5.20	7.75	9.09
AR	4.96	7.22	8.19
ARDL	4.10	6.42	7.16
COINT	4.05	6.25	6.75
Diesel			
AO	3.83	5.77	7.02
AR	3.80	5.75	6.94
ARDL	2.89	4.74	6.34
COINT	2.90	4.76	6.28
Fuel Oil			
AO	7.55	10.13	11.79
AR	7.43	9.76	11.27
ARDL	5.40	7.98	10.73
COINT	5.03	7.69	9.81

components and to the NAE series directly. The numbers in the table are the RMSEs and each section of the table shows the errors for a specific component. In each section, the first two rows show the errors from the two benchmarks — the AO forecast and the AR forecast. The third and fourth rows of each section show the errors for the ARDL forecast and the cointegration model forecast.

The first section presents the results for the NAE series, the market driven energy series constructed in the paper. The ARDL forecasts and the cointegration forecasts are more accurate than both benchmarks as the reported errors are clearly smaller. Although the ARDL and cointegration models have similar forecast power, the cointegration forecasts are marginally more accurate at all forecast horizons. The improvements in forecasts power relative to the AO benchmark using the cointegration approach are 21%, 19% and 21% at the one, two and three-month forecast horizons respectively. This is a meaningful improvement in forecast performance.

A similar picture emerges in the remaining sections of the table, which detail forecast performance for the three main components of the NAE series. For each component, the two econometric forecasts outperform the benchmark forecasts. In addition, the forecasts using the cointegration model are slightly more accurate than those of the ARDL model. Improvements in forecast power relative to the benchmark are greater at the shorter horizons. Using the cointegration approach, the one-month forecast of fuel oil is 33% more accurate than the AO benchmark whereas at the three-month horizon, the greatest improvement in forecast power is for the petrol forecast, which is 26% more accurate than the benchmark. The results indicate that the models are very useful when it comes to predicting the energy series or its components relative to a standard benchmark.

5.2 Forecast Aggregation

Turning to the issue of forecast aggregation, the individual forecasts from the two estimation techniques were combined with naive forecasts for the components that are not modelled explicitly to form a second forecast for the NAE series. Table 3 presents the results of this exercise. The numbers in the table refer to the ratio of the RMSE from the disaggregate approach versus the aggregate approach. A value less than one indicates that the combination of individual forecasts is more accurate than using the same approach to forecast the NAE series directly. The results for

Table 3: Forecast Combination Results

	Month 1	Month 2	Month 3
NAE			
ARDL	0.93	0.98	1.01
COINT	0.92	1.00	0.97

the ARDL approach show that forecast combination can improve forecasts at the one-month horizon by 7% relative to a direct forecast. At months two and three, the forecast errors are broadly similar using indirect versus direct forecasts. For the cointegration model, there is an 8% improvement at the one-month horizon, no difference at two months and a 3% improvement at three months. Overall, the results suggest some role for forecast aggregation, as the results are encouraging at the one-month horizon and, although the gains may be modest at other horizons, it is rarely the case that the combined forecasts are less accurate.

5.3 Refined Prices

In this section, we consider an alternative data source for the input to the forecasts. Up to now, the oil prices used in the models and forecasts were international crude oil prices. In this section, data on the refined price of oil products is used in place of international crude oil. The difference between the two can be thought of as the cost of refining crude oil into a product suitable for retail distribution. Refined prices for gasoline and diesel are available and these represent the prices that refineries charge retailers for gasoline and diesel. Although the price paid at the pump, and in turn reflected in a consumer price index, will also incorporate the profit margin of the retailer and any local taxes, the refined price is closer to the retail price than the price of international crude oil and may help in the construction of more accurate forecasts.

Direct forecasts of the NAE series and its components are constructed in the same manner as before. The refined price of gasoline is used as the input in the NAE and the petrol price forecasts while the refined price of diesel is used to construct the diesel and fuel oil forecasts. The results are presented in Table 4 and are analogous to those in Table 2. The first two rows of each section present the RMSE values for the benchmarks. These are identical to the numbers presented in Table 2 but are replicated here for convenience. As was the case with the forecasts based on crude oil prices, the refined price forecasts are more accurate than the two benchmarks considered at all time horizons.

If we compare the ARDL forecasts using refined prices in Table 4 to those using crude prices from Table 2, there are considerable improvements in forecast power at the one-month horizon for all items. The results are mixed at the two-month horizon. The forecasts for the aggregate NAE series and its diesel component are more accurate but petrol and fuel oil are less accurate. At the three-month horizon, only the forecast for diesel is more accurate. These results would appear to suggest that the benefits to using refined prices are confined to the short forecast horizons. However, if we focus on the aggregate forecast of the NAE series, the improvements in forecast power at months one and two are quite large whereas the forecast at month three is only marginally less accurate, so the evidence in favour of using refined prices in this case is quite compelling.

For the cointegration model forecasts, there are some similarities in the results. At month one, all forecasts are more accurate when using refined prices relative to crude prices. At month two, the forecasts for the NAE series and for fuel oil are more accurate but petrol and diesel are less accurate while all forecasts using refined prices are less accurate for the three-month forecast. In Table 2, we saw that forecasts from the cointegration approach were generally more accurate than those from the ARDL approach. The reverse tends to be true with refined prices. The results using refined prices also suggest that forecast aggregation is unlikely to improve upon the direct forecast, as the component forecasts are less accurate than the direct forecasts in the majority of cases.

Table 4: Direct Monthly Forecasts with Refined Prices

	Month 1	Month 2	Month 3
NAE			
AO	4.44	6.52	7.57
AR	4.33	6.29	7.12
ARDL	2,90	4.71	6.23
COINT	3.19	5.26	6.91
PETROL			
AO	5.20	7.75	9.09
AR	4.96	7.22	8.19
ARDL	3.51	6.46	8.72
COINT	3.51	6.54	8.49
Diesel			
AO	3.83	5.77	7.02
AR	3.80	5.75	6.94
ARDL	2.64	4.51	6.09
COINT	2.85	4.94	6.80
Fuel Oil			
AO	7.55	10.13	11.79
AR	7.43	9.76	11.27
ARDL	5.13	7.44	10.96
COINT	4.53	7.97	10.40

6. Long-Term Forecasts

Despite the focus on short-term forecasts, in this section we generate long-term forecasts over a one-year horizon. Long-term forecasts need to be based on the future path of oil prices and the exchange rate. Forecasting oil prices is now even more difficult than usual given their current volatility. In this paper, the future path of oil is first taken from futures markets and the exchange rate is assumed to remain unchanged from its current level. The actual path of oil prices and exchange rates could turn out to be significantly different and, for this reason, the forecast is subject to a large degree of uncertainty. This is why short-term forecasts were favoured in the paper.

The forecasts are constructed in early January 2009. At this time, the most recent data for the Non-Administered Energy series is available for November 2008. The inflation rate of the series is forecasts for December 2008-December 2009. Exchange rate data and oil price are available for December 2008, the first month in the forecast horizon. For the remainder of the forecast horizon, the exchange rate is assumed to remain unchanged from its monthly average in December 2008 and the assumed path of oil prices is taken from futures data. The futures suggest that the price of oil will rise over the forecast horizon from $46 in January 2009 to $63 by December 2009. The forecast is depicted in Figure 3.

Figure 3: Forecast of Inflation of NAE Series

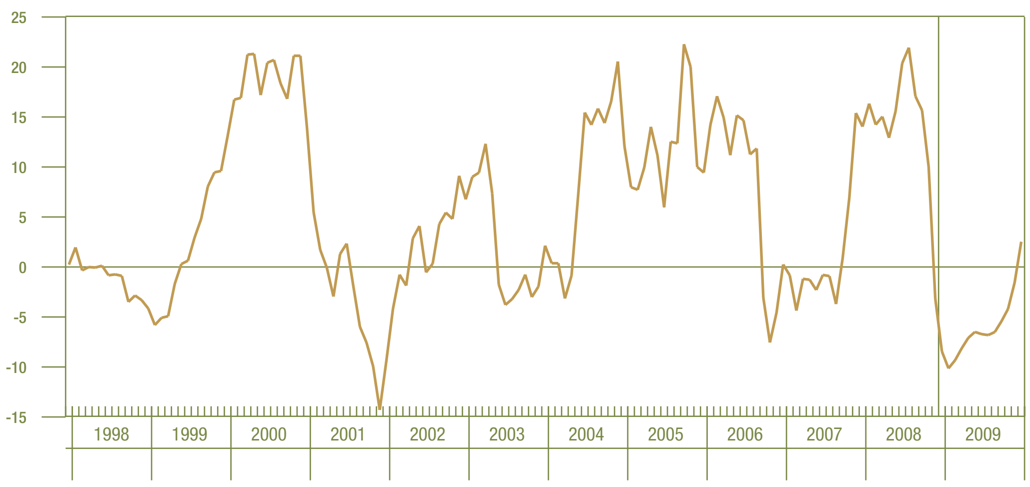

The grid line in the graph signifies the start of the forecast period. Under the assumptions used, it is seen that the inflation rate of the market driven energy component is forecasted to fall quite rapidly in December and January. The inflation rate of the series is expected to rebound from February onwards but it will remain negative for most of 2009. The futures data suggest that NAE energy inflation will not be positive again until December 2009.

7. Summary and Avenues for Future Work

This study provides a means of quantifying the impact of oil price increases on inflation. The exercise is conducted within the context of a small open economy. Overall inflation rates, on an international basis, have been subject to two major influences over the past few years, that of oil prices changes and agricultural commodity price increases. To control for the impacts of agricultural commodity price increases on inflation, the approach adopted here is to focus on energy inflation and, in particular, a measure of energy inflation which is purely market driven. However, once the impact of oil prices on this component is known, it is trivial to calculate the impact on overall inflation.

In focusing on energy inflation and its components, this paper shows that simple econometric techniques significantly outperform standard benchmarks up to three months into the future. The improvements in forecast accuracy range from 20% to 33% at the one-month horizon and from 9% to 26% at the three-month horizon relative to a naïve forecast. By forecasting the constituent parts of the energy series, the issue of forecast aggregation is also considered but any meaningful gains in forecast accuracy are limited to the one-month forecast horizon. Beyond that, it is easier to simply forecast the energy series directly. The paper also investigates whether the use of the price of refined oil products in place of the price of crude oil can improve forecasts. The results indicate that considerable improvements can be made at short time horizons, particularly in the case of the direct ARDL forecast of the energy series. The paper also constructs long-term energy inflation forecasts for the next year although these are quite sensitive to the assumed future path of oil prices and exchange rates.

There are a number of potential avenues for future work. At present, the forecast aggregation procedure only leads to benefits at the one-month horizon. Forecast models could be developed for the items which are currently forecast using naive methods. In addition, local taxes constitute a large percentage of the retail price of petrol and diesel in Ireland. By taking explicit account of this in the model set-up, further improvements in forecast accuracy may be possible. A further consideration is that the retail price of petroleum products may respond asymmetrically to price increases and price decreases. A model that takes allows for this could yield further dividends in terms of forecast performance although a longer time series of data may be necessary to consider this issue.