The European Journal of Finance,
Vol. 14, No. 6, September 2008, 453–468

Routledge
Taylor & Francis Group

Forecasting inter-related energy product prices

M.E. Malliaris[a*] and S.G. Malliaris[b]

[a] *School of Business Administration, Loyola University, Chicago, USA;* [b] *School of Management, Yale University, Connecticut, USA*

Five inter-related energy products are forecasted one month into the future using both linear and nonlinear techniques. Both spot prices and data derived from those prices are used as input data in the models. The models are tested by running data from the following year through them. Results show that, even though all products are highly correlated, the prediction results are asymmetric. In forecasts for crude oil, heating oil, gasoline and natural gas, the nonlinear forecasts were best, while for propane, the linear model gave the lowest error.

Keywords: neural network; regression model; inter-related products; forecasting

1. Introduction

Is it possible to forecast the price movements of related commodities, and if so, how do the symmetries between commodities translate into symmetries in forecast technique? To address these questions, we analyze price movements in five related energy products between December 1997 and November 2002, and find that while the products themselves are strongly related, some prove much easier to forecast than others. Since these products often come from the same raw materials, and since energy consumers can often use technological means to substitute use of one product with another, we would expect strong symmetries in the product forecasts. As we will see, this is not the case.

There is substantial uncertainty about energy prices in the future (Department of Energy 2004). This uncertainty in energy prices commands a great deal of foreign and domestic political attention, and facilitates an active market ranking second only to financial products in amount of trading on futures contracts. As a result, energy commodity price risk has a dominant role in the energy industry (Energy Information Administration 2001, 2002). Herard and Taylor (2003) emphasize that a first step in minimizing earnings volatility in industry is to stabilize margins; a major reason for the margin volatility is the price change in natural gas. Therefore, by accurately forecasting prices into the future, we can create an effective way of managing this risk. This paper looks at a short-term forecast (one month) and compares the ability of two models to handle that forecast. There are about 248 trading days per year, or, around 21 per month. More specifically, we use information from five energy markets to forecast each of those markets 21 trading days from the

*Corresponding author. Email: mmallia@luc.edu

ISSN 1351-847X print/ISSN 1466-4364 online

DOI: 10.1080/13518470701705793
http://www.informaworld.com

day the forecast is made. Data were specifically kept simple as we wanted to focus on the effect of recent price change on pricing one month away.

2. Energy production

In this paper, we study the price movements of crude oil (CL), heating oil (HO), gasoline (HU), natural gas (NG), and propane (PN). Of these five energy products, two – CG and NG – are raw energy products, mined from the ground. CL, the most actively traded commodity in the world, can be refined to give varying amounts of HO, HU, or PN. Natural gas can also be refined into propane, though at lower yields than given by CL. Despite these lower yields, PN is commonly sourced both from CL and NG (Energy Information Administration 2006). A graphical breakdown of these component relationships is shown in Figures 1 and 2 (Energy Information Administration 2005; Roanoke Gas Company 2005).

At present, natural gas accounts for about a quarter of US energy use, but large consumers of energy (factories, power plants) often absorb an up-front cost to implement systems allowing, with minimal change-over cost, the use of either fuel oil or NG (New York Mercantile Exchange 2001) as their main energy source as a hedge against fluctuating market prices and uncertain availability. Because of this substitute relationship, demand for either energy source is more elastic than one would otherwise predict – and, as one would expect, one of the primary determinants of the price of NG is the price of oil (Energy Information Administration 2004).

Table 1 shows the expected high correlations between the above-mentioned energy markets. Academic and industry research and position papers differ as to their explanations for the causes and effects of both absolute energy prices and the relationships therein. Asche, Gjolberg, and

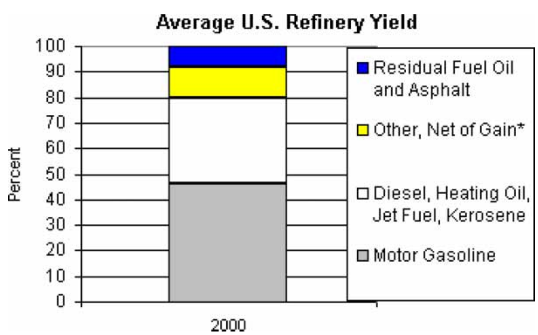

Figure 1. Yield from a barrel of crude oil.

TYPICAL NATURAL GAS COMPONENTS

COMPONENTS	PERCENT BY VOLUME
Methane	89.5 - 92.5
Ethane	5.1 - 2.0
Propane	2.1 - 0.7
Butanes	1.6 - 0.5
Nitrogen	0.6 - 1.9
Carbon Dioxide	1.1 - 2.4
Total	100.0 - 100.0

Figure 2. Components of natural gas.

Table 1. Correlation between energy market prices.

	CL	HO	PN	HU	NG
CL	1	–	–	–	–
HO	0.959721	1	–	–	–
PN	0.842248	0.881154	1	–	–
HU	0.964905	0.926191	0.847288	1	–
NG	0.669869	0.731288	0.677979	0.657551	1

Volker (2003) found a long-run relationship between CL and gas oil prices, but no similar relationship between CL and fuel oil. Brown, Yucel, and Thompson (2003) found that rising oil prices hurt economic activity and preceded nine of the ten recessions following World War II. Current measures indicate that the potential for an energy crisis is historically high (Williams and Alhajji, 2003). OPEC, often a target of criticism about rising prices in energy products, acknowledges a relationship between HU, HO, and CL, but believes the relationship is neither direct nor proportional, and claims that rising prices are due to taxes rather than OPEC policies (Heselbarth 2000).

3. Literature review

There are two strands of relevant literature. The first is the forecasting literature. Since the ability to forecast prices generally depends on the violation of semi-strong form market efficiency in the underlying markets, much of this literature adraws from behavioral finance. Generally, these papers propose a trading strategy based on a behavioral bias or market friction, and confirm the accuracy of their forecast by showing that a trader following the strategy will earn positive abnormal returns. Classic papers in this literature include those of de Bondt and Thaler (1985), who documented long-term (3–5 year) predictable price reversals, and Jegadeesh and Titman (1993), who identified persistent medium-term momentum in stock returns. More recently, Cohen and Frazzini (2006) used such a forecasting strategy to show that investors neglect to take into account public information disclosure by related firms when forming their price forecasts. Specifically, they find that if a firm reports poor performance, one can earn abnormal returns by subsequently shorting the stocks of that firm's *suppliers*: the suppliers' share prices don't respond to the negative event for several weeks.

The second relevant strand of literature deals with the use of neural networks in finance. McNelis (2005) gives examples of the current application of neural networks in bond pricing, production forecasting, inflation and deflation, credit card default, bank failures, and volatility forecasting, among others. Kovalerchuk and Vityaev (2000) present a comprehensive overview of the development and results of major algorithmic models used in finance for prediction, including regression, neural networks, decision trees, ARIMA, and fuzzy logic.

Early attempts to use neural networks to predict stock price movements met with mixed success. White (1988) used a simple neural network to predict IBM stock price movements, and though he was unable to reject the efficient market hypothesis, his simple network was "capable of extremely rich dynamic behavior". Malliaris and Salchenberger (1993, 2002) showed success in forecasting short-term S&P 500 price changes and in estimating option prices. Neural networks have also proved useful in forecasting bankruptcy probabilities. Perez (2006) reviews a series of academic works examining neural networks' success in this type of forecasts. In one example of academic techniques translating into the business world, Burrell and Folarin (1997) describe Chase

Manhattan Bank's decision to use neural networks as part of their credit risk score determination when making corporate lending decisions. Sun et al. (2005) use a radial basis function neural network for forecasting stock prices with good success. Kodogiannis and Lolis (2002) use several variations of neural networks to successfully forecast exchange rates. This case looks at an example with many correlated factors.

Neural networks have been successful in forecasting because of their ability to forecast difficult types of relationships well, given the correct inputs and training set. See, for example, Trippi and Turban (1996) or Smith and Gupta (2002) for collections of papers using neural networks in financial problems. Olson and Mossman (2003) found, in financial models, that they outperform regression models and thus yield greater profitability when trading.

Neural networks are particularly useful in cases where data follow non-linear patterns. Ostermark et al. (2004) and Wu, Lee, and Tan (2006) take advantage of this trait. The former use several non-linear techniques, including neural networks, to model the Finnish Banking and Finance branch index on the Helsinki Stock Exchange. Echoing our own results, they find strong non-linearity in their data. However, where we find that the neural network outperforms linear models in energy markets, Ostermark et al. find that neural networks perform indifferently when forecasting the Finnish branch index. Wu, Lee, and Tan (2006) study delisted firms in the Australian materials industry, and use neural networks to select corporate failure-related features. Gronholdt and Martensen (2005) demonstrate neural networks use in customer satisfaction and loyalty analysis.

4. Data and variables

4.1 *Data*

The data consist of daily spot prices for CL, HO, HU, NG, and PN. All the data are original data, provided by Barchart at www.barchart.com, in absolute values. HO is in cents per gallon, e.g. for 1/3/1994 the price was \$0.453/gallon. PN is in cents per gallon, e.g. for 1/3/1994 the price was \$0.2500/gallon. NG was measured in dollars per MMBtu, e.g. the price for 1/3/1994 is \$2.050/MMBtu. HU was in cents per gallon, e.g. for 1/3/1994 the price is \$0.3970/gallon. Finally, CL is in dollars per barrel, e.g. for 1/3/1994 the price was \$14.54/barrel.

The initial data set contained daily spot prices for the period between January 3, 1994 and December 31, 2002. Inspection of the data over time showed that the relationship between the five markets changed significantly between 1994 and 2002. They became more correlated in movement (see Appendix 1 for graphs illustrating, for example, this relationship between HU and HO for the month of February in 3 years). Anticipating that energy markets will continue to become more rather than less inter-related, and attempting to build a stable model, we used data from December 1997 to November 2002 for this study.

4.2 *Input and output variables*

The input variables for each model consisted of the daily closing spot price, percent change in daily closing spot price from the previous day, standard deviation over the previous 5 trading days, and standard deviation over the previous 21 trading days, for each of the five markets.

For the CL variables, for example, these variables were labeled as CL, CLchg, CLStdDv5, and CLStdDv21. This gave us a total of 20 initial input variables. In addition, the neural network model used a cluster indicator. This is a non-numeric variable indicating a cluster group to which each

Table 2. Correlation of today's spot prices with prices 21 trading days away.

	CL	HO	PN	HU	NG
CLplus21	0.9293	0.9012	0.8029	0.9021	0.6383
HOplus21	0.9242	0.9204	0.8208	0.8898	0.6546
PNplus21	0.8185	0.8329	0.8820	0.8115	0.5684
HUplus21	0.8869	0.8702	0.9265	0.8482	0.6440
NGplus21	0.7160	0.7708	0.7346	0.6810	0.6599

row belonged. The output variable was the daily spot price one month into the future (21 trading days). Again using CL as an example, this variable was labeled CLplus21. The variables used included both raw and derived fields. The raw data help the models fix a current price location, while the derived fields enable the models to discover an underlying relationship that is constant over time. That is, regardless of the current location of the price, the change relationships within and between the commodities could be learned by the models. Models were built for each of the five markets. Table 2 shows the correlation of each energy commodity price today with each price 21 trading days into the future. These values range from a low of 0.6599 to a high of 0.9293.

As the table indicates, these correlations remain high into the future with natural gas showing the lowest overall match. Notice that most of the correlations of one product's prices today with another product's prices 21 days away are slightly lower than those shown between those same products in Table 1, with the exception of NG. For example, the correlation between HO and CL is 0.9597, while the correlation between HO and CLplus21 is 0.9012 and between CL and HOplus21 is 0.9242. Though slightly lower, they are still good enough to encourage us to attempt a model.

4.3 *Training and validation sets*

The data set ran from December 1997 to November 2002, and contained 25 columns. This was split into disjoint training and validation sets.

Twenty of the columns were used for input. Of the remaining five columns, one was selected at a time as the output for each model. The remaining four were unused for that model. For the non-linear model, another input column was generated by the K-Means algorithm, a cluster identifier.

The training set contained 4 years of data, from December 1997 to November 2001, and had 1001 rows. The remaining year of data, December 2001 to November 2002 was used for the validation set. This set contained 247 rows.

The models were judged by their results on validation data that they had not seen during training and which occurred after the end of the training period.

5. Models and methodology

5.1 *Models*

Two model types, one linear, and one non-linear were built for the forecasts. A multiple regression model was built for each of the five energy markets with the spot price 21 trading days into the future as the dependent variable. These models were developed in Microsoft Excel and run with the Excel Data Analysis option. As variables indicated lack of importance to the model, or

Table 3. Variables used in each final regression model.

Variables	CL	HO	HU	NG	PN
CL	x	x	x		x
CLchg	x	x			x
CLStdDv5		x			x
CLStdDv21	x	x	x	x	
HO		x		x	x
HOchg	x				
HOStdDv5	x		x	x	x
HOStdDv21		x	x		x
HU		x	x		
HUchg	x	x			x
HUStdDv5	x	x		x	x
HUStdDv21				x	
NG	x	x	x	x	x
NGchg			x		
NGStdDv5					x
NGStdDv21	x	x	x		x
PN	x	x		x	
PNchg			x		x
PNStdDv5				x	x
PNStdDv21		x	x	x	x

multicollinearity, they were deleted and the model was re-run. A final model for each market was then used for the validation set forecasts. The variables used in each of the multiple regression models are shown in Table 3.

The number of variables used in each of the final regression models ranged from nine to fourteen. Only NG appeared in every model. The variables appearing in at least four of the five models were CL, CLStdDv21, HOStdDv5, HUStdDv5, NGStdDv21, and PNStdDv21. Five out of these six variables appearing in four models are based on movement within the market the past week or month.

Following the development of each of the regression models, the SPSS data mining package Clementine was used to develop a second set of models. These models were non-linear and used two processes in sequence. The data was first run through a K-Means clustering algorithm. In a clustering algorithm, the numerical records are grouped by similarity and a center for the group is calculated. This center is the arithmetic mean of the records in the cluster. The cluster can be described not only by its mean, but also by its distance from other clusters.

The K-Means algorithm in Clementine generated five distinct clusters. Three of these clusters were fairly small, and the remaining two clusters were large. This indicates two common, but distinct, types of patterns of behavior, and three much less likely, but also distinct patterns. The number of rows of the training set per cluster and the center of each cluster (that is, the mean of the records in that cluster for each of the 20 input variables in the training set) are shown in Table 4.

The distance between cluster means, also called cluster proximity, is shown in Table 5. The greater the distance, the more distinct the clusters are. In addition to generating centers for each variable per cluster, the K-Means algorithm identifies the cluster to which each row belongs and generates a node to inspect data and assign it to a cluster. These cluster assignments are

Table 4. Cluster centers and population.

Variable	Cluster-1	Cluster-2	Cluster-3	Cluster-4	Cluster-5
	482 Records	17 Records	384 Records	58 Records	60 Records
CL	15.992	28.231	27.873	28.213	31.167
CLchg	0	−0.01	0.001	0	0.002
CLStdDv5	0.342	0.82	0.59	0.6	0.827
CLStdDv21	0.676	2.382	1.196	0.863	1.304
HO	0.424	0.929	0.765	0.774	0.91
HOchg	0	−0.005	0.001	0	−0.002
HOStdDv5	0.009	0.032	0.02	0.017	0.024
HOStdDv21	0.017	0.06	0.041	0.027	0.041
HU	0.461	0.727	0.794	0.879	0.869
Huchg	0.001	0.001	0.001	−0.003	−0.002
HUStdDv5	0.011	0.021	0.023	0.027	0.026
HUStdDv21	0.023	0.085	0.044	0.057	0.038
NG	2.166	9.236	3.375	4.408	6.686
NGchg	0	0.02	0.002	−0.005	−0.001
NGStdDv5	0.072	0.764	0.102	0.127	0.391
NGStdDv21	0.144	1.313	0.213	0.246	0.821
PN	0.286	0.776	0.503	0.594	0.547
PNchg	0.001	0.016	0.001	0.013	−0.008
PNStdDv5	0.006	0.031	0.013	0.042	0.022
PNStdDv21	0.013	0.059	0.03	0.13	0.062

Table 5. Distance between clusters.

	Cluster-1	Cluster-2	Cluster-3	Cluster-4	Cluster-5
Cluster1	–				
Cluster2	1.996	–			
Cluster3	0.916	1.421	–		
Cluster4	1.253	1.402	0.638	–	
Cluster5	1.376	0.961	0.691	0.744	–

non-numeric. Other data sets can be fed through a trained K-Means node to generate a cluster assignment for each row in them.

The resulting cluster assignment, Cluster1–Cluster5, became an additional input into the neural network model. This input was symbolic data. Neural networks can accept both numeric and non-numeric data. A neural network model was developed for each of the five markets.

Each neural network model used 21 inputs (the 20 original fields, plus the non-numeric cluster identifier), one hidden layer with 20 nodes, and one output node. The neural network model generates a sensitivity analysis that lists the model variables in order of their effect on the model. Table 6 contains this output for each of the five models. Thus, the variables used in the neural network models are listed in their order of importance to the model in this table. Note that in four of the models, HO price was a prominent variable.

If we look at the four most important variables in each of the models, we find little agreement among them. In particular, the CL model has no variable in the top four in common with any other model. HO is one of the top two variables of each of the other models. The NG model has

Table 6. Order of variable significance in neural network models.

CL	HO	HU	NG	PN
Cluster	HO	CL	HO	HO
PNStdDv5	NGStdDv5	HO	PN	CL
Huchg	NGStdDv21	NGStdDv21	CL	NGStdDv5
NGchg	PN	CLStdDv5	NG	PN
HU	PNStdDv5	Cluster	HUStdDv5	NGStdDv21
HUStdDv5	CL	NG	PNStdDv5	PNStdDv5
NGStdDv21	Cluster	PNStdDv5	PNStdDv21	HUStdDv21
CL	PNStdDv21	NGStdDv5	NGStdDv21	CLStdDv5
NG	HU	PNStdDv21	Cluster	HU
HOStdDv21	CLStdDv5	HUStdDv21	HUStdDv21	Cluster
PNStdDv21	Huchg	HU	CLStdDv21	Huchg
CLStdDv21	HOStdDv21	HOStdDv21	CLStdDv5	HUStdDv5
CLchg	HUStdDv5	PNchg	NGStdDv5	NGchg
HUStdDv21	CLStdDv21	NGchg	Huchg	Hochg
Hochg	HUStdDv21	PN	NGchg	PNchg
HOStdDv5	Hochg	Huchg	HOStdDv5	PNStdDv21
NGStdDv5	HOStdDv5	CLchg	HU	HOStdDv21
HO	NG	CLStdDv21	Hochg	CLStdDv21
CLStdDv5	CLchg	Hochg	PNchg	HOStdDv5
PN	PNchg	HUStdDv5	CLchg	CLchg
PNchg	NGchg	HOStdDv5	HOStdDv21	NG

no StDev variables in the top four while all the other models do. Thus, even though the products are inter-related, the price forecasts emphasize different variables.

5.2 *Methodology*

A regression model and neural network model was developed for each of the five energy commodities to forecast the price one month into the future. Data was used, for the training models, from December 1997 to November 2001, for a total of 1001 row of input data.

The regression models were refined until a model using only variables showing no multicolinearity were used. The non-linear approach used first a K-Means clustering algorithm, then a neural network to develop a forecasting model. The networks were trained using a random 50% of the data, with the other 50% used as a testing set to prevent the network from overtraining. When a network is allowed to develop with all of the training data, it frequently does very well on the training data, but does less well on data it has never seen (in our case, we will use the validation set to finally judge this). To prevent this, data is randomly selected, by the algorithm, from the training set, held out from the training process, and used to periodically compare the performance of the network on each piece of the training set. When the network begins to do less well on the testing set than the training set, Clementine stops training the network. This results in a network that is able to generalize better when predicting new data.

Next, a validation set, comprised of a year's worth of data that occurred after the training data, was run through each of the models. This validation set data, for the non-linear model, ran first through the clustering algorithm and each row was assigned a value indicating the cluster to which it belongs. The data then flowed into the trained neural network and a forecast was generated.

Models were compared, for each of the five energy products, by looking at the validation set mean absolute error and the mean-squared error of each regression and neural network model. These statistics from the validation sets were used as the measures of ability of the linear and non-linear models to forecast the five energy markets.

In addition, as a null hypothesis, a simple forecast was generated by assuming that, on average, prices do not change between today and 21 days from now. This null is consistent with the hypothesis that prices follow a random walk with no drift. The same comparison statistics were calculated for this simplistic view of prediction.

6. Models results on training sets

The regression statistics generated by Excel on the training data for each of the models are shown in Table 7. As can be seen in the table, all models did well on the training set, with the lowest R^2 going to PN.

Table 7. Regression statistics.

	CL	HO	HU	NG	PN
Multiple R	0.948154	0.949655	0.919218	0.913553	0.884675
R Square	0.898997	0.901845	0.844962	0.83458	0.78265
Adjusted R Square	0.897977	0.900552	0.843396	0.833077	0.779564
Standard Error	2.199953	0.066021	0.079985	0.6553	0.070904

Table 8. Regression coefficients.

Variable	CL	HO	HU	NG	PN
Intercept	1.2802	−0.0265	0.0574	−0.4755	0.0049
CL	0.7676	0.0217	0.0120		0.0076
CLChg	−6.2214	−0.2938			−0.1580
CLStdDv5		0.0253			−0.0146
CLStdDv21	1.1483	0.0192	0.0216	−0.3455	
HO		0.3671		3.8124	0.4179
HOChg	5.2500				
HOStdDv5	29.2962		0.7437	−4.0060	0.9232
HOStdDv21		−0.5076	0.9226		−0.5630
HU		−0.2205	0.2709		
HUChg	7.8516	0.2477			0.1479
HUStdDv5	−32.4310	−0.7473		5.3742	−1.0396
HUStdDv21				9.9400	
NG	0.6215	0.0185	0.0460	0.7399	0.0191
NGChg			−0.0773		
NGStdDv5					0.0946
NGStdDv21	−4.3548	−0.1398	−0.2250		−0.2425
PN	5.0062	0.1433		−1.6273	
PNChg			0.1381		0.0881
PNStdDv5				−2.8196	0.2912
PNStdDv21		−0.4951	−0.2878	−13.0017	−0.2469

Table 9. Relative importance of inputs in neural network models.

Input	CL	HO	HU	NG	PN
CL	0.0329	0.1606	0.4019	0.1462	0.2507
CLchg	0.0218	0.0172	0.0338	0.0034	0.0131
CLStdDv21	0.0245	0.0487	0.0321	0.0641	0.0309
CLStdDv5	0.0087	0.0890	0.1709	0.0572	0.1193
HO	0.0131	0.3612	0.3828	0.2424	0.3594
Hochg	0.0183	0.0246	0.0217	0.0156	0.0577
HOStdDv21	0.0265	0.0717	0.0775	0.0016	0.0477
HOStdDv5	0.0162	0.0241	0.0123	0.0435	0.0200
HU	0.0386	0.1369	0.0822	0.0170	0.1158
Huchg	0.0568	0.0820	0.0421	0.0512	0.0986
HUStdDv21	0.0213	0.0433	0.0884	0.0826	0.1374
HUStdDv5	0.0383	0.0686	0.0211	0.1249	0.0925
NG	0.0289	0.0228	0.1445	0.1365	0.0032
NGchg	0.0406	0.0057	0.0675	0.0436	0.0690
NGStdDv21	0.0374	0.2743	0.3026	0.0904	0.1712
NGStdDv5	0.0145	0.2779	0.1159	0.0562	0.2345
PN	0.0028	0.2301	0.0570	0.1940	0.2315
PNchg	0.0023	0.0101	0.0739	0.0081	0.0517
PNStdDv21	0.0247	0.1405	0.0989	0.1097	0.0485
PNStdDv5	0.0593	0.1843	0.1392	0.1098	0.1406
Cluster	0.1992	0.1446	0.1546	0.0830	0.1130

Table 8 gives the coefficients of each of the regression equations. Blank cells indicate unused variables for that model. The coefficients from model to model emphasize the difference placed on each variable between models.

Table 9 shows the relative importance of each of the variables in the non-linear neural network model. With a neural network, the sensitivity analysis values vary from 0 to 1. The higher the number attached to an input variable, the greater the importance of that variable in determining the value of the output variable.

Variables whose importance are greatest in the neural network forecast have values over 0.1 in the table. Notice that NG, while important in every regression model, is more important in only two of the neural network models. Also, in the CL model, the cluster affinity variable is much greater in significance than any other variable in the model. Comparison of these two tables indicates a very different approach taken by the two model types to forecasting each product.

Table 10. Training data results with neural network.

	CL	HO	HU	NG	PN
Minimum error	−11.524	−0.199	−0.317	−2.461	−0.215
Maximum error	10.403	0.399	0.480	4.746	0.389
Mean error	−2.152	−0.036	−0.024	−0.157	0.01
Mean absolute error	4.732	0.099	0.096	0.641	0.061
Standard deviation	5.183	0.111	0.124	0.98	0.087
Linear correlation	0.774	0.859	0.789	0.792	0.817
Estimated accuracy	93.761	95.436	96.080	97.064	95.675

Table 11. Comparison of simple model, regression and neural network errors.

	Average absolute error			Mean squared error		
	Simple	Regression	Neural net	Simple	Regression	Neural net
CL	1.973	2.126	1.120	6.013	6.653	2.269
HO	0.051	0.055	0.035	0.004	0.005	0.002
HU	0.057	0.053	0.029	0.006	0.004	0.001
NG	0.388	0.414	0.218	0.240	0.242	0.075
PN	0.041	0.061	0.080	0.003	0.006	0.009

Table 12. Percent of correct direction of forecasts.

	Regression (%)	Neural net (%)
CL	40	79
HO	63	72
HU	65	83
NG	68	81
PN	68	69

The neural network results on the training data are shown in Table 10. This table gives the estimated accuracy of the model along with typical statistics generated by the network. The estimated accuracy for each model compares well with the regression results. On the training data results, shown in Tables 7 and 10, both models appear to do rather well.

7. Model results on validation data

After the development of the models on the training set data, the next step was to run the validation set data through each of the models. The validation set data occurred in time after the end of the training data, so it is a good measure of the ability of each of the models to forecast into the future. As a measure for each regression and neural network model, we used the mean absolute error and the mean squared error. For comparison purposes, a "Simple Model Forecast" was made by assuming that the market in 21 days would be the same as it was today. The same error quantities were calculated for this model.

Table 11 shows the mean absolute error and mean-squared error for each of the models developed. The best model, that is, with the lowest error, for each product is shown in bold type. With the exception of the models for PN, the neural network models gave a forecast with smaller error over the year. For forecasting PN, the best results came from the Simple Model. So, neither the regression nor neural network approaches were able to do well in forecasting PN with the given variables.

Table 12 gives a different way of looking at the overall results. If we ignore the actual amount of the forecasts and look only at the direction predicted (that is, up or down), then the neural networks significantly outperformed the regression models in all but the PN case. With PN, the neural network just slightly edged out the regression model.

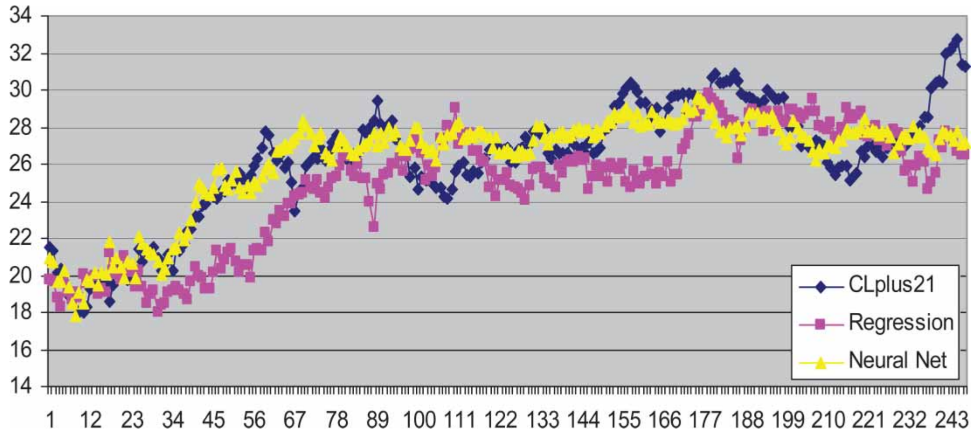

Figure 3. Crude oil: Regression and neural network forecasts compared with actual prices for the validation set time period.

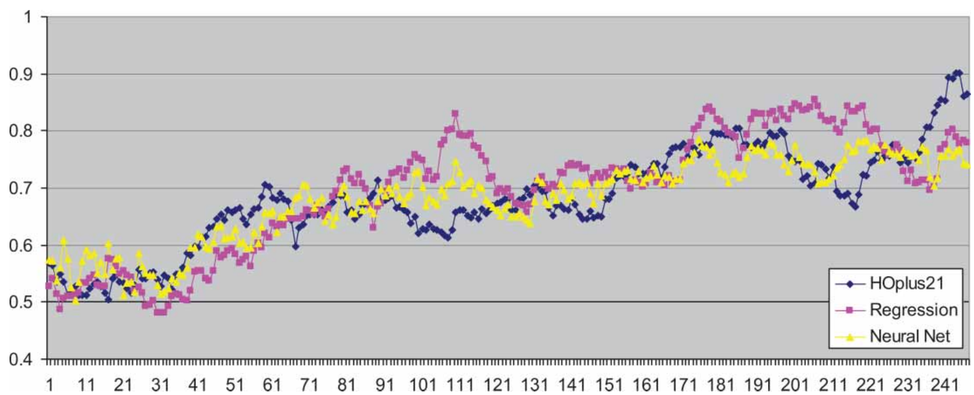

Figure 4. Heating oil: Regression and neural network forecasts compared with actual prices for the validation set time period.

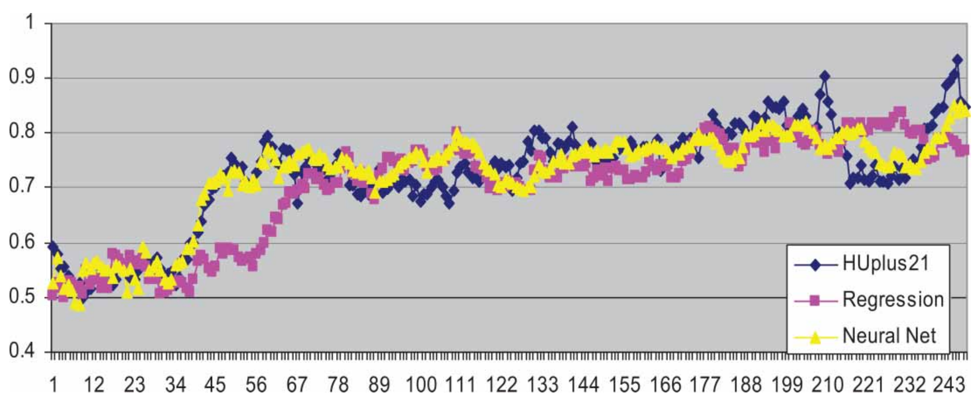

Figure 5. Gasoline: Regression and neural network forecasts compared with actual prices for the validation set time period.

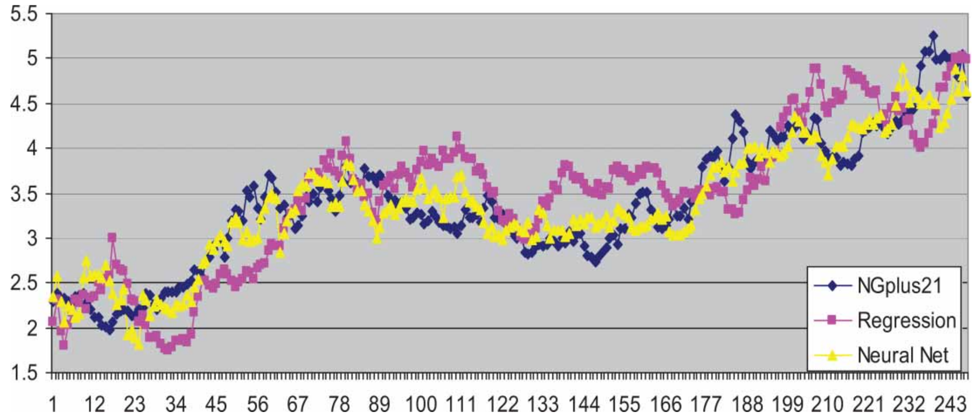

Figure 6. Natural gas: Regression and neural network forecasts compared with actual prices for the validation set time period.

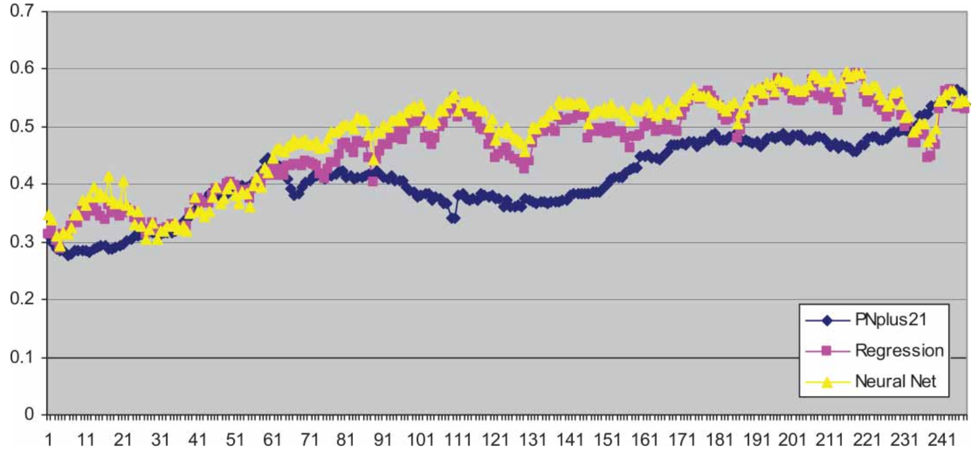

Figure 7. Propane: Regression and neural network forecasts compared with actual prices for the validation set time period.

Figures 3–7 show the relationship of the actual prices, forecasted regression prices and forecasted network prices. As can be seen, other than for PN, the neural network forecast is closer to the actual price than that of the regression model.

8. Conclusion and suggestions for future research

While the non-linear models derived by the neural network provided superior forecasting in a majority of cases, there remain undeniable asymmetries in the predictive abilities of the three models examined. For CL, HO, HU, and NG, the neural network gave the best results, consistently boasting a mean-squared error less than half that of the regression or simple predictions. However, with PN, the neural network gave the least accurate prediction. Also surprising were the results of the simple model, which predicted that the commodity price 21 trading days into the future would hold unchanged from the present day's price. Except in the case of HU, the mean-squared error of the simple prediction was lower than that of the regression model.

One can draw a number of conclusions from these results. First, it is clear that in a number of cases, there is enough information contained in a simple set of price data to allow effective forecasting. That is to say, while the neural network had no extraneous knowledge (say, of news items regarding energy consumption, or foreign wars, or what season it might be), it was nonetheless able to make a reasonably accurate prediction for four out of five energy products.

Second, an ability to predict the price of a given source good does not necessarily imply an ability to predict the price of such a good's byproducts. This asymmetry is exhibited by the neural network model's inability to effectively price PN futures. PN, as we recall from the introduction, is produced by processing or refining NG or CL; the model is generally capable of predicting prices for both these raw materials. But when we try to turn the same model to predicting the price of PN itself, not only does the model produce less-than-stellar results, but it would be bested both by a standard regression and by a trader following the simplistic idea that prices would remain constant.

Lastly, we observe that traditional statistical techniques for understanding and extracting information about trends are often less than ideal in market situations, despite their otherwise broad acceptance by the profession. Regression analysis might be a common tool in analysis of market movements, but the results show that in energy markets, regressions are not appreciably better at predicting future movements than models making the blanket assumption of a simple market.

This paper has concerned itself with an examination of the performance of models simulating the actions of a trader concerned with only the most basic of information: present and historical commodity prices. As we have seen, while such simple information can often provide comparatively excellent predictions, some surprising asymmetries remain. Future work along these same lines might begin to explain the root causes of these asymmetries by examining more closely the relationships between these various energy products, considering specifics such as time-to-market, refining and processing costs, change-over costs for energy users switching from one fuel to another, and volume of usage and cyclical or seasonal factors.

References

Asche, F., O. Gjolberg, and R. Volker. 2003. Price relationships in the petroleum market: An analysis of crude oil and refined product prices. *Energy Economics* 25, no. 3: 289–301.

Brown, S., M. Yucel, and J. Thompson. 2003. Business cycles: The role of energy prices. Working Papers 03–04, Federal Reserve Bank of Dallas.

Burrell, P., and B. Folarin. 1997. The Impact of neural networks in finance. *Neural Computing & Applications* 6: 193–200.

Cohen, L., and A. Frazzini. 2006. Economic links and predictable returns. Yale Working Paper.

De Bondt, W., and R. Thaler. 1985. Does the stock market overreact. *Journal of Finance* 40: 793–805.

Department of Energy. 2004. Annual Energy Outlook 2004 with Projections to 2025. http://www.eia.doe.gov/oiaf/archive/aeo04/index.html

Energy Information Administration. 2001. Energy price impacts on the U.S. economy. www.eia.doe.gov/oiaf/economy/energy_price.html

Energy Information Administration. 2002. Derivatives and risk management in the petroleum, natural gas, and electricity industries, U.S. Department of Energy, October. http://www.eia.doe.gov/oiaf/servicerpt/derivative/index.html.

Energy Information Administration. 2004. Residential natural gas prices: Information for consumers, October, 2004. http://www.eia.doe.gov/neic/brochure/oil_gas/natgas04/Chapter1.htm.

Energy Information Administration. 2005. Average U.S. refinery yield. http://www.eia.doe.gov/pub/oil_gas/petroleum/analysis_publications/oil_market_basics/ref_image_simple.htm

Energy Information Administration. 2006. Propane prices, what consumers should know. http://www.eia.doe.gov/bookshelf/brochures/propane06/propane.html.

Gronholdt, L., and A. Martensen. 2005. Analysing customer satisfaction data: A comparison of regression and artificial neural networks. *International Journal of Market Research* 47: 121–30.

Herard, B., and D. Taylor. 2003. Reducing earning volatility: A short course in energy cost risk management. *Chemical Market Reporter*, 8 September.

Heselbarth, R. 2000. Oil prices hitting homes. *Contractor*, July.

Jegadeesh, N. and S. Titman. 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48: 65–91.

Kodogiannis, V., and A. Lolis. 2002. Forecasting financial time series using neural network and fuzzy system-based techniques. *Neural computing and applications* 11: 90–102.

Kovalerchuk, B., and E. Vityaev. 2000. *Data Mining in Finance*. Hingham, Massachusetts: Kluwer Academic Publishers.

Malliaris, M., and L. Salchenberger. 1993. A neural network model for estimating option prices. *Applied Intelligence* 3: 193–206.

———. 2002. Using neural networks to discover patterns in international equity markets. In *Neural networks in business: Techniques and applications*, Eds. K. Smith and J. Gupta. Hershey, Pennsylvania: Idea Group Publishing.

McNelis, P. 2005. *Neural networks in finance.* Amsterdam: Elsevier Academic Press.

New York Mercantile Exchange. 2001. Risk management with natural gas futures and options. http://www.nymex.com/broch_main.aspx.

Olson, D., and C. Mossman. 2003. Neural network forecasts of Canadian stock returns using accounting ratios. *International Journal of Forecasting* 19: (July–September): 453–65.

Ostermark, R., J. Aaltonen, H. Saxen, and K. Soderlund. 2004. Nonlinear modelling of the Finnish Banking and Finance branch index. *The European Journal of Finance* 10: 277–89.

Perez, M. 2006. Artificial neural networks and bankruptcy forecasting: A state of the art. *Neural Computing and Applications* 15: 154–63.

Roanoke Gas Company. 2005. The origin of natural gas. http://www.roanokegas.com/aboutus/origin.html

Smith, K., and J. Gupta. 2002. *Neural networks in business: Techniques and applications*. Hershey, Pennsylvania: Idea Group Publishing.

Sun, Y., Y. Lians, W. Zhang, H. Lee, W. Lin, and L. Cau. 2005. Optimal partition algorithm of the RBF neural network and its application to financial time series forecasting. *Neural Computins and Applications* 4: 36–44.

Trippi, R., and E. Turban. 1996. *Neural networks in finance and investing*. rev. ed. Chicago, Illinois: Irwin.

White, H. 1988. Economic prediction using neural networks: The case of IBM daily stock returns. *IEEE International Conference on Neural Networks*: 451–58.

Williams, J., and A. Alhajji. 2003. The coming energy crisis? *Oil and Gas Journal* 101 (February).

Wu, W., V. Lee, and T. Tan. 2006. Data preprocessing and data parsimony in corporate failure forecast models: Evidence from Australian materials industry. *Accounting and Finance* 46: 327–45.

Appendix 1. Prices of gasoline and heating oil, February 1994, 1998, and 2002