MULTIPLE HORIZONS AND INFORMAITON IN THE DOE'S ENERGY SUPPLY FORECASTS: A DIRECT TEST FOR INFORMATION CONTENT

Mark R. Manfredo, Arizona State University (Polytechnic Campus)
Dwight R. Sanders, Southern Illinois University

ABSTRACT

The Department of Energy's (DOE) Energy Information Administration (EIA) provides energy supply forecasts over a number of forecast horizons. This research examines the informational content of the DOE's quarterly supply forecasts for crude oil, natural gas, coal, and electricity made for one-, two-, three-, and four-quarters ahead. We specifically consider the information provided in longer-term forecasts relative to the information provided by short-term forecasts. Results suggest there is valuable information provided in the EIA's forecasts from one- to four-quarters ahead. However, forecasts are often not rational, suggesting the use of a composite forecast incorporating the information from both nearby and more distant horizon forecasts.

INTRODUCTION

Private industry participants—such as producers, utilities, and traders—as well as governments and political groups use energy production forecasts to manage expectations concerning prices and levels of business activity. The Department of Energy's (DOE) Energy Information Administration (EIA) accommodates the industry's need for energy outlook information by "providing high-quality, policy-independent energy information to meet the requirements of Government, industry, and the public in a manner that promotes sound policymaking, efficient markets, and public understanding" (Caruso, 2005). Indeed, the traditional objective of any public market outlook effort is to increase profits, utility, or social welfare through more effective economic decisions (Freebairn, 1978). Accurate public information can result in improved decision-making by private forecasters, while also reducing market price variation (Smyth, 1973). Conversely, systematic errors in forecasts could lead to a misallocation of scarce resources (Stein, 1981). Since the EIA's energy-related forecasts are important to a number of stakeholders, it is important that their performance be well understood (Aaron, 2000). At the most rudimentary level, forecasts need to provide information to decision makers. Although other academic notions of forecast performance—namely bias, rationality, and efficiency—are important (see Pons, 2000), forecasts must contain information to accomplish the stated goals of the public forecasting agency and to further the interests of their public clientele.

While there is a long and rich literature of public forecast evaluation in agricultural commodity markets (e.g., Pettee 2000; Freebairn, 1978; Sumner and

Mueller, 1989; Carter and Galopin, 1993; Baur and Orazem, 1994; Garcia; et al., 1997) and for macroeconomic variables (see Granger, 1996; Mills and Pepper, 1999; Schuh, 2001), there is generally a paucity of research examining the performance of public forecasting efforts in the energy sector, in particular with regards to their information content. For instance, while Lynch (2002) provides an excellent overview of long-term crude oil supply forecasting and notes that "oil supply forecasting has been done very badly in the past" with a tendency for forecasters to repeat methodological errors (p. 387), little if anything is said about their information content. Similarly, Floris, et al. (20010 note that production forecasts can vary widely depending on the forecaster's initial model and parameter choices. In a review of their own annual forecasts, the EIA provides insight into the performance of their forecasts through explanatory analysis and examination of forecast accuracy measures (Sanchez, 2006). While this information is helpful in understanding the EIA's forecasting procedure and errors, little insight is provided regarding information content, in particular at longer forecast horizons. Bentzen and Linderoth (2005), however, do shed some light into the performance of longer term energy consumption forecasts documenting that energy consumption projections have only "weakly" improved since the 1970s and that the forecasts tend to degrade quickly at longer horizons. Shlyakhter, et al. (1994) also provide a method for calculating confidence areas around the EIA's annual forecasts. Indeed, each of these studies provides valuable insight into the overall performance of publicly available energy forecasts. However, each of these studies fails to examine the unique information content of EIA forecasts at multiple forecast horizons.

Therefore, in this research we examine the information content of the EIA's supply forecasts. The supply forecasts examined include crude oil, natural gas, coal, and electricity. In particular, we consider the information content in these supply forecasts across multiple forecast horizons (one- to four-quarters) incorporating the methodology proposed by Vuchelen and Gutierrez (2005). The method of Vuchelen and Gutierrez is used to directly test if multiple horizon forecasts provide any information beyond the one-period ahead forecast. For instance, if the EIA is forecasting a 3.0% increase in natural gas production one-quarter ahead, and a 3.5% increase two-quarters ahead, then does the forecasted 3.5% increase two quarters hence really provide any useful information beyond the 3.0% forecasted for one-quarter ahead? It may indeed be the case that the longer horizon forecast provides unique information relative to the shorter horizon forecast. However, it may also be the case that the two-quarter ahead forecast is just a random adjustment to the one-quarter ahead forecast.

This research makes contributions on several fronts. First, the research provides critical insight into the information content of the EIA's energy supply forecasts. Second, the research method is unique relative to other studies examining the information content of public forecasts because it specifically addresses the incremental value added by forecasts beyond one period. Indeed, understanding the information content of multiple-period ahead forecasts is important for both forecast users and the EIA. If multiple-period ahead forecasts are found to provide little or no information relative to the one-period ahead forecasts, then it may be more appropriate for forecast users to simply extrapolate the one-period forecast to the desired forecast horizon. Furthermore, if certain multiple-period ahead forecasts are found to be just random adjustments to the one-period ahead forecasts, the EIA may be able to improve their forecasting efforts by allocating their forecasting efforts differently, or concentrating their efforts on markets where multiple-period horizons add the most value. Ultimately, the EIA can use the information from this research to refine their forecasting methods, or better allocate their scarce resources devoted to forecasting energy supply. Finally, and most importantly, practitioners and forecast users can use the information from this research to better utilize EIA production forecasts, pin pointing those which provide the greatest incremental information. In doing, economic decision-making can be improved.

The remainder of this paper is organized as follows. First, we present the direct test for the information content of multiple-period ahead forecasts as proposed by Vuchelen and Gutierrez (2005). Next, the EIA production forecast data used in this study are discussed, followed by results and discussion.

METHODS

In evaluating the performance of forecasts, researchers often focus on the concept of forecast rationality. A rational or optimal forecast is one which is unbiased and efficient (Diebold and Lopez, 1998) in that it does not consistently over- or under-estimate the actual value (unbiased) and utilizes all information available to the forecaster (efficient). Forecast rationality has traditionally been tested with the following regression:

$$A_{t+1} = \alpha + \beta F_t^{t+1} + u_{t+1}, \tag{1}$$

where, A_{t+1} is the realized value at time t+1, F_t^{t+1} is the forecast for time t+1 made at time t, and u_{t+1} is the error term (Mincer and Zarnowitz, 1969). This test maintains a joint null hypothesis that the forecast is both unbiased (α=0) and efficient (β=1). Furthermore, an efficient forecast is also characterized by an i.i.d. error term with no serial correlation in u_{t+1}. However, Holden and Peel (1990) have shown that this traditional joint null hypothesis is only a sufficient, but not necessary, condition for rationality. Therefore, a rejection of the null hypotheses in equation (1) does not lead to clear alternative statements about forecast properties. Consequently, Granger and Newbold (1986) suggest that researchers modify the traditional regression test in equation (1) focusing on forecast error terms (see Pons, 2000). Indeed, the tests presented in equation (1) are important in determining the optimal properties of a forecast. However, they provide little insight into the information content of the forecasts. This is particularly true when forecasts of the same variable are available for multiple periods ahead.

In examining the information content of the EIA's energy supply forecasts at different horizons, we apply the methods of Vuchelen and Gutierrez (2205) who develop a direct test for information content in multiple-horizon forecasts. In deriving their test, Vuchelen and Gutierrez first, express multiple k-period ahead forecasts (F_t^{t+k}) as simply the sum of consecutive adjustments to the most recent realized or actual observation (A_t). Thus, the one-period ahead forecast can be decomposed into the following components,

$$F_t^{t+1} = A_t + (F_t^{t+1} - A_t), \tag{2}$$

and for the two-period ahead forecast,

$$F_t^{t+2} = A_t + (F_t^{t+1} - A_t) + (F_t^{t+2} - F_t^{t+1}). \tag{3}$$

In equation (2), the one-period ahead forecast can be expressed as an adjustment to the current level. In other words, the one-period ahead forecast (F_t^{t+1}) is equal to the current level (A_t) plus the forecasted change from the current level (F_t^{t+1}-A_t). Similarly in equation (3), the two-period ahead forecast is equal to equation (2) plus the forecasted change in the following period which is merely the difference between the two-period and one-period ahead forecasts (F_t^{t+2} - F_t^{t+1}). Vuchelen and Gutierrez then substitute the decomposition in equation (2) into the traditional test of forecast rationality in equation (1) which yields the following regression equation:

$$A_{t+1} = \theta + \kappa A_t + \lambda(F_t^{t+1} - A_t) + u_t. \tag{4}$$

This resulting regression provides a framework for examining the forecasts optimal properties as well as its information content. An unbiased and efficient forecast is characterized by θ=0 and κ=λ=1, in which case equation (4) simplifies to the original forecast rationality test in equation (1). However, an informative one-step ahead forecast only requires that $\lambda \neq 0$. That is, $\lambda \neq 0$ implies that F_t^{t+1} is providing incremental information relative to just using A_t as a forecast by itself. Hence, a forecast need not be optimal to provide incremental information. In general, the Vuchelen and Gutierrez test is more revealing than traditional measures of forecast rationality in the sense that both forecast optimality and information content can be tested simultaneously.

The test as presented above in (4) only considers the information content of one-period ahead forecasts. However, Vuchelen and Gutierrez demonstrate that the test is easily extended to examine multi-period ahead forecasts. The test equation for two-period ahead forecasts is developed by substituting equation (3) into (1),

$$A_{t+2} = \gamma + \delta A_t + \eta(F_t^{t+1} - A_t) + \varepsilon(F_t^{t+2} - F_t^{t+1}) + u_{t+1}. \tag{5a}$$

In (5a), forecast rationality is tested under the null that γ=0, and δ=η=ε=1. Under the null hypothesis, equation (5a) simplifies to the two-period ahead version of equation (1). But in this case, equation (5a) tells us the amount of information that F_t^{t+2} provides relative to the most recent observation A_t, and the one-period ahead forecast F_t^{t+1}. If it is found that γ=η=ε=0 and δ=1, then the forecaster may as well use the most recent observation A_t, as the forecast (e.g., a naïve forecast). In other words, neither the one-period (F_t^{t+1}) or two-period ahead (F_t^{t+2}) forecasts provide any additional information relative to the naïve forecast (A_t). Whereas, if δ=η=1 and γ=ε=0 then the forecaster may as well just use the one-period ahead forecast F_t^{t+1}. If ε=0, then the two-period ahead forecast is not providing any incremental information relative to the one-period ahead forecast. This process is more easily seen by substituting equation (2) into equation (5a),

$$A_{t+2} = \gamma + \phi F_t^{t+1} + \mu(F_t^{t+2} - F_t^{t+1}) + u_{t+2}. \tag{5b}$$

In equation (5b), the null hypothesis of μ=0 suggests that the two-period ahead forecast, F_t^{t+2}, adds no incremental information relative to the one-period ahead forecast, F_t^{t+1}. Through repeated substitution, a direct test for the information content of k-period ahead forecasts can be generally expressed as:

$$A_{t+k} = \beta_1 + \beta_2 F_t^{t+k-1} + \beta_3 (F_t^{t+k} - F_t^{t+k-1}) + u_{t+k}\,. \qquad (5c)$$

Although not elaborated upon by Vuchelen and Gutierrez, equation (5c) implies an optimal weighting scheme for a composite forecast of F_t^{t+k} and F_t^{t+k-1} (Harvey and Newbold, 2000; Clemen, 1998). In essence A_{t+k} is a linear combination of F_t^{t+k} and F_t^{t+k-1}, with optimal weights of β_3 and (β_2-β_3), respectively. The null hypothesis that the k-period ahead forecast (F_t^{t+k}) adds no incremental value to the k-1 period ahead forecast (F_t^{t+k-1}) is equival ent to testing that its' weight in the composite forecast is zero, β_3=0.

TABLE 1.
SUMMARY OF HYPTOTHESIS TESTS,

$$A_{t+k} = \beta_1 + \beta_2 F_t^{t+k-1} + \beta_3 (F_t^{t+k} - F_t^{t+k-1}) + u_{t+k}$$

Hypothesis Test	Explanation
β_2=β_3=1, β_1=0	Failure to reject null hypothesis suggests that k-period ahead forecasts are rational.
β_2=β_3	Failure to reject null suggests that the weight on the k-1 ahead forecast (β_2-β_3) is zero in the implied composite forecast.
β_3=0	Failure to reject null suggests that there is no information contained in the incremental k-ahead forecast horizon.
β_3=1	Failure to reject null suggests that the k-ahead forecast is properly scaled.

Many alternative hypotheses can be tested using this method. These alternative hypothesis are summarized in table 1. For instance, if β_1=β_3=0 and β_2=1, then F_t^{t+k-1} is, by itself, the best predictor of A_{t+k}. In this research, we will focus on three primary questions. First, are the k-period ahead forecasts rational (β_1=0, β_2=β_3=1)? If so, then they can be used by themselves as the best predictor. Second, does the k-1 ahead forecast receive any weight in the implied composite forecasts, β_2=β_3? A failure to reject the null hypothesis, β_2=β_3, suggests that the weight (β_2-β_3) on the k-1 ahead forecast is zero. But, it may still be the case that the k-period ahead forecast also fails to provide information. So, the third test is against the null hypothesis that there is no information contained at the incremental forecast horizon (β_3=0), where β_3 is the composite forecast weight on the k-period ahead forecast. Finally, the k-period ahead forecast is tested for proper scale, β_3=1. That is, the k-period ahead forecast may provide information, $\beta_3 \neq 0$, but it may need to be scaled to provide a correctly calibrated forecast ($\beta_3 \neq$l). Of course, a rational k-period ahead forecast implies that the k-1 ahead forecast gets a zero weight (β_2=β_3), the k-period forecast provides information ($\beta_3 \neq 0$) and is properly scaled (β_3=1). Clearly, the direct test of Vuchelen and Gutierrez represents a flexible empirical method, allowing for considerable insight as to the amount and nature of the information content of the EIA's energy supply forecasts.

DATA

Using the direct test of Vuchelen and Gutierrez, quarterly production forecasts published by the EIA in the *Short-Term Energy Outlook* (*STEO*) are investigated. A broad range of energy forecasts are examined in order to compare the results across energy sectors. Specifically, supply forecasts are examined for crude oil, natural gas, electricity, and coal, and are taken from the U.S. Supply and Demand: Base Case tables presented in the *STEO*. For crude oil, we focus on the quantities labeled "total crude oil supply", and focus on "total supply" for natural gas, coal, and electricity. In these tables, total supply quantities include domestic production, imports, adjustments to strategic reserves, and other stock adjustments. Therefore, these forecasts represent the EIA's expectations for energy supplies that will be made available to the marketplace for consumption.

The *STEO* is published on a monthly basis, and is usually released between the 2nd and the 12th of each month. Each monthly report contains *quarterly* forecasts for at least four-quarters ahead. Given this, we focus on forecasts for one-quarter ahead, two-quarters ahead, three-quarters ahead, and four-quarters ahead. Because the supply forecasts are for calendar quarters, forecasts for each energy commodity are collected from the January, April, July, and October reports. So, for example, from the January report one- through four-quarter ahead forecasts are collected. The corresponding one-quarter ahead forecast would cover the months of January, February, and March, the two-quarter ahead forecasts would reflect expectations of supply for the months of April, May, and June, and so forth. From the April report, the one-quarter ahead forecasts would reflect expectations for the months of April, May, June, and the two-quarter ahead forecasts would be for the months of July, August, September, and so forth. The actual or realized supplies are then taken from subsequent releases of the *STEO* reports, and are matched to the quantities that the EIA is attempting to forecast. The sample period is from the second quarter of 1997 (1997.2) to the first quarter of 2006 (2006.1), resulting in 37 observations of one-quarter ahead forecasts and realized values for each energy product examined (crude oil, natural gas, coal, and electricity). An observation is lost for each additional horizon, resulting in a sample of 36, 35, and 34 observations for the two-, three-, and four-quarter ahead forecasts, respectively.

Energy supplies are known to demonstrate seasonal patterns as a result of natural fluctuations in production and in anticipation of seasonal demand loads. For instance, electricity production in the U.S. usually expands in the summer months to satisfy increased cooling loads as a result of warmer temperatures. Therefore, the analyses focuses on seasonal differences defined as the log-relative supply growth change from the same quarter of the prior year. Given this, AS_{t+1} is defined as the actual supply level in quarter t+1, and FS_t^{t+1} is the one-period ahead supply forecast for quarter t+1. The change in actual supplies is defined as $A_{t+1}=\ln(AS_{t+1}/AS_{t-3})$, and the forecasted supply change is $F_t^{t+1}=\ln(FS_t^{t+1}/A_{t-3})$. Thus, changes reflect the percent change in the quarterly supplies from the prior year. As defined, these variables are used in the direct test as shown in equation (5c).

RESULTS

The information content at each forecast horizon (one- through four-quarters ahead) for each of the energy products is evaluated using the direct test of Vuchelen and Gutierrez. As discussed by Vuchelen and Gutierrez, Equation (4) can be estimated using standard OLS procedures. However, versions of equation (5c) with k>1 are characterized by overlapping forecast horizons, which will result in correlated forecast errors and subsequent biased and inconsistent standard errors. To correct this problem, we follow the lead of Brown and Maital (1981) and use the OLS coefficient estimates, but correct the variance-covariance matrix using the methods proposed by Hansen (1982) and demonstrated by Hansen and Hodrick (1980).

At each horizon, equation (5c) is estimated and the following restrictions are tested to reveal the information content and rationality of the k-period ahead forecast. First, the k-period ahead forecast is tested for rationality under the null hypothesis, $\beta_2=\beta_3=1$, $\beta_1=0$. A failure to reject the null hypothesis suggests that the k-period ahead forecast contains unique information and meets the requirements for rationality. Moreover, if the rational null hypothesis is not rejected, then the k-period ahead forecast can be used by itself. Second, we focus on the composite use of the k-1 and k-period ahead forecast by testing that ($\beta_3=\beta_2$), where a failure to reject the null indicates a zero weight on the k-1 ahead forecast. Finally, the minimum requirement for information contained in the k-period ahead forecast is tested under the null that $\beta_3=0$, and we also test if it is properly scaled with a value of unity ($\beta_3=1$). Collectively, these tests will reveal how, if at all, practitioners should utilize the multiple-step ahead forecasts provided by the EIA.

Crude Oil

The results for the crude oil supply forecasts are presented in table 2. The parameter estimates and the corrected standard errors are shown in the first three columns, and the p-values from the parameter restrictions are displayed in the remaining columns. As shown under the first set of parameter restrictions, the null hypothesis of rationality for the k-period ahead forecast is rejected at each horizon at the 5% level. So, at no horizon, should a forecast user rely exclusively on the EIA energy supply forecast for that horizon. Indeed, at horizons of one and two quarters, a composite forecast using the k-1 period ahead forecast along with k-period ahead forecast is implied by failure to reject the hypothesis that $\beta_2\neq\beta_3$. For k=1, the composite forecast would have a weight of 0.5237 (β_2-β_3) on A_t and 0.5614 (β_3) on the one-quarter ahead forecast. At this horizon, the forecast is providing unique information ($\beta_3\neq0$), but it is not properly scaled ($\beta_3\neq1$); hence, a composite forecast is implied by the direct test. At the two-quarter horizon (k=2), a similar result is found. Again, the two-quarter ahead forecast is not rational, and a composite forecast using the one and two-quarter ahead forecast is implied by the rejection of the null hypothesis that $\beta_2=\beta_3$. However, at the two-quarter horizon, $\beta_3=0$ is not rejected at the 5% significance level (p-value =0.1389), which suggests that the two-quarter ahead forecast provides little information relative to the one-quarter ahead forecast. Indeed, at the k=2 horizon, the weight on the one-quarter ahead forecast is 0.4353 (β_2-β_3), while the weight is 0.3425 (β_3) on the two-quarter ahead forecast which is statistically insignificant. For three- and four-quarter horizons, the k-1 forecasts do not receive any weight in the composite forecast ($\beta_2=\beta_3$), and the k-period ahead forecast does not reject the null of no incremental information ($\beta_3\neq0$) at the 5%

significance level. For example, for k=3, a forecast user can ignore the two-quarter ahead forecast, and the three-quarter ahead forecast needs to be scaled by 0.3824, which is of marginal statistical significance. That is a forecast for a three-quarter ahead growth in crude oil supply of 5% should be adjusted to 1.912% (0.05 x 0.3824). Given these results, the evidence suggests that there is statistically significant information at the one-quarter horizon. However, little information is provided at multiple horizons (k>1). Additionally, at no horizon are the forecasts fully rational. Collectively, the crude oil results suggest that EIA supply forecasts provide little incremental value beyond the one-quarter horizon.

TABLE 2
CRUDE OIL SUPPLY FORECAST TEST,

$$A_{t+k} = \beta_1 + \beta_2 F_t^{t+k-1} + \beta_3 (F_t^{t+k} - F_t^{t+k-1}) + u_{t+k}$$

	Coefficient Estimates			Hypothesis Tests			
Horizon	**β_1**	**β_2**	**β_3**	**$\beta_2=\beta_3=1, \beta_1=0$**	**$\beta_2=\beta_3$**	**$\beta_3=0$**	**$\beta_3=1$**
k=1	-0.0026 (0.0032)[a]	1.0851 (0.1634)	0.5614 (0.1662)	0.0002[b]	0.0002	0.0019[c]	0.0126
k=2	-0.0023 (0.0051)	0.7778 (0.2595)	0.3425 (0.2257	0.0000	0.0029	0.1389	0.0065
k=3	-0.0046 (0.0068)	0.6747 (0.3066)	0.3824 (0.2080)	0.0003	0.2467	0.0756	0.0057
k=4	-0.0057 (0.0069)	0.6029 (0.2356)	0.3361 (0.1662)	0.0000	0.2255	0.0523	0.0003

Note: For horizon k=1, F_t^{t-1} is replaced by A_t.
[a]Standard errors are in parenthesis.
[b]P-value from Chi-squared test for stated restriction
[c]P-value from two-tailed t-test on stated restriction.

Natural Gas

The natural gas supply forecasts (table 3) display comparable results regarding the information content and quality of the EIA's multiple-period ahead forecasting efforts. At each forecasting horizon, the k-period ahead forecast is not rational and should not be used by itself or without a weighting or scale adjustment. As noted, rationality ($\beta_2=\beta_3=1$, $\beta_1=0$) of the k-period ahead forecast is rejected at each horizon. However, close inspection of the individual coefficients reveals that the forecasts tend to be biased, with β_1 large relative to its standard error. Absent the bias, the one-quarter horizon forecast, by itself, may provide the best forecast. That is, the weight on the k-1 ahead forecast (β_2-β_3) is not statistically different from zero at the 5% level. So, at this horizon, $\beta_3\neq0$ and we cannot reject that $\beta_3=1$; thus, the one-quarter ahead forecast can be used with a bias (intercept adjustment). If the EIA is forecasting a 5% increase in natural gas supplies, then the bias corrected forecast is 3.08% (0.05 – 0.0192). At the two-quarter horizon (k=2), rationality is again rejected, but at this horizon a composite forecast is implied ($\beta_2\neq\beta_3$). At the two-quarter horizon, the k-quarter ahead forecast provides unique information ($\beta_3\neq0$), but it should be combined with the one-quarter horizon forecast with weights of 0.6077 and 0.4843, respectively.

At the three- and four-quarter horizons, the k-1 ahead forecast receives no weight in the implied composite forecast ($\beta_2=\beta_3$). At these horizons, the k-ahead forecast does provide important information ($\beta_3\neq0$). Moreover, a "weight" of unity ($\beta_3\neq1$) cannot be rejected at the 5% level. Therefore, beyond two-quarters ahead, the natural gas supply forecasts can be used by themselves with just an adjustment for the intercept or bias term (β_1).

TABLE 3
NATURAL GAS SUPPLY FORECAST TEST,

$$A_{t+k} = \beta_1 + \beta_2 F_t^{t+k-1} + \beta_3 (F_t^{t+k} - F_t^{t+k-1}) + u_{t+k}$$

	Coefficient Estimates			Hypothesis Tests			
Horizon	**β_1**	**β_2**	**β_3**	**β_2=β_3=1, β_1=0**	**β_2=β_3**	**β_3=0**	**β_3=1**
k=1	-0.0192 (0.0086)[a]	1.2749 (0.2761)	1.0029 (0.2676)	0.0153[b]	0.0648	0.0007[c]	0.9915
k=2	-0.0238 (0.0070)	1.0920 (0.2260)	0.6077 (0.1648)	0.0000	0.0243	0.0008	0.0234
k=3	-0.0204 (0.0114)	0.6977 (0.2640)	0.6483 (0.1832)	0.0005	0.7843	0.0013	0.0642
k=4	-0.0205 (0.0126)	0.5174 (0.2715)	0.6151 (0.1969)	0.0000	0.5955	0.0039	0.0601

Note: For horizon k=1, F_t^{t-1} is replaced by A_t.
[a]Standard errors are in parenthesis.
[b]P-value from Chi-squared test for stated restriction
[c]P-value from two-tailed t-test on stated restriction.

Coal

The coal supply forecast results (table 4) are remarkably similar across horizons. That is, at each horizon, rationality is rejected (β_2=β_3=1, β_1=0), the weight on the k-1 ahead forecast is zero (β_2=β_3), the k-period ahead forecasts do generally provide unique information ($\beta_3 \neq 0$), but they do not receive a unitary weight ($\beta_3 \neq 1$). The lone exception is at the k=3 horizon, where the informational value of the three-quarter horizon is not statistically significant at the 5% level. At this horizon, there appears to be very little informational content within the EIA supply forecasts. Otherwise, for the coal forecasts, the prescribed usage by practitioners is to use the k-period ahead forecast exclusively with a weight or scale equal to β_3. Importantly, there is also a statistically significant bias for k=2 and k=3. So, at these horizons, the k-period ahead forecast needs a bias adjustment as well. For instance, a DOE forecast for a 10% increase in coal supplies two quarters hence (k=2), would be adjusted to 4.5% (0.10 x 0.3600 +0.0090).

TABLE 4
COAL SUPPLY FORECAST TEST,

$$A_{t+k} = \beta_1 + \beta_2 F_t^{t+k-1} + \beta_3 (F_t^{t+k} - F_t^{t+k-1}) + u_{t+k}$$

	Coefficient Estimates			Hypothesis Tests			
Horizon	**β_1**	**β_2**	**β_3**	**β_2=β_3=1, β_1=0**	**β_2=β_3**	**β_3=0**	**β_3=1**
k=1	0.0071 (0.0054)[a]	0.4906 (0.2342)	0.3824 (0.1758)	0.0058[b]	0.5069	0.0370[c]	0.0013
k=2	0.0090 (0.0043)	0.1885 (0.1798)	0.3600 (0.1275)	0.0000	0.1908	0.0081	0.0000
k=3	0.0099 (0.0043)	0.1155 (0.1437)	0.2007 (0.1048)	0.0005	0.5396	0.0647	0.0000
k=4	0.0054 (0.0041)	0.2618 (0.1359)	0.3541 (0.1152)	0.0000	0.3447	0.0045	0.0000

Note: For horizon k=1, F_t^{t-1} is replaced by A_t.
[a]Standard errors are in parenthesis.
[b]P-value from Chi-squared test for stated restriction
[c]P-value from two-tailed t-test on stated restriction.

Electricity

The one-quarter ahead and two-quarter ahead electricity forecasts are the only energy supply forecasts that provide a rational forecast (table 5). There is a failure to reject the rationality conditions at the 5% level for each of these forecasts. Therefore, the one- and two-quarter ahead forecasts can be used by themselves as the best predictors. At no horizon, however, is there strong evidence for using a composite forecast ($\beta_2 \neq \beta_3$). Indeed, at the two- and three-quarter horizons, the results suggest that the k-period ahead forecast should be used with a weight or scaling of β_3 because the forecast provides information ($\beta_3 \neq 0$) but β_3 does not equal unity. For instance, at the four-quarter horizon, the DOE forecasts should be used by themselves, but scaled by a factor of 0.5218.

TABLE 5
ELECTRICITY SUPPLY FORECAST TEST,

$$A_{t+k} = \beta_1 + \beta_2 F_t^{t+k-1} + \beta_3 (F_t^{t+k} - F_t^{t+k-1}) + u_{t+k}$$

	Coefficient Estimates			Hypothesis Tests			
Horizon	**β_1**	**β_2**	**β_3**	**$\beta_2=\beta_3=1, \beta_1=0$**	**$\beta_2=\beta_3$**	**$\beta_3=0$**	**$\beta_3=1$**
k=1	0.0017 (0.0059)[a]	0.7931 (0.2366)	0.6488 (0.2121)	0.3217[b]	0.3528	0.0044[c]	0.1073
k=2	0.0055 (0.0059)	0.4769 (0.2210)	0.7060 (0.1876)	0.0677	0.1336	0.0007	0.1268
k=3	0.0139 (0.0070)	0.0821 (0.2987)	0.4688 (0.1934)	0.0104	0.0646	0.0214	0.0099
k=4	0.0055 (0.0064)	0.4576 (0.2550)	0.5218 (0.1683)	0.0195	0.7356	0.0042	0.0080

Note: For horizon k=1, F_t^{t-1} is replaced by A_t.
[a]Standard errors are in parenthesis.
[b]P-value from Chi-squared test for stated restriction, [c]P-value from two-tailed t-test on stated restriction.

Discussion of Results

Collectively, the results for crude oil, natural gas, coal, and electricity supplies lead to a couple of general conclusions. First, other than the one- and two-quarter ahead electricity forecasts, there is no evidence that the DOE multiple-horizon forecasts are rational ($\beta_2=\beta_3=1$, $\beta_1=0$). Second, the multiple-horizon forecasts in crude oil appear to contain the least amount of information, with only the one-quarter horizon showing strong evidence of marginal information. In the other markets, there is generally evidence that valuable information is contained in multi-period ahead forecasts that should not be ignored by forecast users. However, the results also indicate that the information provided is not rational and must either be used in a composite forecast with the k-1 period ahead forecast, or it must be appropriately adjusted for bias and scaling effects. It is important that forecast users understand and implement these adjustments to garner the best possible predictor from the DOE forecasts.

In general, the findings support the efficacy of the DOE's forecasting efforts at multiple horizons. Broadly speaking, the most information is contained in the multi-period electricity forecasts, while the crude oil supply forecasts seem to provide little incremental information at each horizon. Without knowing the exact model used by the EIA in forecasting crude oil supplies, it is difficult to determine why the more distant forecast horizons do not contain unique information relative to the one quarter forecast. However, given the volatility of the crude oil market and the numerous economic and political factors that influence crude oil supplies, the EIA may be focusing their efforts on the one quarter forecasts and making reasonable adjustments to this forecast for longer horizons. For the other energy supplies, the fact that unique information is indeed contained in the multi-period forecasts suggests that these forecasts are not mere adjustments to the one quarter forecast. That is, the forecasts used by the EIA are specified such that they adequately capture information that is unique to a particular forecast horizon. Although the forecasts are not entirely rational—that is, they should often be used in a composite forecast with those forecasts at alternative horizons—the data clearly suggests that the multi-period ahead forecasts are providing some unique information, especially for natural gas, electricity, and coal.

SUMMARY AND CONCLUSIONS

The Department of Energy's Energy Information Administration provides the public with forecasts of energy supplies in their *Short-Term Energy Outlook* publication. Public outlook information is designed to aid in the efficient allocation of resources by providing businesses, governments, forecasters, and other entities with valuable information to aid in decision making. The importance of accurate public outlook information is apparent in the energy sector, especially considering the systematic impact of the energy sector on the entire economy.

In this research we examine one primary aspect of the DOE's overall outlook efforts by determining the efficacy of supply forecasts for crude oil, natural gas, coal, and electricity. We specifically focus on the information content of these forecasts over multiple forecast horizons. Indeed, while there is a considerable amount of literature examining the performance of publicly available forecasts in other arenas, namely agricultural commodities and macroeconomic variables, little effort has specifically focused on the performance of public forecasts in the energy sector. In particular, there is little if any published academic research examining the information content of multi-period forecasts for the energy sector—a sector that has a tremendous systematic impact on the economy as a whole. Therefore the results of this research provide an important look into the performance and information content of the EIA's multiple-horizon forecasting efforts, and help to fill a gap that has existed in the literature.

In determining the information content of these supply forecasts, we use the methods put forth by Vuchelen and Gutierrez (2005) for testing information content in multiple-period ahead forecasts. For each of the supply forecasts (crude oil, natural gas, coal, and electricity) we examine the information content in one-quarter, two-quarter, three-quarter, and four-quarter ahead forecasts. In particular, the Vuchelen and Gutierrez test considers the information content of the more distant forecast relative to the previous quarter's forecast. For instance, if the two-quarter ahead forecast is shown to provide no incremental information relative to the one-quarter ahead forecast, then only the one-quarter ahead forecast is needed. While not elaborated by Vuchelen and Gutierrez, we demonstrate how their test can be used to consider the possibility that a composite forecast may be warranted. Considering the above example of a one- and two-quarter ahead forecast, there may be cases where both the one-quarter and two-quarter ahead forecasts could be used in a composite forecast which would could potentially yield a more accurate forecast in a mean-square error framework than using the two-quarter ahead forecast on its own.

By and large, the results show that the EIA's energy supply forecasts at multiple horizons do indeed contain unique information. The exception seems to be in the crude oil supply forecasts, where information is only present at the one-quarter horizon (5% level). However, without knowing the exact forecasting procedures used by the EIA for crude oil, it is difficult to speculate why the more distant forecast horizons contain no unique information relative to the one quarter forecast. In contrast, the electricity forecasts tend to be the most rational as well as providing unique information at each horizon. While the evidence strongly suggests that several of the energy supply forecasts are not rational, the empirical procedures supply the appropriate adjustments, whether through a composite forecast, scaling, or bias corrections. Thus forecast users should look to make these types of adjustments when utilizing DOE forecasts. Moreover, the DOE should also consider this result, and perhaps adjust their forecasts accordingly. Still, the DOE appears to be fulfilling their goal of providing valuable and incremental information to the energy industry through their multiple-horizon forecasts.

REFERENCES

Aaron, H.J. 2000. "Presidential Address—Seeing Through the Fog: Policymaking with Uncertain Forecasts." *Journal of Policy Analysis and Management*. 19: 193-206.

Baur, R.F., and P.F. Orazem. 1994. "The Rationality and Price Effects of U.S. Department of Agriculture Forecasts of Oranges." *Journal of Finance* 49: 681-695.

Bentzen, J., and H. Linderoth. 2005. "Has the Accuracy of Energy Demand Projections in the OECD Countries Improved Since the 1970s?" Working Paper, The Aarhus School of Business Working Paper, Department of Economics, Denmark (http://www.hha.dk/nat/WPER/01-5_jb.pdf), retrieved December 20.

Brown, B.W., and S. Maital. 1981. "What do Economists Know? An Empirical Study of Experts' Expectations." *Econometrica* 49, 491-504.

Carter, C.A., and C.A. Galopin. 1993. "Information Content of Government Hogs and Pigs Reports." *American Journal of Agricultural Economics* 75: 711-718.

Caruso, G.F. 2005. Retrieved December 28 Energy Information Association website, http://www.eia.doe.gov/neic/aboutEIA/guidelines.html.

Clemen, R.T. 1998. "Combining Forecasts: A Review and Annotated Bibliograhpy." *International Journal of Forecasting* 5: 559-583.

Diebold, F.X., and J.A. Lopez. 1998. "Forecast Evaluation and Combination." *Handbook of Statistics 14: Statistical Methods in Finance*, G.S. Maddala and C.R. Rao, eds. Amsterdam: North-Holland.

Floris, F.J.T., M.D. Bush, M. Cuypers, F. Roggero, and A-R Syversveen. 2001. "Methods for Quantifying the Uncertainty of Production Forecasts: A Comparative Study." *Petroleum Goescience* 7: 87-96.

Freebairn, J.W. 1978 "An Evaluation of Outlook Information for Australian Agricultural Commodities." *Review of Marketing and Agricultural Economics* 46: 294-314.

Garcia, P., S.H. Irwin, R.M. Leuthold, and L. Yang. 1997. "The Value of Public Information in Commodity Futures Markets." *Journal of Economic Behavior and Organization* 32: 559-570.

Granger, C.W.J. 1996. "Can We Improve the Perceived Quality of Economic Forecasts?" *Journal of Applied Econometrics* 1: 455-473.

Granger, C.W.J., and P. Newbold. 1986. *Forecasting Economic Time Series* (2nd ed.). New York: Academic Press.

Hansen, L.P. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50: 1029-1054.

Hansen, L.P., and R. Hodrick. 1980. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy* 88: 829-853.

Harvey, D.I. and P. Newbold. (2000). "Tests for Multiple Forecast Encompassing." *Journal of Applied Econometrics*. 15: 471-482.

Holden, K., and D.A. Peel. 1990. "On Testing for Unbiasedness and Efficiency of Forecasts." *The Manchester School* 15: 120-127.

Lynch, M.C. 2002. "Forecasting Oil Supply: Theory and Practice." *Quarterly Review of Economics and Finance* 42: 373-389.

Mills, T.C., and G.T. Pepper. 1999. "Assessing the Forecasters: An Analysis of the Forecasting Records of the Treasury, the London Business School and the National Institute." *International Journal of Forecasting* 15: 247-257.

Mincer, J., and V. Zarnowitz. 1969. "The Evaluation of Economic Forecasts." *Economic Forecasts and Expectations*. National Bureau of Economic Research, New York.

Pettee, I.E.W. 1936. "Long-Term Commodity Price Forecasting—1850 to 1930." *The Journal of Business*. 9: 95-113.

Pons, J. 2000. "The Accuracy of IMF and OECD Forecasts for G7 Countries." *Journal of Forecasting* 19: 53-63.

Department of Energy. 2005. "Annual Energy Outlook Forecast Evaluation." Retrieved December 28 from Energy Information Association website http://www.eia.doe.gov/oiaf/analysispaper/forecast_eval.html.

Sanchez, E. 2006. "Annual Energy Outlook Forecast Evaluation." Retrieved January from U.S. Department of Energy web site. http://tonto.eia.doe.gov/FTPROOT/forecasting/forecast_eval2k2.pdf.

Schuh, S. 2001. "An Evaluation of Recent Macroeconomic Forecast Errors." *New England Economic Review* January/February: 35-56.

Shlyakhter, A.I., D.M. Kammen, C.L. Broido, and R. Wilson. 1994 "Quantifying the Credibility of Energy Projections from Trends in Past Data." *Energy Policy* February: 119-130.

Smyth, D.J. 1973. "Effect of Public Price Forecasts on Market Price Variation: A Stochastic Cobweb Example." *American Journal of Agricultural Economics* 55: 83-88.

Stein, J.L. 1981. "Speculative Price: Economic Welfare and the Idiot of Chance." *Review of Economics and Statistics* 63: 223-232.

Sumner, D.A., and R.A. Mueller. 1989. "Are Harvest Forecasts News? USDA Announcements and Futures Market Reactions." *American Journal of Agricultural Economics* 71: 1-8.

Vuchelen, J., and M-I Gutierrez. 2005. "A Direct Test of the Information Content of the OECD Growth Forecasts." *International Journal of Forecasting* 21: 103-117.