
To: Patrick Barickman
From: Jon Wilkey
Date: November 2, 2016
RE: October 2016 monthly report on price forecasting for the R model

During the month of October work on this project focused on analyzing the various methods of creating energy price forecasts (EPFs) identified in literature review conducted during the month of September. A summary of this analysis is presented below, categorized by EPF type (structural, time-series, and user-input) and concluding with a comparison of the relative merits of each approach. Based on these results, the recommended methods for generating EPFs are given in the table below:

EPF Method	Number of Years into the Future that EPF is Being Made			
	0 – 1	1 – 5	5 – 10	10 – 20+
Structural		X	X	
Time-Series			X	X
User-Input	X			

Work for the month of November is expected to focus on adding options for each type of EPF method to the R model, documenting model changes in the User Manual, and final delivery of the R model repository.

Structural EPFs

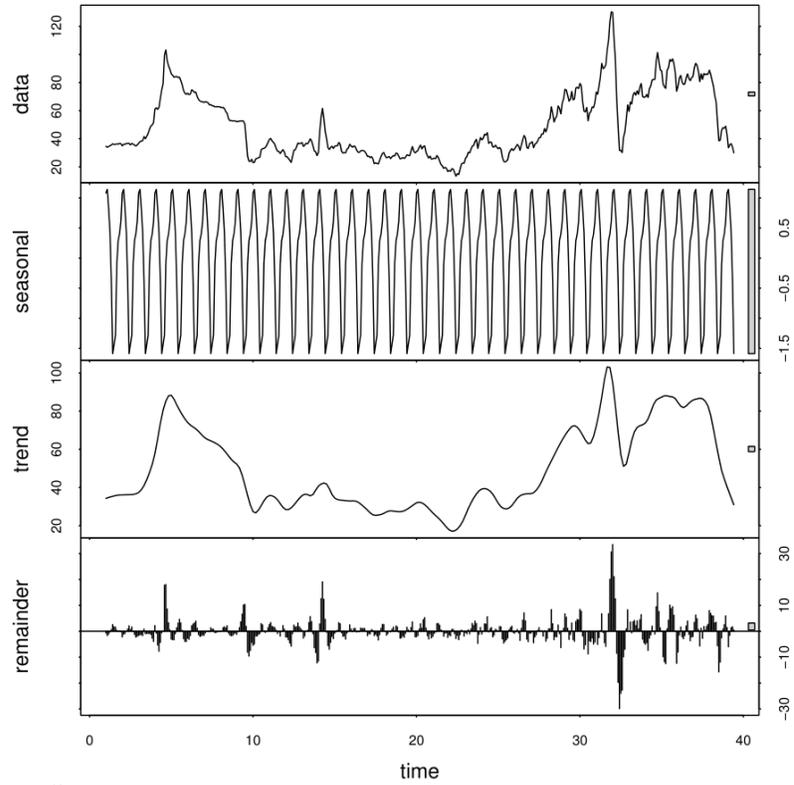
EPFs in this category attempt to model the dynamics of supply and demand that ultimately lead to observed energy prices. Of the long-term, publicly accessible EPFs reviewed in last month's report, only U.S. EIA and U.K. Dept. of Energy and Climate Change (DECC) attempt to apply independent structural EPFs that could be used in the R model. However, given that the DECC has only published four EPFs (2012 – 2015), there is little basis for assessing the accuracy of the DECC's EPFs and the R model's current EIA-based methodology remains the best EPF in this category.

Time-Series EPFs

In a time-series forecast, the future values of some quantity is predicted solely based on the past values of that same quantity. In general, the quantity to be forecasted is assumed to be comprised of (a) seasonal variation, (b) an underlying trend, and (c) noise. There are a number of methods for fitting the underlying trend, but in general the most widely used is the ARIMA method (autoregressive integrated moving average), which utilizes the past values (p), differences in past values (d), and the moving average of the fitting error (q) as predictor variables. ARIMA models are classified based on how many of each type of these predictor variables (p, d, q) they use, and certain combinations are known to have characteristic behaviors (e.g. they decay to zero, result in a random walk around a constant mean, etc.).

To applying the time-series approach to our specific oil and gas price history, we begin by taking the decomposition of past oil and gas prices into the seasonal/trend/noise components as shown in Figure 1 on the following page.

(a)



(b)

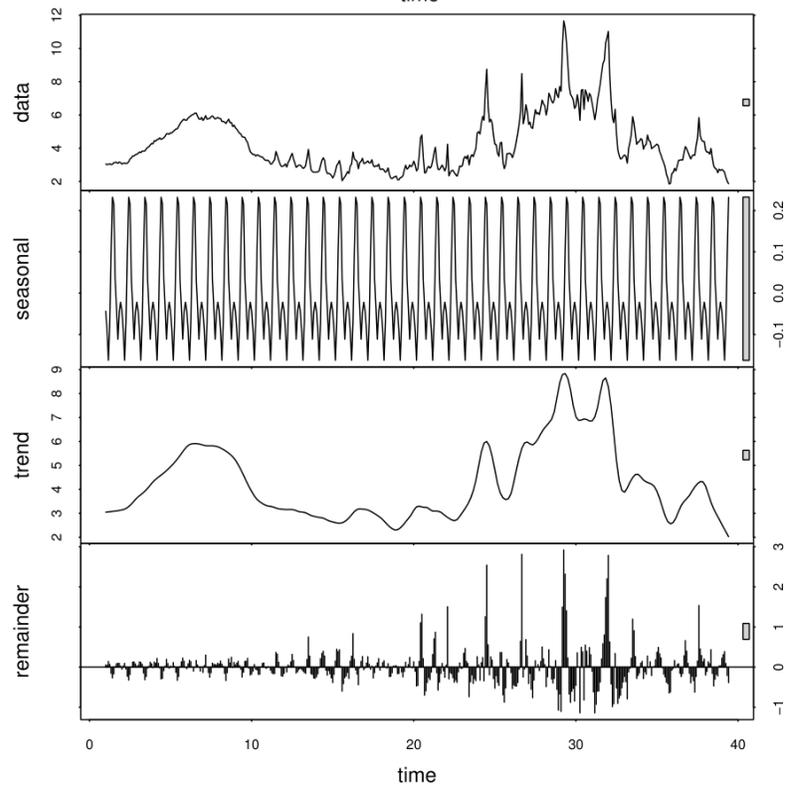


Figure 1: Time-series decomposition of (a) oil (\$ / bbl) and (b) natural gas (\$ / MCF) prices in the Uinta Basin. Note that the "time" values on the x-axis are years since 1977.

From these decompositions, we can draw the following conclusions:

1. Seasonal variation in oil prices ($\pm \$1.5/\text{bbl}$) and natural gas prices ($\pm \$0.2/\text{MCF}$) is negligible.
2. Over their entire price history, the trend for both oil and gas is flat. There are clearly periods where prices went up during the early 1980s and mid 2000s, but long-term there is no clear upward/downward trajectory.
3. There is substantial noise in the remainder term ($\pm \$30/\text{bbl}$ for oil and $+\$3 -\$1 / \text{MCF}$ for gas).

Consequently, we can expect that an ARIMA time-series fit to the price data will project a constant future price (as the median forecast) with substantial variation around that constant value.

The best method I've found for fitting ARIMA(p,d,q) models in R is using the forecast¹ package, which allows the user to either fit a specified ARIMA model or to automatically select the best ARIMA model based on AIC (Akaike information criteria) values (a measure of the goodness of fit for a statistical model). An example twenty year forecast for both oil and gas prices is shown in Figure 2, which was generated by testing various ARIMA models on the oil and gas price histories from July 1977 (the start of the oil and gas price history data set) to December 1995.

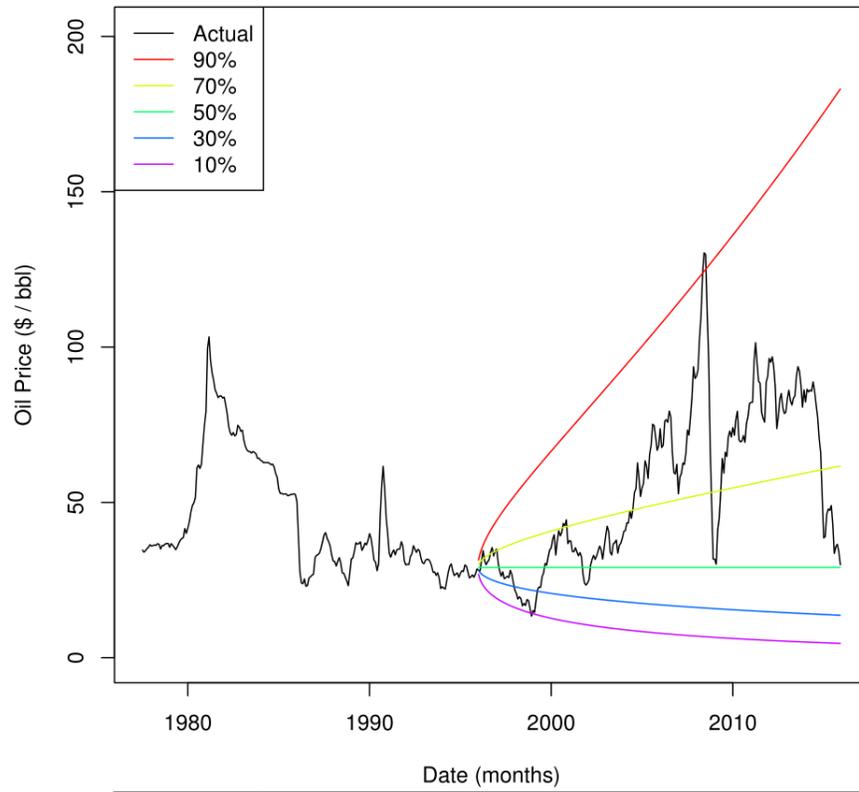
It should be noted that the auto-ARIMA fit function can produce substantially different fits based on the selection of training data, and that the best fit for a given set of training data (as measured by AIC values) may not actually be the best match to the test data used for cross-validation. For example, the auto-ARIMA function found that the best fit for oil prices given the 1977-1995 data was a (0,1,1) model, which produces the results shown in Figure 2(a) that fully encompass the observed test prices within the 10th-90th percentiles. For natural gas, the best auto-ARIMA fit for over that same period is a (1,1,4) model (see Figure 3), which produces a narrower price forecast than most of the observed natural gas prices during the test period. For that reason, the use of a (0,1,1) model as in Figure 2(b) is preferable. When the auto-ARIMA is fit to the entire price history data set (1977-2015), we again get a different set of best-fit models for oil (1,1,0) and gas (0,1,2) prices. I'll be performing additional work over the coming month to clarify whether a specific ARIMA fit should be used or if the model should be allowed to select the ARIMA model automatically.

User-Input EPFs

The final option for handling EPFs is to allow the user to specify the price directly. The R model can already support this approach as-of the version deliver to UDAQ in July, however it could be further expanded to support additional options such as directly specifying the number of wells drilled each month.

¹ Hyndman RJ and Khandakar Y (2008). "Automatic time series forecasting: the forecast package for R." *Journal of Statistical Software*, **26** (3), pp. 1-22. <http://www.jstatsoft.org/article/view/v027i03>

(a)



(b)

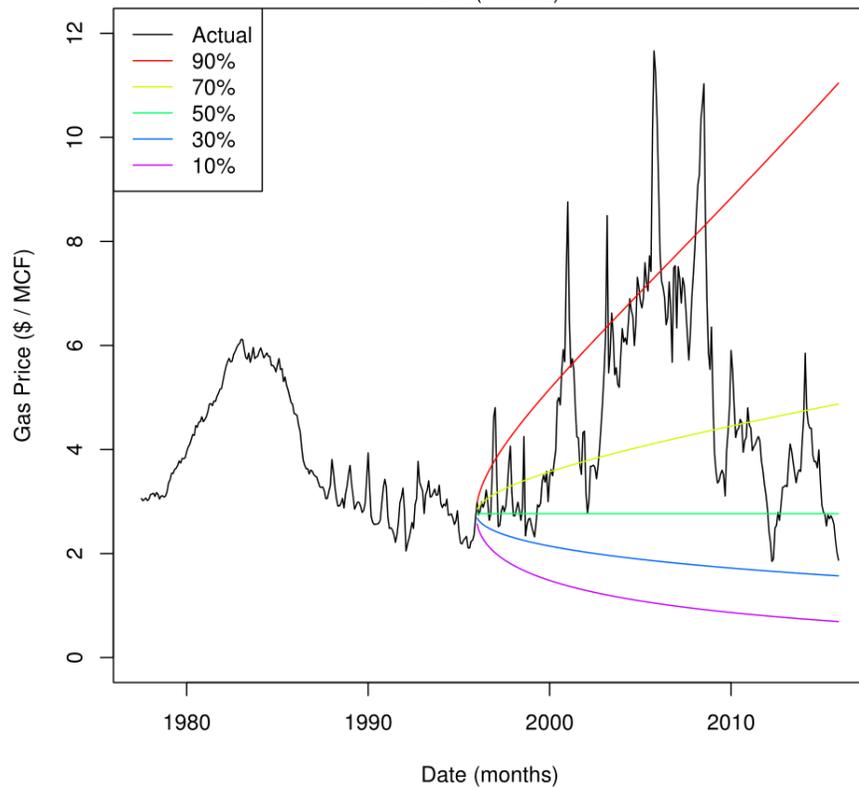


Figure 2: Twenty year ARIMA(0,1,1) forecasts of (a) oil and (b) natural gas prices.

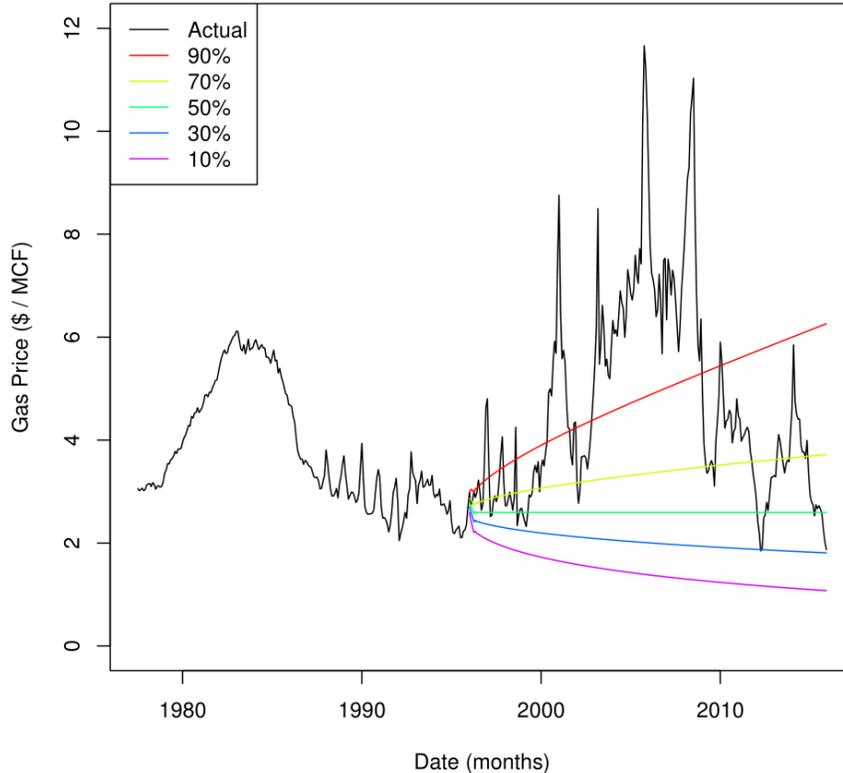


Figure 3: ARIMA(1,1,4) forecast for natural gas prices using the same training set as Figure 2(b). Note that the result produces a narrower forecast than actually observed during the majority of the 20-year test period. The (1,1,4) model is the best fit to the training data, but performs worse than the (0,1,1) model in Figure 2(b) during the test period.

Comparison of EPF Methods

The best choice of EPF model depends on the time-frame over which model predictions are being made:

1. Annual forecast (≤ 1 year)
 - a. User-input is highly preferable. Within a 1 year period it's reasonable to make any of the following simplifying assumptions to eliminate the uncertainty and complexity of EPFs from the R model:
 - i. Hold current oil and gas prices constant
 - ii. Assume prices are equal to EIA's forecast (on average EIA's relative error 1 year out is only 15%)
 - iii. Directly specify the number of new wells you expect to see in the Uinta Basin over the coming year.
2. Short-term forecast (1 – 5 years)
 - a. Up to five years, EIA's structural forecast is preferred.
 - i. There's substantial uncertainty in EPFs over a 5-year period, so user-input is likely to be wrong unless the user has access to more detailed and direct knowledge about future development than is publicly available.

- ii. There are enough observations about EIA's forecasting error to reasonably incorporate that error into their structural price forecast.
 - iii. EIA's structural forecast can predict trends up/down which the time-series approach tends to ignore (given the past price history).
- 3. Medium-term forecast (5 – 10 years)
 - a. Time-series is slightly preferred.
 - i. There are enough observations about EIA's error rate to use the structural forecast, but the error rates are skewed because the 1999-2005 forecasts from EIA are all essentially the same forecast.
 - ii. However, EIA has made some forecasts (such as their 2008 forecast) that have performed reasonably well over the medium-term.
 - iii. There are no downsides to the time-series approach other than (again) the tendency of the ARIMA fits to predict a flat constant price.
- 4. Long-term forecast (10 – 20 years)
 - a. Time-series approach is highly preferable.
 - i. At twenty years, all models will be wrong, but as the cross-validation tests with the ARIMA(0,1,1) fit show in Figure 2, at least the time-series approach can fully encompass the range of observed prices.
 - ii. There are too few observations to use the current EIA-based EPF method.
 - 1. Beyond 16 years it's mathematically impossible to use (at least two observations of the observed error rate are required in order to estimate a distribution)
 - 2. Even where mathematically possible, the error rates are not representative beyond 10 years (at which point every forecast from EIA is the identical 1999-2005 forecast)
 - 3. At some point in the (distant) future there will be enough information to use this approach.