Bidding and Winning in GOM Deep Water Lease Sales

John D. Grace, Scott Morris and Tony Dupont
Earth Science Associates, Long Beach, CA

Deep water dominates exploration and production in the US Gulf of Mexico, as it has for two decades. In 1999, its oil output reversed the Gulf-wide decline that started in 1986. The area covered by at least 200 meters (656 feet) of water has yielded the Gulf's largest discovery and over a dozen other giants each with estimated ultimate recovery over a quarter-billion barrels of oil equivalent. Exploratory success further and further down the continental slope has motivated engineering advances that continually break worldwide water depth records. Despite the crash of oil prices in 2015-2016, and sky-rocketing competition for investment from onshore shale plays, deep water oil production in 2018 was over 1.6 million barrels per day *and rising* – contributing more than one of every six barrels of growing US domestic supply.

Mineral rights to the Gulf's resources, beyond a narrow littoral, belong to the US federal government. Since significant offshore activity began in the late 1940s, the government has leased its rights to oil and gas firms. After 1982, the Gulf has operated under a system of "area-wide" lease sales – usually two a year. Companies submit sealed bids for blocks that are typically three miles square. In shallow water, winners receive the right to explore for five years; in deep water, the primary term is ordinarily twice that. Although bidders may formally submit joint bids for a single block, they are prohibited from otherwise cooperating to choose blocks for bidding or in setting bid amounts.

Since 1983, companies have exposed $27 billion in deep water bids, winners have paid a total of $19 billion for leases and of that amount, left $7 billion, or over a third, on the table.[1] The industry bought a total of 12,278 leases covering 6,834 blocks with its investment. Of these blocks, 1,044 have been drilled and presently 1,608 are covered by current leases. As of mid-2018, about a quarter of blocks covered by current leases had produced oil and gas.

To win at minimum cost, companies devote significant attention to evaluate both the likelihood that other companies will compete for blocks of interest to them and how much those competitors might offer. As a rule, the lease goes to the highest bidder. Therefore, if a company judges the likelihood of a competing bid on a block as very small, it can submit a low offer with little fear of being out-bid. When a company *does* expect a contest, the best economic outcome is to win the block with a bid that is only $1 greater than the next highest offer.

Risked economic returns to prospects on a block impel each company's decisions to bid and how much to expose. Highly proprietary geoscience, engineering and economic analysis constitute most of the pre-bid evaluations. The public never sees this. However, decades of bidding, wining, losing, exploring and

---

[1] All dollar values used in this article are constant (or "real") 2017 dollars, using the gross domestic product deflator reported by the US Federal Reserve Bank of St. Louis, accessed on July 6, 2018 and updated on September 10, 2018 at https://research.stlouisfed.org/fred2.

producing sketch outlines of companies' strategies, reveal the characteristics of blocks that do or do not attract bids and how much money firms think it takes to win.

In this research, we analyzed all 24,392 deep water bids submitted since 1983 with two goals: 1. to forecast the probability that any block offered by the government will attract at least one bid in a future sale and, 2. to predict the amount of the high (winning) bid for each open block. We focused and extensively tested our models on the 3,867 deep water bids from 2009 to 2018 – those most relevant for predicting upcoming sales. Modeling the bid/no-bid decision was conducted by company and the bid-amount decisions were estimated for all companies collectively.

The Bid/No-bid Decision

Before each lease sale, the federal agency that manages the Gulf of Mexico (GOM) publishes a list of blocks available. It's usually all blocks not leased at the time of the sale. However, we do not include blocks in such deep water or so distant to be uninteresting over the next decade. For companies, the first decision is which blocks warrant study? Then, given a budget, firms rank alternatives and pick a subset to bid. We tested about three dozen variables in search of those that could explain companies' bid/no-bid decisions (Table 1).

The Gulf-wide variables are designed to capture the general drivers of bidding decisions in the GOM (e.g., oil price) and local attributes of open blocks are its characteristics that influence the probability of bidding (e.g., water depth, does the block contain or is near a past lease, a discovery or wells?). Finally, we recompute the local variables for *each company* analyzed (Table 2). For instance, when computing the Chevron Bid/No-bid model, for each block, we would compute not just if it was close to a past lease – but a past *Chevron* lease.

From a modeling perspective, the probability that a company will bid on a specific block emerges from the interaction of each block's suite of attributes and the positive and negative preferences of each individual company toward each of those attributes, as revealed by their past bidding decisions.

Table 1
Predictive Variables Used in the Bid/No-bid Models

Gulf-wide Variables	Variables Locally Computed for Each Block (but not company-specific)		Variables Locally Computed for Each Block (separately for each top company)	
Oil price	Fields	Newly Available Blocks	Past Leases	Blocks not Currently Leased
Change in oil price	Blocks not Currently Leased	Discoveries	Current Leases with Fields	Relinquished Leases
Region of GOM	Wells	Relinquished Drilled Leases	Past Bids	Relinquished Drilled Leases
	Platforms	Relinquished Leases	Current Leases	
	Current Leases	Past Bids		
	Past Leases	Water depth		

For practical reasons, we limited the company-wise analysis to the 20 top players over the last decade (Table 2). We also created a "composite" 21st company, aggregating the next cohort of 25 lower-activity companies below the top 20 bidders. These company lists are revised annually based on activity (e.g., Noble, Stone and Cobalt are not included in the 2019 analysis).

Table 2
High- and Medium-Activity Bidders in Deep Water GOM, 2009-2018

Company	Money Exposed (Million $2017)	Number of Bids	Company	Money Exposed (Million $2017)	Number of Bids
High-Activity Companies			**Medium-Activity Companies (cont)**		
Shell	$1,419	285	Ecopetrol	$212	121
BP	$777	336	Stone	$195	87
Statoil/Equinor	$895	189	Cobalt	$179	108
Chevron	$889	284	Venari	$175	62
ExxonMobil	$914	239	Red Willow	$170	79
BHP	$618	223	Total	$158	120
Anadarko	$443	185	Murphy	$156	64
Medium-Activity Companies			Repsol	$138	74
LLOG	$306	115	Ridgewood	$124	67
Noble	$237	84	Houston Energy	$59	89
Hess	$265	69			

The best way to illustrate how companies evaluate the block attribute variables is by tornado plots. Using ExxonMobil and LLOG in 2014 as examples (Figures 1a and 1b), block attributes from Table 1 are listed along the vertical axes, the most influential variables are at the top. The bars move left of zero showing a company's negative preferences toward an attribute in bidding and to the right for a positive preference. The longer the bar in either direction – the greater the influence.

Note in Figures 1a (ExxonMobil) and 1b (LLOG) that not all explanatory variables appear for both companies. Different variables are significant for different companies, as they focus on different factors, which can change from year to year. As well, in 2014, the set of variables that influenced ExxonMobil's bid/no-bid decisions was a larger set than for LLOG.

The statistical analyses in the project were run in *R* and all data was drawn from Earth Science Associates' *GOM³*, sourced by data submitted by bidding companies to the US government. The model can be run on past years for validation and insight and it can be used to predict the next sale. In either event, a training period of at least three years prior to the forecast year must be defined by the user. To forecast the next sale, the user must also predict oil price. After the model executes, output including reports, tables, maps and map files for loading into *GOM³*, or a geographic information system (GIS), are returned.

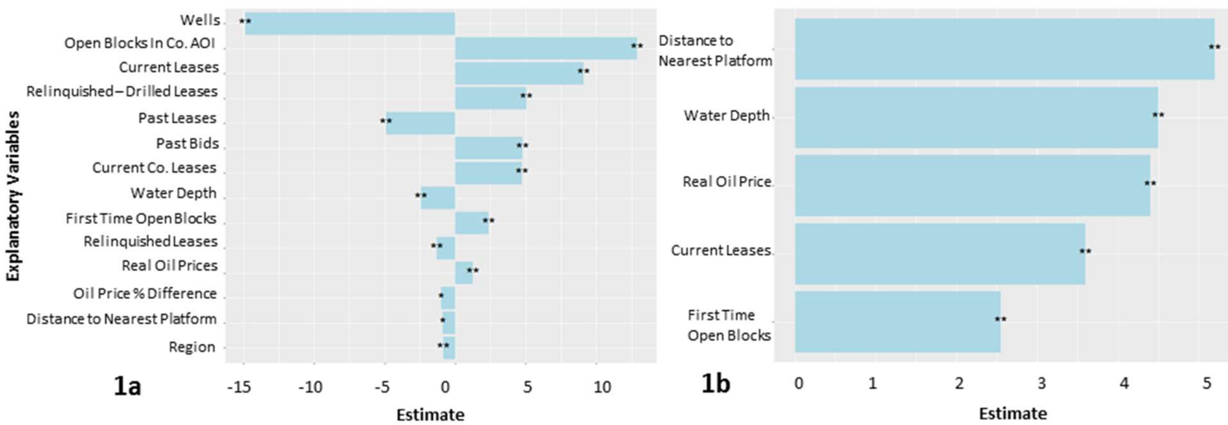

Figures 1a and 1b show tornado plots from the bid/no-bid model for ExxonMobil (1a) and LLOG (1b) for the 2014 sales (trained on 2009-2013). The independent variables in the analysis are listed on the left of each chart. Bars to the left of zero reflect a negative effect of a variable on the probability that the company will bid on a block possessing that attribute. Positive bars indicate the reverse. Variables that are statistically significant at a 10% level have one asterisk and those that are significant at a 5% level have two asterisks.

Bid/No-bid Model Performance

How well did the bid/no-bid model work? Clearly, with no access to bidder's proprietary data, analysis or strategies, any outside forecasts will inherently be imperfect. The purpose of our research is to establish consistent guidance based on a transparent, objective methodology and public data. It provides neutral benchmarks – a decision aid and not a decision maker.

Using the 2014 sales as an example, we ran the model with an oil price of $101/bbl and used a training period of 2009-2013. We applied each of the 21 company models separately to estimate the probability of each company biding on each block. This led to five classifications of blocks: the best ones, Group A in Figure 2, which we assess as having a 50-100% of getting at least one bid. Group B included the next best – those with a 25-50% chance of being bid and on down by half until reaching the lowest blocks with < 5% chance of being bid, Group E.

Our forecasts of the likelihood of each open block being bid has two timeframes. In Figure 2 the short run target is forecasting the next sale, here, 2014. Of the 3,725 open blocks modeled, 115 were classified highest, as Group A; of those, 50 were actually bid in 2014. While getting almost half right in the short run was a very good result, the process is not really over once the immediate next-sale result is complete.

Many of the best blocks in 2014 that remained unbid were left because bidders' budgets were not big enough to immediately pick up more of the Group A blocks. Sales over the following three years demonstrated this. Between 2015 and 2017, 15 more Group A blocks from the 2014 forecast were also bid, making a long run total of 65 bids of the 115 Group A blocks, or 57% accuracy. This is within the 50-100% target for the Group A class and is a major success. In Figure 2, the short run success for each Group is shown in solid colors and above, in the hatched portion of the same bar, the long run success is shown. The total accuracy is printed above.

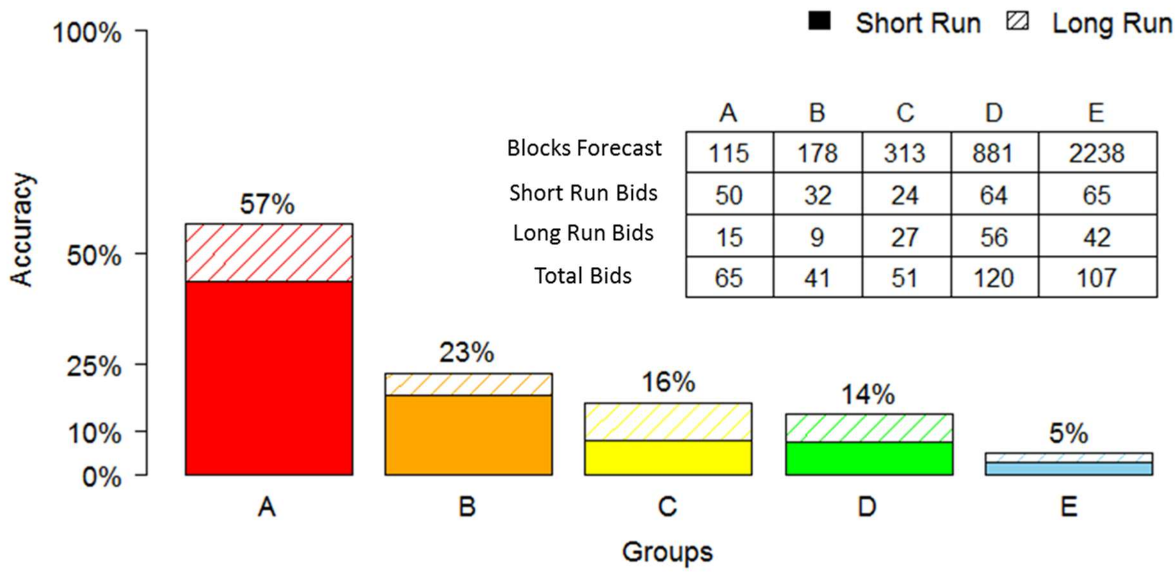

	A	B	C	D	E
Blocks Forecast	115	178	313	881	2238
Short Run Bids	50	32	24	64	65
Long Run Bids	15	9	27	56	42
Total Bids	65	41	51	120	107

Figure 2 shows the results of the bid/no-bid model for a test case run for 2014, based on a training period from 2009-2013 and an oil price of $101/bbl. The model grouped open blocks into five classes, based on the likelihood of receiving at least one bid. Group A are the best blocks, forecast to be bid from 100% to 50% of the time; Group B blocks have a 50% to 25% chance of a bid; Group C from 25% to 12%; Group D from 12% to 5 and lowest, Group E, less than 5%. The solid colors in each column reflect out of the number blocks in that class, how many were actually bid in the next sale (2014, abbreviated as "SR Bids"). The hatched area above indicates how many *additional* blocks in that Group were actually bid over the following three years (e.g., 2015-2017, or "LR Bids").

Keeping in mind that we have no proprietary information, the forecast performance of the 2014 bid/no-bid model was a victory. Moreover, we think that long-run growth in accuracy would have been *much* higher (especially for the best blocks) had 2015 and 2016 not been the bottom of the oil price crash.

More generally, Figure 2 speaks to three basic measures of model performance:

1. How much better is our model at identifying blocks that will be bid than random selection? For the 2014 sales, about 6% of the blocks in the study area received bids and we classified 115 blocks as Group A – the most likely to be bid. Under a random process we would have expected 7 (6%) of these blocks to receive bids. However, 50 of the Group A blocks were picked up by companies (in the short run) which is over 7 times higher than would have been expected without our model.

2. Does the model correctly classify blocks into the five "quality" groups (A through E), as reflected by the rate at which they are actually bid? For over 150 years, explorationists have recognized that there are exponentially fewer *good* places to look for oil than *bad* ones. Figure 2 reflects that

quality continuum and exponentially declines from the 115 (2.2%) classed as top blocks (Group A) down to the 2,238 (42%) assessed as having only a <5% chance of being bid (Group E).

3. Is improvement of model accuracy in the long run proportional to forecasted block quality? If the model is assessing fundamental differences in prospectivity, the bidding rate of blocks assigned to the top groups will grow over time. This is what happened; whereas, the growth in number of blocks bid in Group E, the bottom of the barrel, hardly changed between the short and long runs.

Stepping back from performance for a single forecast period, how did the bid/no-bid model perform over time? Given our focus on 2009 – 2018 and requiring a minimum three-year training period, the results are shown in Figure 3.

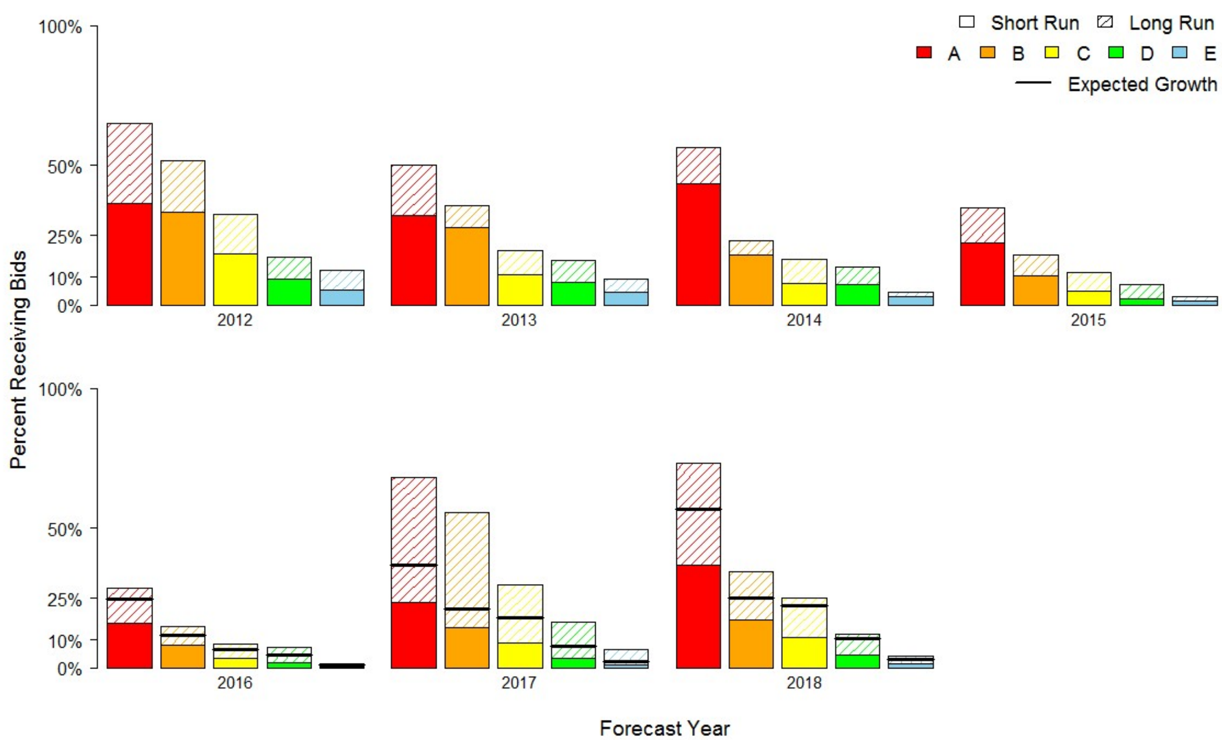

Figure 3 shows the performance of the bid/no-bid model for 2012 through 2018. The grouping of blocks by probability of receiving at least one bid is explained in Figure 2. For 2012-2015, we have all three years of long run results after each forecast year. For 2016-2018, all three of the years of the long run have not yet occurred. So, we estimated the amount by which long-run forecast accuracy is expected to increase based on 2012-2015 performance and represented it by thick horizontal black lines.

Generally, the models performed very well in 2012 through 2014. Then, however, price crashed and the principles that worked before 2015 fell apart. Model performance in 2015 was significantly poorer and in 2016, the bottom of the price collapse, overall model accuracy fell by about 50% compared to 2012-2014. In 2017 and 2018, price recovered, the industry slowly returned to the bidding behavior of 2012-2014 and the models in both years significantly outperformed their *pre-crash* successes.

The Bid Amount Decision

Once a decision is made to submit a bid on a block, the question is how much to offer. Bidding too little risks loss to a competitor (or being rejected by the government as insufficient – see the next article in this series); bidding too much leaves money on the table. Like Goldilocks, the magic is in finding the bid amount that is "just right."

To forecast bid amounts, we employed the same general type of model and used roughly the same set of explanatory variables as in the bid/no-bid modelling. Global variables, like oil price, play a key role (in the first instance, through budgets); water depth remains important and so do block-specific variables.

The bid amount model runs collectively for all companies, it also produces a tornado plot showing which variables were important to the *industry's* decisions on bid amounts (Figure 4). In the spring 2018 sale, the largest positive influence on bid amount was the block being open for the first time and being close to both current and past leases. The paradoxical result is the negative relationship between prevailing oil prices and average bid amount. A reasonable interpretation is that, even by the spring of 2018, the recovery of prices from the crash was not high or trustworthy enough to fire the intensity of competition that inflates bid amounts across the board, as it did in 2012-2014. In fact, the spring 2019 sale saw remarkably low bid amounts, in spite of price expectations in the $60s.

Bid Amount Model Performance

The distribution of bid amounts is, and always has been, highly skewed to the rich end. Most bid amounts are close to the median for the sale, which is usually around $900,000**.** The low half of the distribution is between about $600,000 and $900,000. The wild end, however, runs to over $100 million for a single block.

Figure 5 gives a broad view of bid amount forecasts and actual bids from 2012 through 2018. Two conclusions arise: First is that the central parts of the distribution of actual bids (represented by the medians and interquartile boxes) and the model forecasts are very close and improved over time. Second, the high outliers in the distribution of actual bid amounts are *consistently* underrepresented by the results of the bid amount model. While we forecast bid amounts into the millions, our approach cannot forecast bids of tens of millions of dollars for a block, never mind predicting exactly which lucky blocks will attract that kind of money.

The redemption of this limitation on predicting outliers is that, by definition, they happen very rarely. From 2012 to 2018, 85% of bid amounts were under $4 million, carrying errors in the forecasted bid amount that are 50%, or less, of the actual bid amount. As actual bid amounts climb into the tens of millions of dollars, to stratospheric numbers above $100 million, the forecast errors go astronomical. However, the number of blocks attracting that money is a *tiny* fraction of the total population of bids.

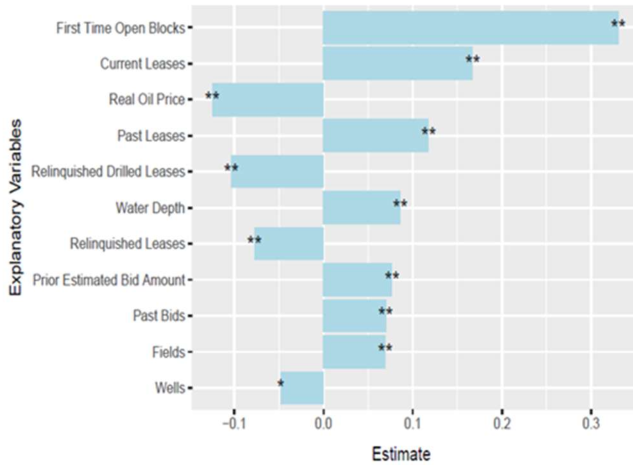

Figure 4 shows the tornado diagram from the model estimating bid amounts for all relevant open blocks analyzed in advance of the spring 2018 sale. Interpretation of this plot, for all companies collectively, is the same as for Figure 1.

Conclusions

The extreme complexity of interpreting geoscience and engineering analyses of blocks, then fitting them into a company's strategy and portfolio, is hard enough. So, to help provide an objective and transparent context of the competition bidders face, we developed two mathematical models to forecast the likelihood of bidding and bid amount by block.

In assessing their performance, we compared forecast results to actual outcomes over the last seven years of a *highly* volatile oil market and tremendous swings in investment. Most of our forecasts hit their theoretical target ranges, though, especially for the good and best blocks – usually at the low end of those ranges (e.g., 57% long-run accuracy for the Group A forecast in 2014). An accuracy of 95% would, of course, be much nicer. However, that won't happen, even if proprietary data and analysis were magically available. There is an inherent, inescapable variability in both bid and bid amount decisions. Nevertheless, it is encouraging that the models did better post-crash than the credible job they did between 2012 and 2014.

We expect to further improve accuracy by rolling in open-source geoscience, engineering and economic data. While that will bring a jump in how often we're right, a lot will remain unknown – especially when the bid amount is $400,000 for one block and $40 million for the block next to it.

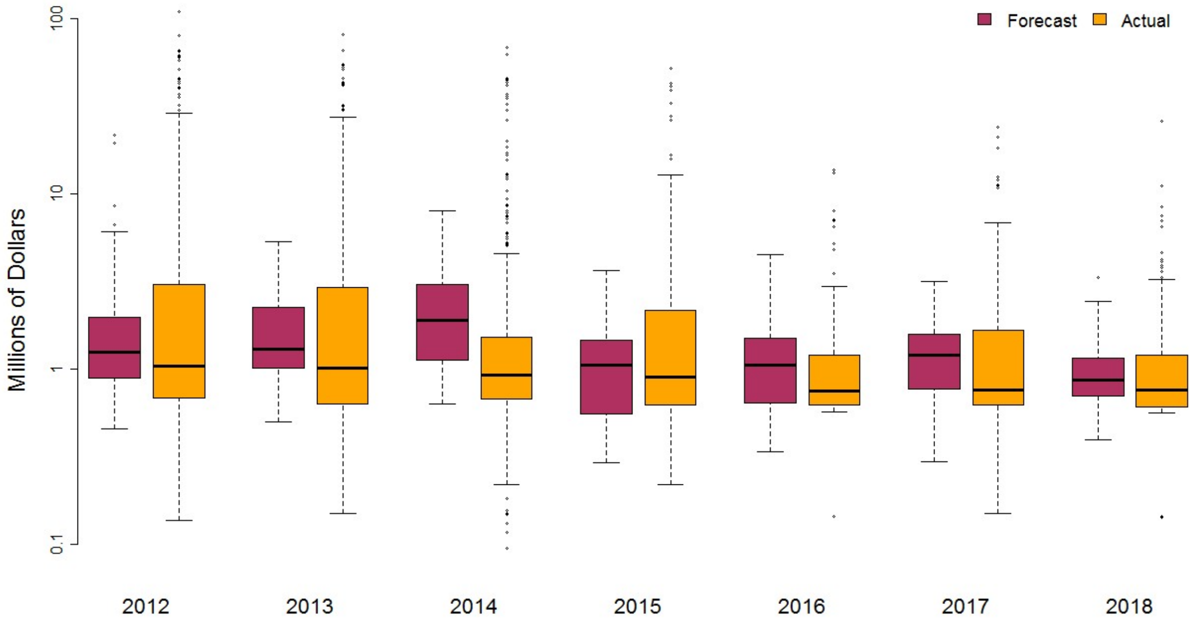

Figure 5 compares the distributions, in box-and-whisker plots, of the forecasted (purple) and actual bid amounts (orange) on blocks receiving bids each year. The thick black lines in each box are the medians of each distribution; the bottoms of the boxes are set at the 25th-percentiles of the distributions and top are the 75th-percentiles. If there are outliers, the "whiskers" extend above and below the boxes by 1.5 times the inter-quartile range. If there are no outliers, they extend to the lowest and highest bids. The dots represent outliers. The vertical axis is logarithmic.

Because we output the results in GIS-ready form, companies can leverage the value of our results. By layering in their proprietary geoscience, engineering and scouting information, they can customize our results with the wealth of proprietary information we cannot incorporate. In that context, quantitative analyses such as ours are ultimately principally a tool in decision makers' box, helping to inform their intuition.