The Lessons of Bid Rejection in GOM Deep Water Lease Sales

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

In the last 10 years, of all open blocks bid in deep water, the government has rejected only 4% of high bids for being too low. That's not much – unless it happens to you.

The federal government awards leases through a largely opaque process. Before each sale, it publishes a list of open blocks (typically thousands) and a minimum legal bid (MLB) for each. These minima do not reflect any analysis: in deep water, $25/acre was assigned for all water depths through 1999 and thereafter for some blocks in water shallower than 1,500 feet; from 1999-2010, for all but the shallowest deep-water blocks, the minimum went to $37.50/acre and, since 2011, it was set at $100/acre.[1]

After the sale, all bids exceeding the MLB follow one of two paths. If the government itself sees *no* value on the block (i.e., no "discernible prospect" nor "economic quantities of risked resources") they judge it "unviable" and promptly award it to the highest bidder. What's to lose?

Yet the majority of blocks bid are judged "viable". They are, therefore, "subject to a full-scale resource and economic evaluation to determine if each tract's highest qualified bid is representative of fair market value." This process yields the lowest high bid for a block that the feds consider "representative of fair market value [FMV]", a measure they call "Mean Range of Values" (MROV). If the high bid exceeds it – the government awards the lease.

But here's where it gets murky. The standard definition of FMV is the price for which the property would sell on an open market. However, a state monopolist selling leases in a sealed first-bid auction to private competitive bidders *is in no sense* an "open market". So, absent an open market reference price, the government takes a different approach. They set the MROV as high as possible, while still allowing the bidders to make a "normal return" on their investment, after paying for the lease, exploration, production, transportation, taxes and royalties. If the lessee ends up with $1 more than the fed's estimate of "normal" return – they feel they should have charged $1 higher for the lease.

Computing normal returns for dozens of companies submitting bids on a total of up to hundreds of open blocks in a sale would be impossible, so the feds don't even try. For more than 80% of all offered deep-water blocks, when it comes time to compute MROV, the government simply enters the arbitrary MLB assigned to blocks before the sale. As a typical example, in the March 2018 sale, of 103 deep water blocks bid, 87 were given MROVs of $576,000, or $100/acre times the area of a standard block, 5,760 acres.

---

[1] In the first paper in this series, all money values were shown in real dollars ($2017). However, government practices are persistently in *current* dollars (e.g., using $25/acre as a minimum bid for decades). Therefore, in this second paper, unless otherwise noted (e.g., Figure 3), all money is in current dollars.

Lesson 1: <u>As long as government practice remains constant</u>, unless your block covers or offsets demonstrated actual or possible production, there is a >80% empirical chance that if your bid is high and more than minimum legal bid, you'll win.

The government actually does analyze about 20% of deep-water blocks on offer (i.e., it sets MROV > MLB). The most important pattern to emerge from this analysis is that they systematically assign higher MROVs for blocks that are proximate to demonstrated production or productive capacity.[2] Figure 1 documents the relationship between location and the amount of a high bid versus the probability it would fall *below* the MROV on the block and therefore would likely be rejected.

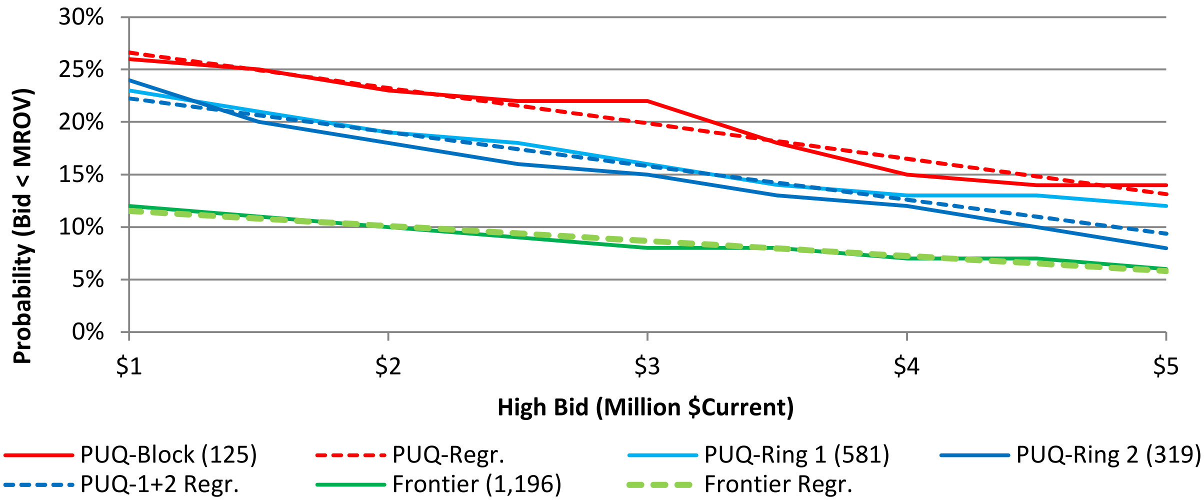

Figure 1 shows the relationship between the amount of a high bid on a block, the relative location of the block and the probability of a specific high bid being less than the MROV for the block. The block groupings, PUQ-blocks, PUQ-Rings 1 and 2 and Frontier are defined in footnote 2. The dotted lines represent linear regression estimates for the PUQ blocks, the blocks of PUQ-Rings 1 + 2 collectively and Frontier blocks. The number of observations in each group are in parentheses. To interpret, a $2 million bid on a PUQ-block would have a 23% chance of being below MROV and the same bid on a Frontier block would only have a 10% chance of being below MROV.

In pursuit of their policy goal to leave lessees no money beyond their calculation of "normal return", the government recognizes that blocks on or near established production are generally more valuable than the Gulf-wide average. Usually, they enjoy lower geologic risk, raising their risked net present value.

[2] Here, a "productive" block has, *or has had,* at least one of the following attributes: 1) a first production date (the "P" of the PUQ abbreviation); 2) it belonged to a production unit (the "U") or 3) it had a qualifying date, meaning logs or well tests once demonstrated commercial hydrocarbons (the "Q"). We also define two sets of neighboring blocks: PUQ-Ring 1 blocks are those eight that immediately ring PUQ Blocks and PUQ-Ring 2 are the set of 16 blocks that immediately ring PUQ-Ring 1 blocks. Blocks outside the union of the three PUQ sets are classified as Frontier.

Moreover, they are often closer to infrastructure, dropping costs for transportation and perhaps development. To capture the fed's "fair" share of a larger value – MROVs must rise.

Though >80% of blocks have MROV=MLB, what is far more complex is the distribution of MROVs *above* the current $576,000 MLB. A model for forecasting the MROVs of open blocks in the next sale is addressed below. However, more generally, the regional trends in MROVs from 2009 – 2018 for both PUQs and Frontier blocks are mapped in Figure 2.

In interpreting Figure 2, note that the grey blocks, representing the most recent MROV=MLB, correspond exactly to block boundaries. However, the colored areas are gridded values designed to show regional trends in the most recent MROVs assigned between 2009 and 2018. If government practice persists, the most likely MROV on a Frontier block is $576,000. Where the block MROV is above the MLB and, in Figure 2, it falls in an area colored dark blue, the local trend in MROVs has been assignments of values between the MLB and $4 million.

In the end, however, there is a crucial component of variance in Frontier MROVs that simply cannot be systematically explained with just public data – the extreme MROVs. In that sense, the distribution of MROVs reflects the extreme right tail of the distribution of companies' bids themselves. As shown in the first article in this series – it is not that uncommon to see one block in a lease sale attract a bid of $700,000 and the immediately adjoining block pull in $18 million. While the rich (right) end of the distribution of MROVs is not *nearly* as extreme as that of bids, there are some very lavish MROVs.

Lesson 2: Location matters. The closer an open block is to a block with demonstrated potential, the higher the MROV. In Frontier areas, MROVs have a highly bimodal pattern: Overwhelmingly, MROVs = MLB but about 20% are assigned up to 70 times higher. Beyond the math, however, if you see great prospects in your data or read in a trade publication that a specific area is hot, those setting the MROVs saw the same articles – and they've also seen your data.

We note what appears to be an interaction between the amounts of rejected bids and MROVs over time. Since 2009, there have been 59 blocks that received additional bids after being rejected. When rebid, nearly all of rejected blocks are accepted in the next sale. However, 12 of these blocks had multiple rejections until finally being accepted (the extreme was five rejections).

Between 2009 and 2018, the government has made $250 million more by rejecting bids since the subsequently accepted bids are typically higher. Only 17% of the rejected-then-accepted bids were *less* than the original bid. The balance is usually more than double the original high bid. Even so, MROVs on rejected blocks tend to float down over time, but overwhelmingly, bidders raise their offers by more than MROVs fall.

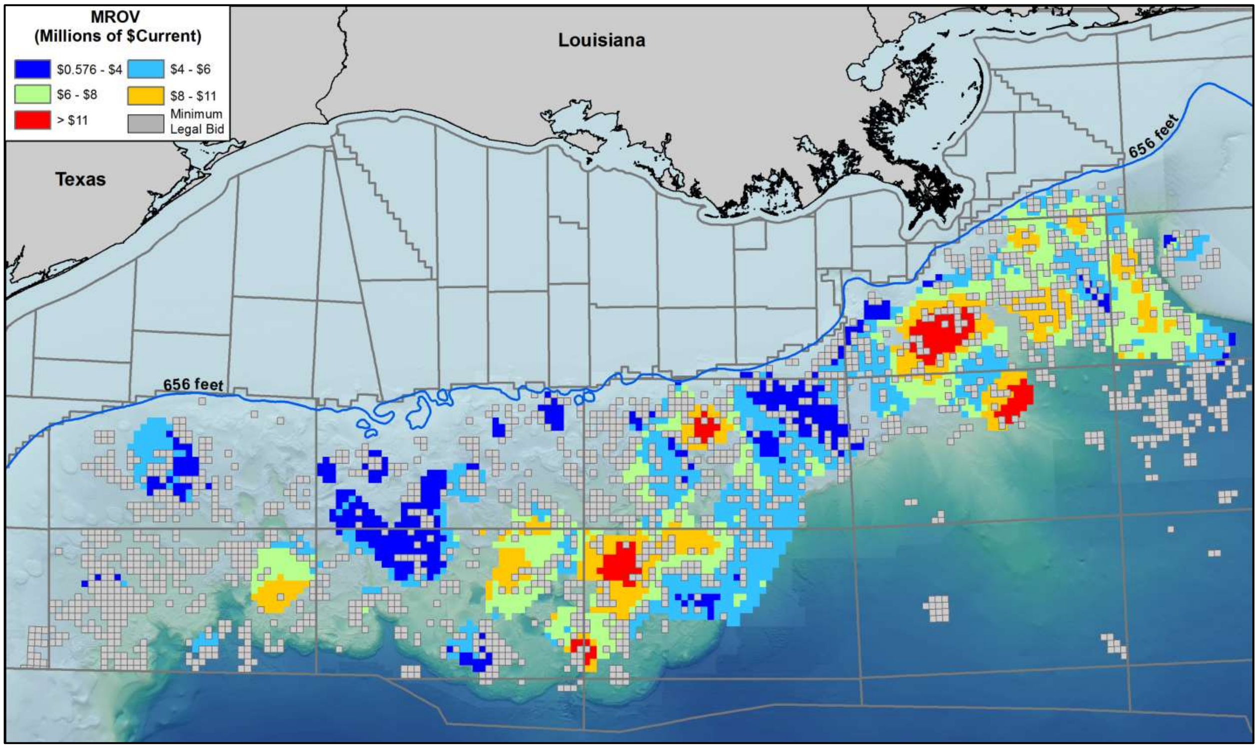

Figure 2 maps the MROVs in deep water from 2009-2018. The colored areas are locally gridded values of MROVs above the MLB for the block and the grey blocks were assigned MROVs equal to the MLB. The areas not covered by either grey blocks or the colored grid have not been assigned MROVs between 2009 and 2018. The geographic limit of deep water in this study is shown by the 656-foot isobath (blue line).

Lesson 3: If your bid was rejected, you really want the block and a higher bid is economic – shoot to meet the government somewhere in the middle.

The price forecasts used by the government when deciding MROV are never revealed – neither are the discount rates, cost escalators or any other parameters that might be usefully employed in auditing the process. However, it appears that the government's price forecasts are not nearly as variable as either the spot price of oil or the planning prices companies use for investments.

This may be well founded. Almost all deep-water leases run 10 years. So, it makes sense for the government to orient to a decade timeframe, rather than just reading the WTI spot price in a newspaper on the morning of the sale. In fact, if we compare the price of oil in any month from 1974 through 2008 to the oil price that prevailed for the decade following that month, a pretty consistent relationship emerges (Figure 3).

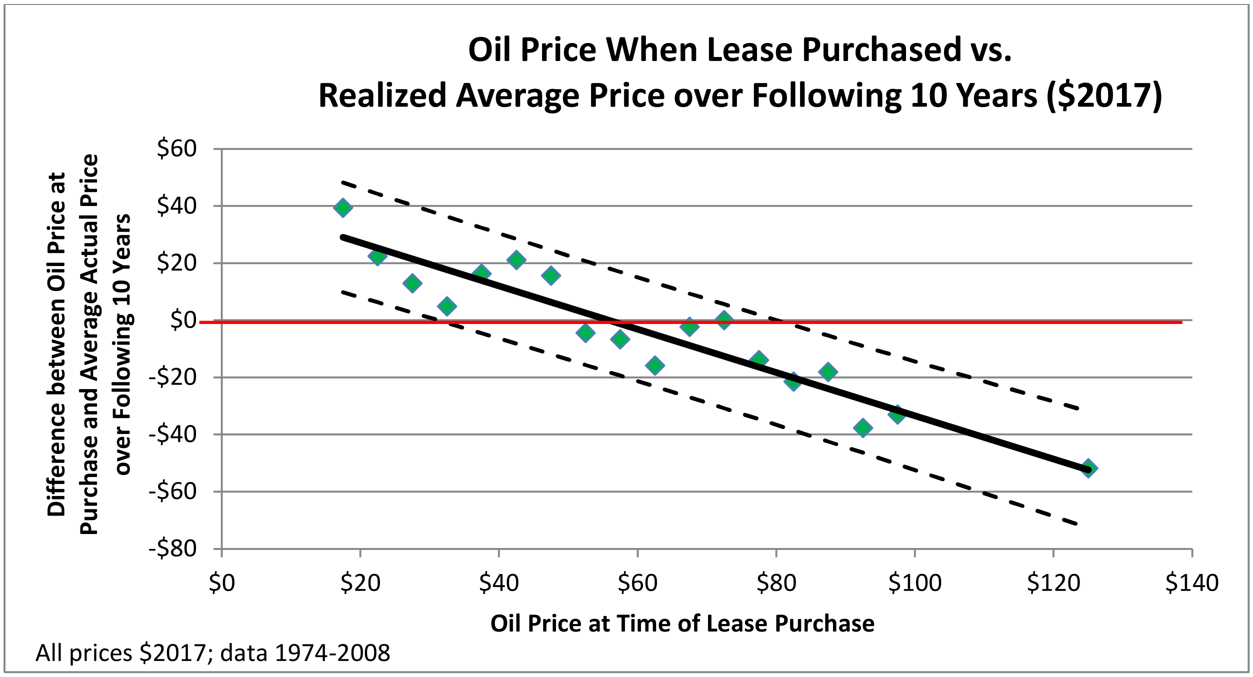

Figure 3 shows the regression-estimated relationship (black line) between the price of oil in any month between 1974 and 2008 (e.g., when a lease is purchased) and the price that prevailed for 10 years thereafter (e.g., over a lease term). The 10%/90% prediction intervals for the regression are shown in dashed lines. Because of the long period covered, note that these data are expressed in constant 2017 dollars.

As an example, take an oil price of $100/bbl. From $100 on the x-axis, read up to the regression line in Figure 3 and left to the y-axis, which equals about -$33. This estimates the long-run (10-year) discount to the $100/bbl price. Therefore, the expected average price for the 10 years after a $100 month is $67/bbls ($100 - $33). The 10%/90% confidence intervals run from approximately $48 to $86/bbl. If that relationship holds, slow changes in price forecasts may make sense.

Lesson 4: Although it can be only indirectly inferred, it appears that, over time, the government uses a fairly stable price forecast when computing MROVs on blocks.

If you are high bidder on a block but shy of the MROV, rejection is not automatic. The regs require a comparison of the high bid to the government's "backstop" reservation price for the block, called ADV (Adjusted Delay Value).[3] If ADV < MROV there are two official rationales:

1. The government is accounting for the time value of money. They can calculate the present value cost of deferring revenues by rejecting in this sale and awarding the block in the next sale. If the MROV-ADV difference actually accounted for the time value of money over a one-year deferral, a reasonable outside bound might be an ADV 10% less than MROV. Yet the actual differences

[3] This is also referred to as DMROV (Delayed MROV).

between the MROV and ADV run from zero to 92% of the MROV and average 36%, implying astronomical discount rates that are just not credible for economic analysis.

2. Theoretically, if the open block adjoins a producing lease, the neighboring lessee could drain some of its hydrocarbons before the next lease sale, devaluing the block. This is probably extremely rare in practice and would be extremely difficult to show.

As both official reasons are problematic, ADV could very well be a fudge factor. Perhaps when the government sees the bids after the sale, they recognize that some of their MROVs were really unjustifiably high.

Lesson 5: A high bid less than MROV is not the final word. However, there is scant publicly available guidance on what data and assumptions go into MROV calculation, even less on ADV or the appeals process afforded to rejected high bidders.

Finally, in trying to divine the government's rejection logic, a large, almost philosophical contradiction becomes clear. The government's express legal objective in setting MROV is to define a minimum payment for a lease below which taxpayers would not get fair market value – the price that would prevail in an open market. Yet that same economic theory also says that MROV, as well as being the *minimum* acceptable price for the government, as the seller, is also and precisely the *same price* that is the *maximum* price which a bidder, as buyer, could pay and earn a "normal return".

This means every time a company pays the government more than the MROV, it is at the cost earning less than a normal return and the government is, in the same transaction at that price, receiving more than fair market value for the lease. For just the 13 leases, out of over 12,000 leases between 1983 and 2018, where the high bid *exactly matched* the MROV, did both sides get "fair market value".

Integrated Forecasts of MROV

These are good rules of thumb but what we could really use is a pre-sale forecast MROV on each open block in the next sale. The government assigning 80% of MROVs to MLB, and the fact that the MROVs of the 20% of blocks with MROV>MLB are exponentially distributed, make this particularly difficult. These problems call for a different forecasting approach.

We employed a decision tree model. Historically, decision trees have been used for analyzing decisions like drilling a prospect. First, do you acquire a lease or not; then, if acquired, to drill or not; then, it is a dry hole or discovery and so on. All paths end in a terminal node describing a specific outcome. As costs, revenues and any probabilities are integrated, they propagate down the path to the terminal nodes, giving the final economics and risks of that particular path through the tree. Collectively, the tree's terminal nodes span the array of outcomes and chart the choices made at each step to reach them.

We employ the decision tree model predictively. For this, we have two training sets: 1. the MROVs assigned for the blocks in previous sales and 2. the values of block attributes at each sale that may have influenced MROV assignment. Traditionally in decision trees, attributes are elicited as answers to questions: Was the block bid in the past? Had it been previously rejected? Had it previously been leased? Did it have previous production? How much was the last MROV assigned? Does it adjoin a current lease? The terminal nodes contain similar MROV outcomes and collectively, they cover the range of historical MROV assignments.

As an example, Figure 4 shows part of a decision tree model of the probability that a block would receive an MROV = MLB. In "growing" this tree, the block characteristic with the greatest impact (at the top of the tree) was: had this block been rejected when bid last time? For a "yes" answer, the model estimated a 0% chance it will be assigned an MROV = MLB in the next sale (top terminal node in Figure 4).

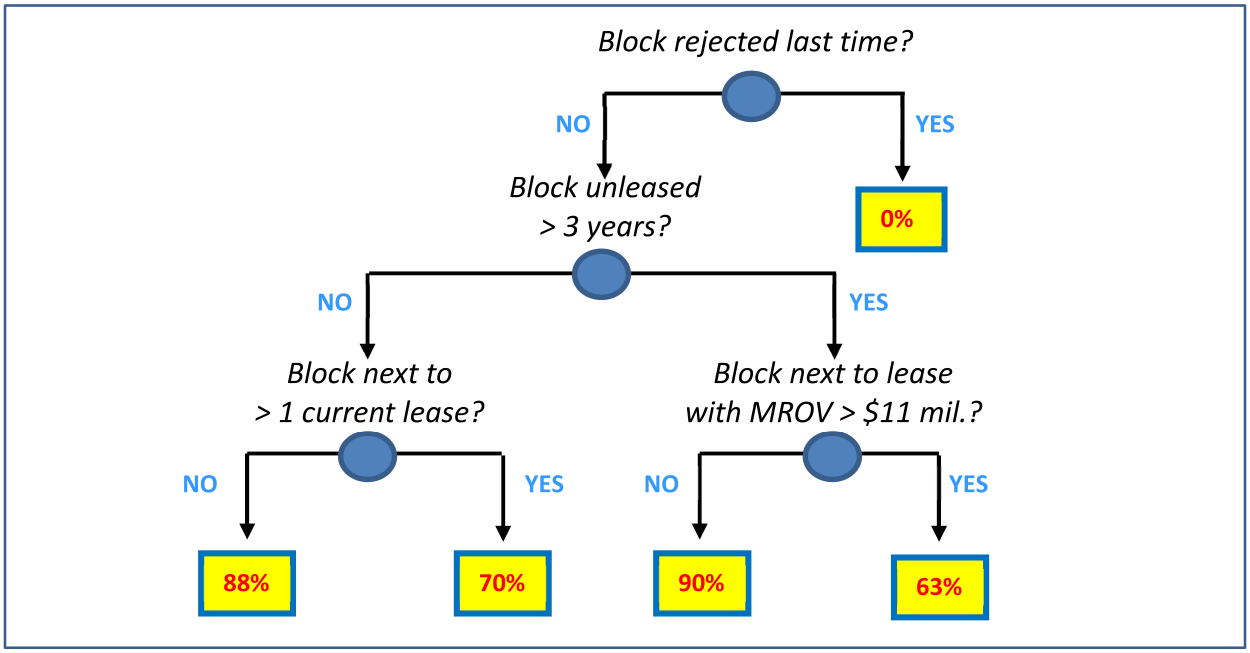

Figure 4 shows part of a decision tree model for estimating the probability that the government would assign a minimum bid amount (MROV) equal to the minimum legal bid (MLB, usually $576,000) to a specific open block in the next sale. The yellow boxes are terminal nodes.

If the block was *not* rejected last time, the "no" path branches to another question: how long had the block been unleased? The value that best partitions similar values of block attributes and MROVs was 3 years unleased. Following each path through optimally posed questions leads to the remaining terminal nodes. To better capture the uncertainty in the analysis, we employed a methodology called "random forests". As the name implies, forests are collections of decision trees.

We predicted the MROVs with two random forest models. First was to estimate the probability that the government would assign MROV = MLB. In Figure 5a, a map of five blocks is shown, where the labels give our estimate of the probability that MROV = MLB. The four green blocks were in fact assigned MROV=MLB. The orange block, with the lowest probability of MROV=MLB was actually assigned an MROV > MLB.

Second was to directly estimate the MROV on each block in the next sale. For the same blocks in Figure 5a, Figure 5b shows our predicted MROV ("Pred") and the MROVs actually assigned by the government

("Act"). For the three out of four blocks that are green in Figure 5a, the predicted and actual MROVs were very close – the errors for three were between $144,000 and $264,000. The error on the northern most green block was higher.

For the orange block, which the government assigned MROV>MLB, the predicted MROV was much higher than the rest. Although the actual MROV was over twice our estimate, our prediction of a MROV = $2.1 million was among the higher predictions for that sale. This is a red flag to bidders that this block would attract a high MROV (and relatively unlikely to be equal to the MLB).

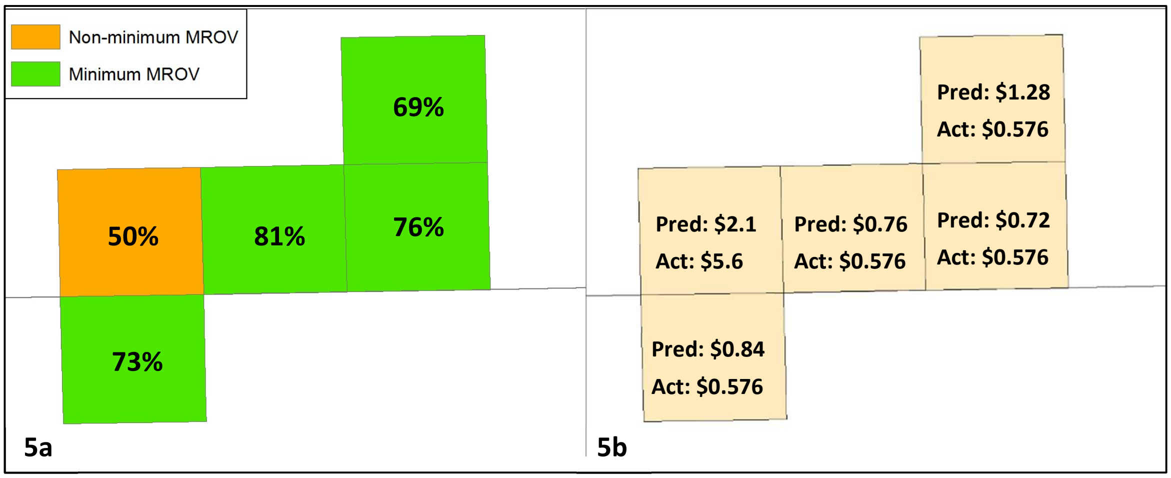

Figure 5 maps the random forest forecasts for the government's principal lease reservation price, MROV, below which bids are generally rejected. The block labels in Map 5a show the forecasted probabilities that the government would assign an MROV=MLB. The green blocks were actually assigned MROV=MLB and the orange block got an MROV > MLB. Figure 5b shows the same five blocks in which the labels compare our estimated MROV (Pred) and actual MROVs (Act) in millions of dollars.

Conclusions

The government's bid rejection process is turbid and complex. Nevertheless, systematic analysis of thousands of bids to better understand why some were rejected provides guidance to bidders to avoid that in the future. It is not a perfect system but a set of predictions based on an objective and transparent methodology based on public data. These are benchmarks.

In this, and the first article in this series, we have tried to clarify the government's opaque processes of bid evaluation and rejection by applying neutral, quantitative analyses. This yielded insights and models that contribute to the tool sets companies have to understand competition on open blocks, what other bidders might offer for them and how the government weighs bids for acceptance. The fuller the foundation of empirical data and analysis, the more efficient the process.