Crude Oil & Gasoline: Supply and Demand

Testing simple ways to forecast how much crude oil and gasoline will be in storage.

Data through: June 2026
Model results: September 16, 2026
Write-up: September 17, 2026

1. Introduction

I wanted to see whether simple rules could help forecast U.S. crude oil and gasoline inventories. Inventories / stocks, are the amounts held in storage.

The two markets are connected by refineries. A refinery takes in crude oil and turns it into gasoline, diesel, jet fuel and other products.

What supply and demand mean here

For crude oil, supply includes oil produced in the United States and oil imported from other countries. Demand mainly means the crude that refineries take in to process. The amount of each product made from crude varies, so extra crude going into a refinery does not all become gasoline.

For gasoline, supply comes from refineries, blending and imports. Blending means mixing ingredients to make fuel. Gasoline demand is measured by product supplied. This is the U.S. Energy Information Administration’s (EIA’s) estimate of how much gasoline leaves the main supply system for domestic use. It is an estimate, rather than a direct count of fuel sold at filling stations.

Stocks rise when more comes in than goes out. This is called a build. They fall when more goes out than comes in, which is called a draw. Trade, regional movements and use outside refineries also affect the balance, so crude and gasoline stocks can still rise or fall together.

	Crude oil	Gasoline
Amount being forecast	Commercial crude in storage; excludes the government reserve	Total motor gasoline in storage; includes finished fuel and blending ingredients
Main sources of supply	Oil extracted from the ground and imports	Refinery and blender production, after their inputs are allowed for, plus imports
What counts as domestic demand	Crude oil taken in by refineries	Product supplied: EIA’s estimate of domestic use
Connection	Raw material going into the refinery	One of several fuels made by refining and blending
Reasons for yearly patterns	Refinery activity, maintenance and available crude	Driving, fuel-blend requirements and refinery supply

The crude target excludes the Strategic Petroleum Reserve (SPR), the U.S. government’s emergency oil stockpile. The gasoline target is total motor gasoline stocks. This includes finished gasoline and blending components—ingredients that can be mixed into gasoline.

2. Data

I used monthly EIA petroleum supply data from January 2007 through June 2026. The data cover five U.S. oil-market regions, called Petroleum Administration for Defense Districts (PADDs): East Coast, Midwest, Gulf Coast, Rocky Mountains and West Coast. I added the regional forecasts to get a U.S. total.

I measured the amount in storage at the end of each month in million barrels. Supply and demand are flows, which is amounts produced, moved or used over time. I expressed them as daily rates and used the number of days in each month when calculating changes in storage.

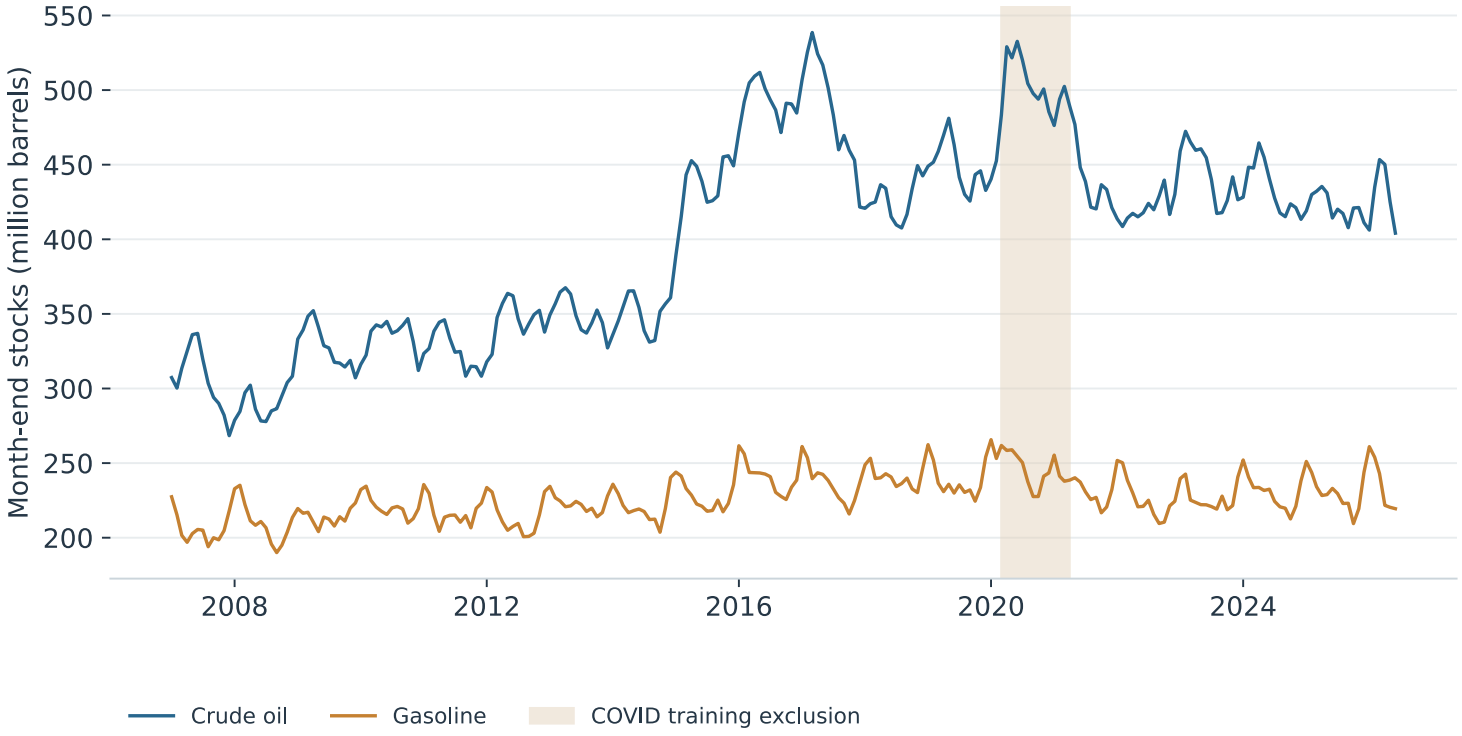

Figure 1. Two inventory pools at different stages of the oil chain. Crude stocks exclude SPR. Gasoline stocks include finished fuel and blending ingredients. The lines show amounts in storage. The shaded COVID period stays in the chart but is left out when training the models. Chart data.

Why the time of year matters

Seasonality means a pattern that tends to repeat during the year. Gasoline use usually rises during the summer driving season. Stocks tend to fall in summer and rebuild in winter. Fuel-blend requirements also change during the year. Crude stocks depend partly on when refineries run more or close for maintenance. Maintenance can reduce both crude use and gasoline production.

EIA describes these patterns in its explanations of seasonal gasoline stocks, driving demand and fuel blends, and crude storage and refinery maintenance. The patterns give me a reason to test seasonal rules.

Note: These models do not include detailed maintenance schedules, outages, weather or shipping movements.



Figure 2. Crude intake and gasoline use follow different seasonal patterns. Average patterns for complete years, 2007–2025. For each year, its own average is set to 100. A value above 100 means that month was above the year's average. I then average across years. COVID observations are included here. Chart data.

How I prepared the data

The crude data include production, imports, exports, movements between regions, refinery use and refinery capacity. Capacity is how much a refinery could process. The gasoline data include refinery and blender production, trade, regional movements, biofuel production and product supplied. Both studies also use reported adjustments and transfers. These are entries used to account for movements or differences in the reported supply balance. I adjusted gasoline production for blending-component inputs to avoid counting the same material twice.

I treated selected unreported flows as zero and marked them. Values withheld by the source stayed missing. Crude transfers before January 2022 are zero because reporting had not begun.

I left March 2020–March 2021 out of model training because COVID changed normal activity. Training means using earlier examples to work out a model's rules. Some stock-model rows use past months as inputs. I also excluded rows that depended on those COVID months, extending their training exclusion through March 2022. The charts still show the full history.

EIA can revise old observations after first publishing them. I used those revised histories and assumed the previous month's data were available at each forecast date. I am not able to recreate exactly what information was public at the time. The results therefore do not show how the models would have performed based on the data that was present at time of release.

3. Models

Checking where the oil goes

> Change in storage = amount coming in − amount going out + unexplained balance difference.

For crude, production and imports bring oil in and refinery use and exports take it out. For gasoline, refining, blending and imports add supply and product supplied and exports take it out. Movements between regions and other accounting entries also matter.

I checked whether the historical data roughly followed this balance. For example, EIA's product-supplied calculation already uses changes in stocks. Its broader crude balance can also include reserve activity that my commercial-stock target excludes.

Some differences remain in the source data. In January 2026, gasoline changes calculated from stock levels differ from reported changes in three regions.

The forecasting methods

A benchmark is a simple forecast used for comparison. The unchanged-stock benchmark assumes next month will have the same amount in storage as the latest observed month. A seasonal-change rule starts with current stocks and adds the usual increase or decrease for the next calendar month. A same-month-last-year rule repeats the amount in storage from that month a year earlier. The two seasonal rules make different predictions.

I also tested models that forecast flows and then calculate the implied stocks. A regression uses a fitted formula to relate earlier inputs to a future value. A decision tree splits earlier examples into groups using a series of rules. Random forests average many trees, boosted trees add trees that try to correct earlier errors. Neural networks learn patterns through many linked calculations. I also tested combinations of regional forecasts.

A one-month forecast gets a new observation before forecasting the next month. A year-ahead forecast instead starts once and predicts a twelve-month path: one value for each of the next twelve months. I compared methods on ten earlier periods in date order for monthly forecasts and six earlier annual paths for year-ahead forecasts.

How I judged the results

The final comparison covers July 2024–June 2026. Crude's current tuning and model selection exclude these months, but earlier models had been evaluated on this period. Some development periods also serve both parameter tuning and model comparison. A new, unseen period would provide a stronger independent test.

For gasoline, I chose the seasonal methods after seeing the final-period results, replacing the fitted methods selected earlier. This can make their scores look better than they will on new data since the same results helped choose the winner.

Mean absolute error (MAE) is the average size of the forecast misses, treating overestimates and underestimates as positive amounts. Lower MAE is better. Root mean squared error (RMSE) is another error score that gives extra weight to large misses. Both scores are measured in the same units as stocks.

Crude and gasoline have different amounts in storage, so I also divided MAE by average observed stocks. This contextualizes how large is it compared with the market’s usual stock level. It is a ratio of two averages, not the average of each month’s percentage error. I also measured how much each method reduced MAE compared with unchanged stocks.

4. Results

The seasonal-change rule helped with next month

Over the same 24 months, the crude seasonal-change rule missed by 6.78 million barrels on average. The unchanged-stock forecast missed by 9.63. Gasoline’s average miss was 3.43 million barrels, compared with 8.47 for unchanged stocks. That reduced average error by 29.6% for crude and 59.5% for gasoline.

The same 24 months: July 2024–June 2026.

Measure	Crude oil	Gasoline
Average miss: seasonal-change rule (MAE)	6.78 million barrels	3.43 million barrels
Average miss: unchanged stocks (MAE)	9.63 million barrels	8.47 million barrels
Reduction in average miss versus unchanged stocks	29.6%	59.5%
Average amount in storage	423.06 million barrels	230.16 million barrels
Average miss / average amount in storage	1.60%	1.49%
Error score giving extra weight to large misses (RMSE)	9.04 million barrels	5.09 million barrels

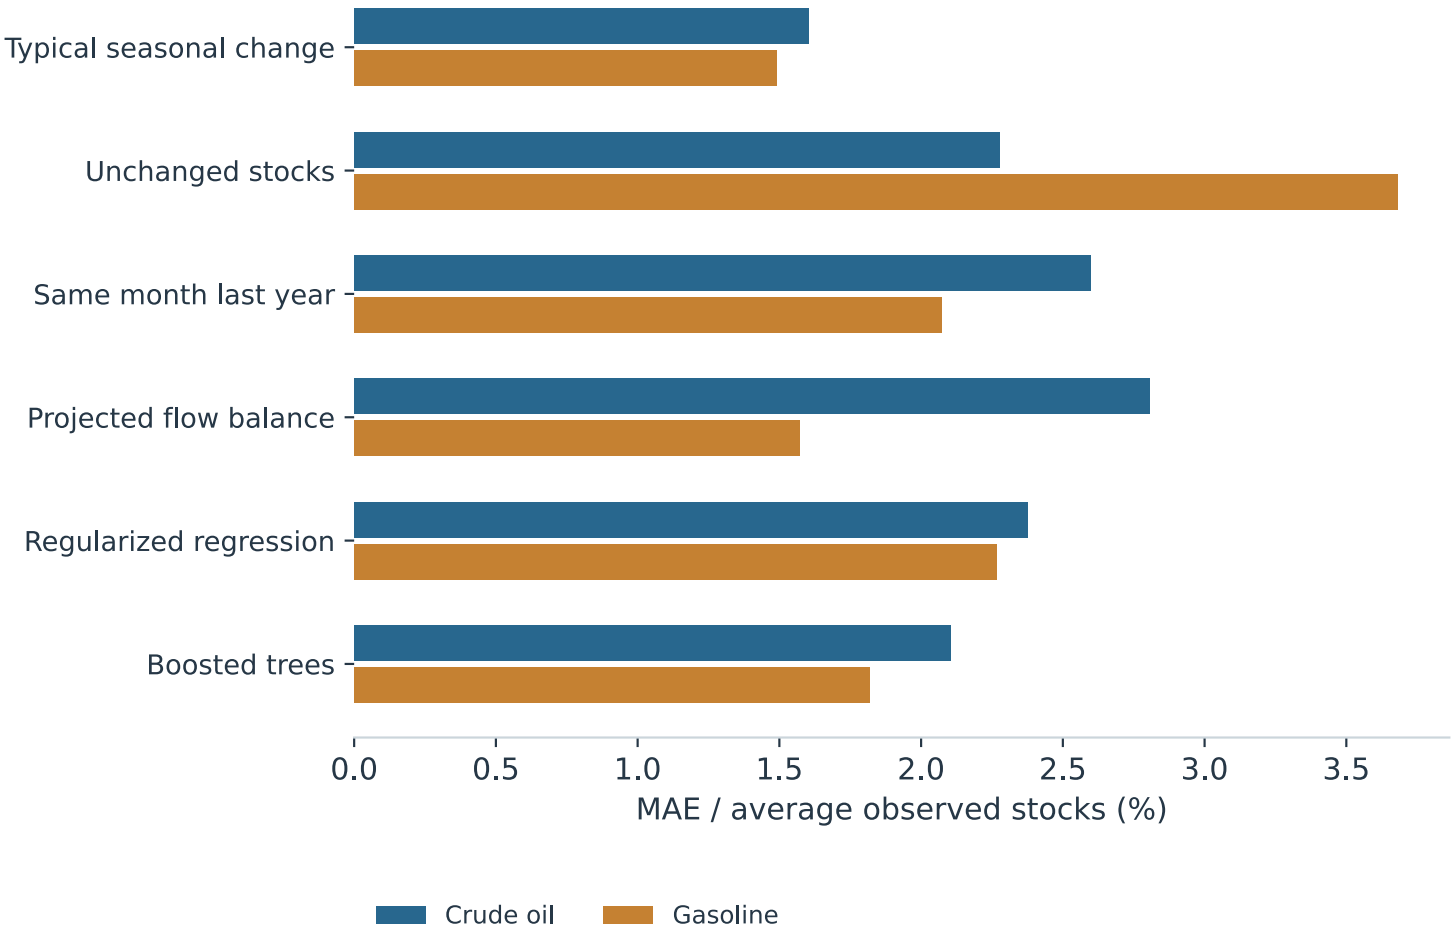

Figure 3. One-month error after accounting for inventory size. Six simple rules and model types tested over the same 24 months, July 2024–June 2026. Each bar is the average miss (MAE) divided by average stocks. Shorter bars are better. This compares the size of errors with the amount in storage. Chart data.

Gasoline’s miss in barrels is roughly half crude’s, but there is less gasoline in storage. Average gasoline stocks were 230.16 million barrels, compared with 423.06 for crude. Dividing each average miss by its average stock level gives 1.49% for gasoline and 1.60% for crude. The difference between the markets is much smaller on this scale.

Gasoline improved more over the unchanged-stock forecast in this sample.

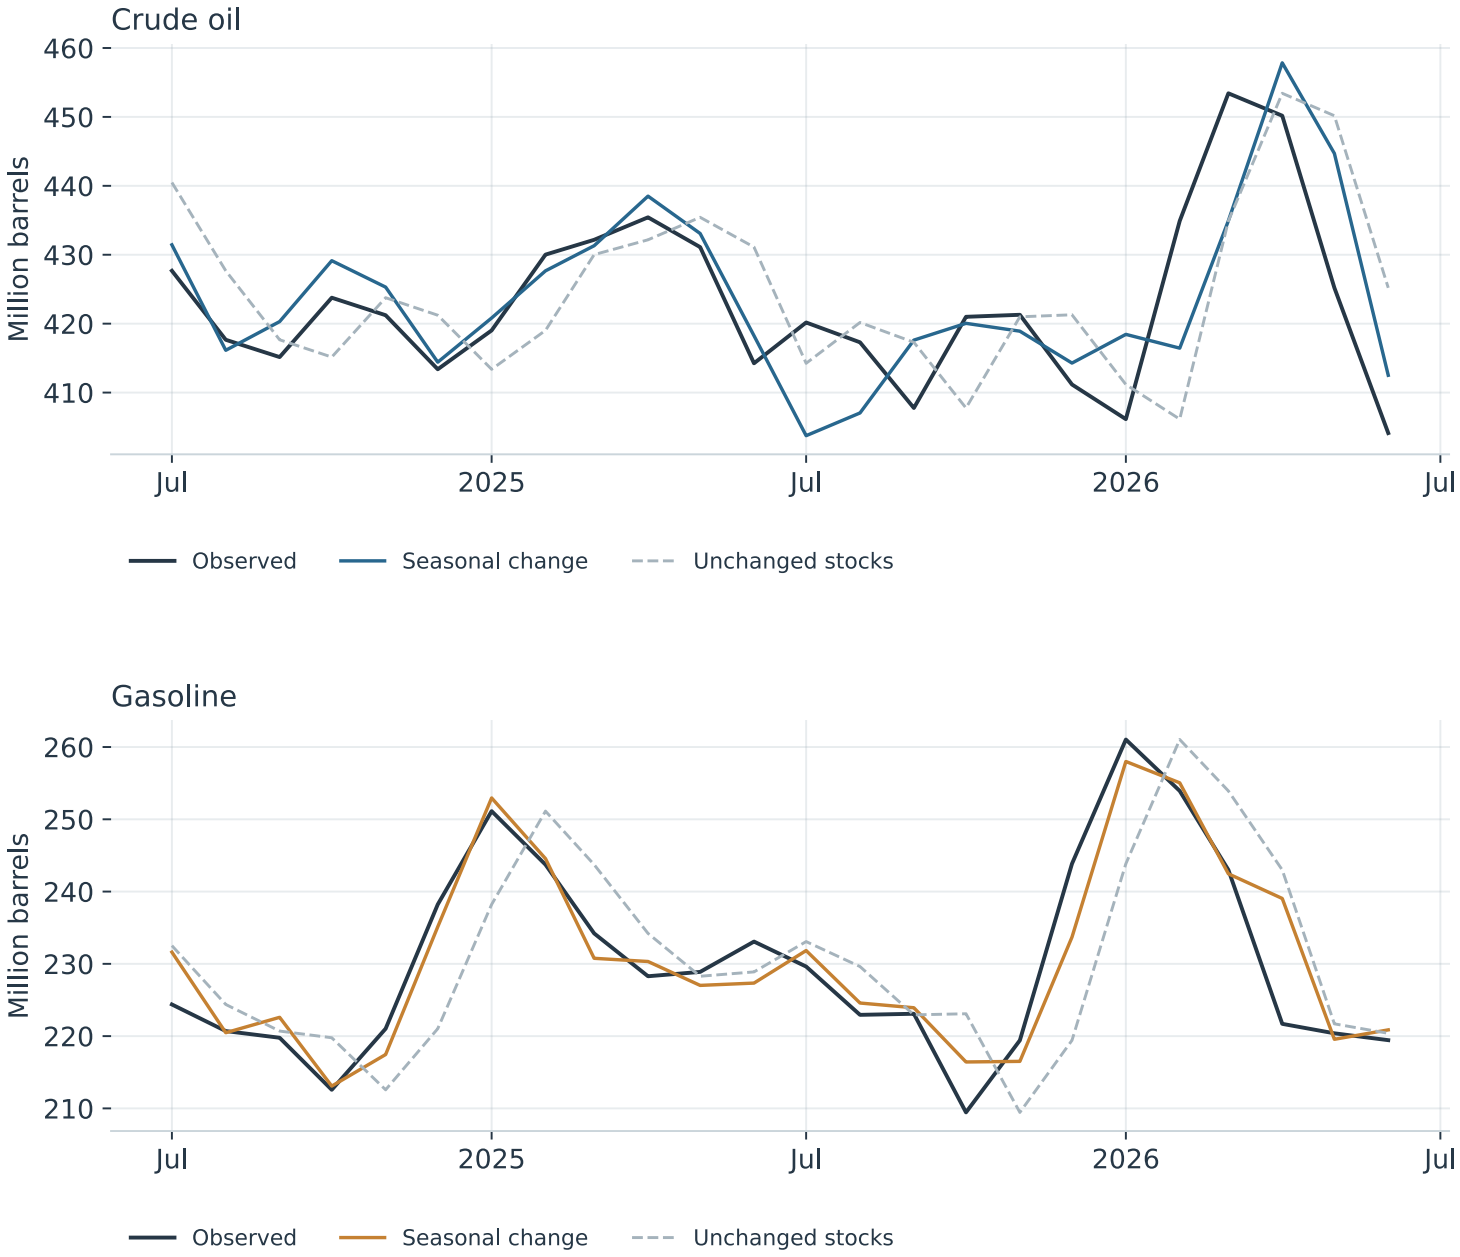

Figure 4. The seasonal rules track the broad path but miss large moves. I start a new forecast each month using the previous month's observation. The lines therefore show repeated one-month forecasts, rather than one forecast made for the whole two-year period. The panels use different vertical scales, so compare the gaps within each panel carefully. Chart data.

An average improvement still leaves large mistakes

Both seasonal rules slightly overestimated stocks on average. Their largest monthly misses were 19.50 million barrels for crude and 17.35 for gasoline.

The average improvement does not mean that every month was forecast well.

Results also differed across regions. Larger regions can have larger misses in barrels. An overestimate in one region can cancel part of an underestimate in another. I calculated the national score after adding all regional predictions, so a good national result does not mean that every region performed equally well.

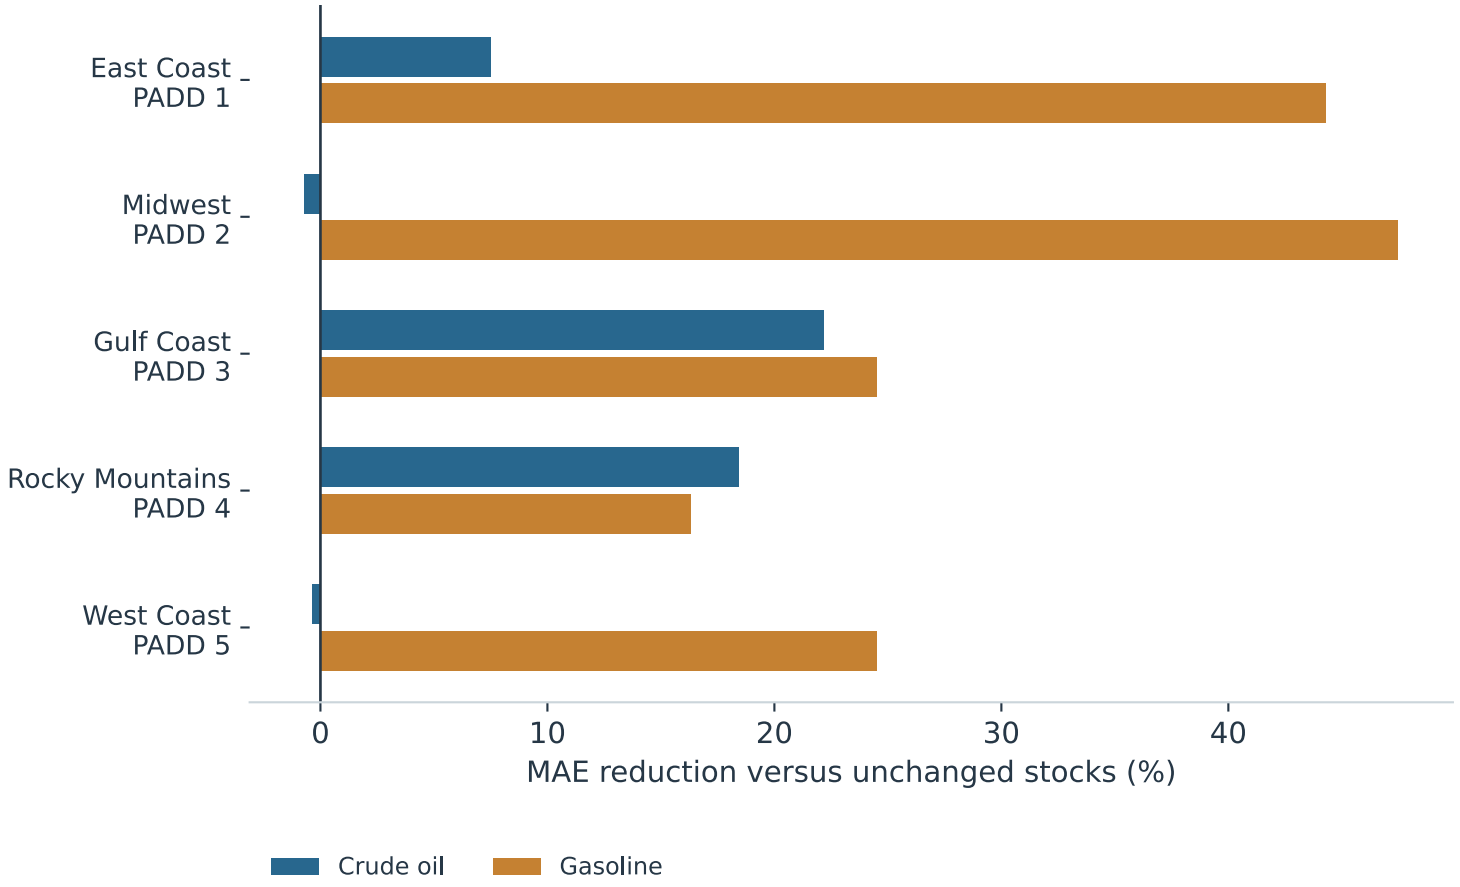

Figure 5. Regional results show where seasonality helps. Each bar compares the seasonal-change rule with unchanged stocks in one region. A positive percentage means the seasonal rule reduced the average miss. The U.S. score comes from adding the regional predictions first. It is not the average of these regional percentages. Chart data.

The year-ahead comparison is less clear

The original crude study used a regression. Its average miss was 9.06 million barrels across 13 overlapping annual paths, giving 156 forecast points. Paths overlap when forecasts started in different months cover some of the same later months. Those shared months are counted more than once, so the points are not all independent examples.

Gasoline's same-month-last-year rule missed by 4.77 million barrels across 2 annual paths, giving 24 points. Its sample was different from crude's. Comparing those two headline errors directly would mix different tests.

I therefore compared crude’s saved forecasts only from June 2024 and June 2025, the same starting dates used for gasoline. Crude’s average miss rises to 12.15 million barrels; gasoline’s stays at 4.77. Relative to average stocks, those misses are 2.87% and 2.07%. Compared with unchanged stocks, they reduce error by 19.5% and 58.8%.

Using the same dates also changes the crude ranking. Repeating last year’s stocks misses by 10.99 million barrels, less than the regression’s 12.15. For gasoline, repeating last year also beats boosted trees, which miss by 5.64. The same-month-last-year rule therefore leads the available methods for both products on these two starting dates.

Product	Method	Average miss (million barrels)	Average miss / average stocks	Error reduction versus unchanged stocks
Crude oil	Constrained regression	12.15	2.87%	+19.5%
Crude oil	Same month last year	10.99	2.60%	+27.2%
Gasoline	Same month last year	4.77	2.07%	+58.8%
Gasoline	Boosted trees	5.64	2.45%	+51.3%

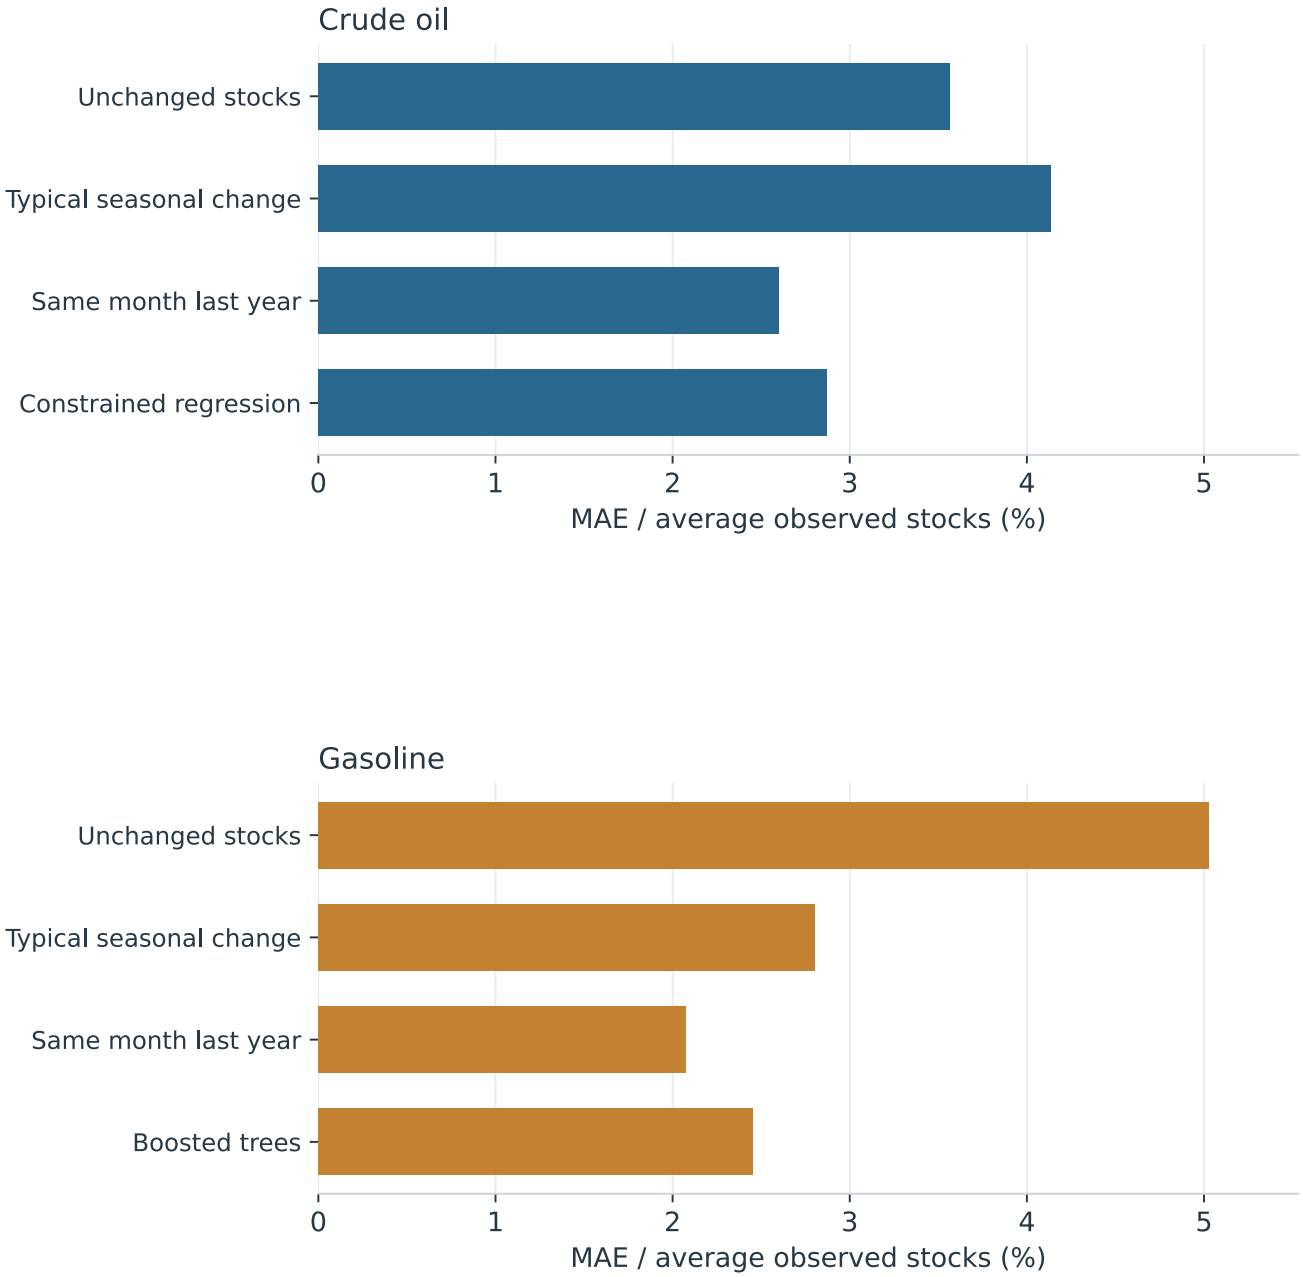

Figure 6. Matching the annual starting dates changes the crude comparison. Both products start their annual forecasts in June 2024 and June 2025, giving 24 monthly results each. I calculated crude's same-month-last-year rule from saved observations available at each starting date. This lets me compare the existing results on the same dates. Chart data.

Recent supply and use were useful inputs

In separate tests, I added each predictor to a simple forecast based on the calendar month. I used six test periods in date order to see whether that extra information reduced the error.

Recent production plus imports helped supply forecasts most in both studies. But production means crude extracted from the ground in one study and gasoline made

by refineries and blenders in the other. The most useful demand inputs were recent refinery intake for crude and recent product supplied for gasoline. In both cases, recent activity helped predict later activity.

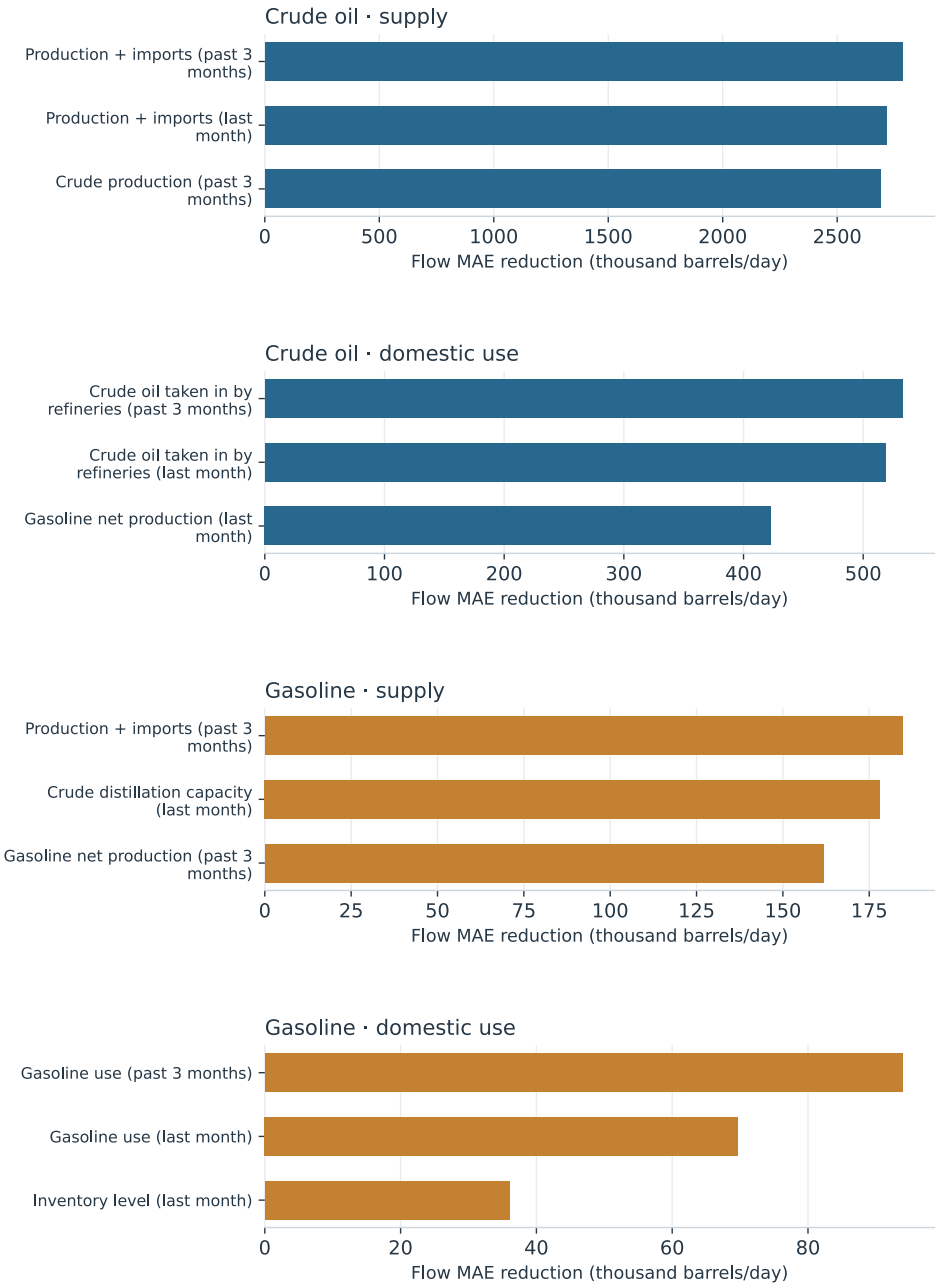

Figure 7. Recent physical flows matter in both studies. I added each input separately to a forecast based only on the calendar month, using six tests in date order. The bars show the three inputs that reduced error most for each task. Units are thousand barrels per day. Chart data.

Some stock information helped only when used with other inputs. Crude stocks compared with their year-earlier level helped its supply and demand models. The previous month’s change in crude stocks also helped gasoline models when combined with other information, even though it was weak on its own.

The outlook made with June 2026 data

The saved outlooks cover July 2026–June 2027 and use data through June 2026. Crude uses its regression; gasoline repeats the previous year's monthly stocks.

The separate forecasts of supply and demand imply different stock levels from the direct stock forecasts. For gasoline's first forecast month, the stock rule predicts a build while the projected flows imply a draw. These gaps come from differences between the models.



Figure 8. The June 2026 outlooks are separate inventory models. Dashed lines show the stock forecasts saved for July 2026–June 2027. Dotted lines show what the separate flow forecasts imply for stocks after adding inflows and subtracting outflows. Chart data.

5. Conclusion

For next month, the seasonal-change rule was useful in both markets. It reduced average error compared with assuming no stock change. Gasoline had a smaller miss in barrels and a larger improvement over that benchmark. Much of the difference in barrels narrowed after allowing for the smaller amount of gasoline in storage.

For a full year ahead, I would be more cautious. Repeating last year’s stocks led the available methods for both products when I used the same two starting dates. That sample is too limited to establish a dependable winner. Crude’s original annual forecasts reuse some of the same months, and gasoline’s methods were chosen after reviewing the reported results.

The lesson is to compare a complicated model with simple rules before trusting it. A lower average error can still hide a large miss in an individual month. Crude use and gasoline supply also need to be linked through refining as they cannot be treated as the same flow, because refineries make several products and gasoline blending adds other ingredients.

My next steps would be to test the chosen methods without changing them after seeing the answers, use data with the publication dates a forecaster would actually face, and add refinery maintenance and weather information. I would also make the stock and flow forecasts agree with a shared refinery balance. Finally, I would estimate forecast ranges to show the risk around each prediction, instead of giving only a single possible value.

Data and code

Crude oil: inventory models and results · supply and demand predictors.

Gasoline: inventory models and results · supply and demand predictors.

Download one-month results · Download year-ahead results.