

A horizontal banner with a dark blue background featuring a pattern of glowing blue dots and lines, resembling a digital or energy field. On the right side of the banner are two white line-art icons: a laboratory flask with a flame above it, and a line graph with an upward-pointing arrow.


Energy Positioning Report

Tuesday, 30 September 2025


Market positioning index – selected energy contracts

	Current week		Input Scores				29/09/2025
	Last Price	Funds	CTA	Slope	Technical	Sentiment	Position Score
ICE Brent	67.97	16	39	40	19	35	30
NYMEX WTI	63.45	-40	28	5	-14	-3	-5
NYMEX RBOB	199.51	15	4	23	-17	-10	6
NYMEX Heating Oil	235.66	28	46	12	-10	23	19
ICE Gasoil	700.25	48	48	32	15	41	36
NYMEX Henry Hub	3.27	-7	-1	-48	-6	30	-13


Position score trend





	Previous week		Input Scores				22/09/2025
	Last Price	Funds	CTA	Slope	Technical	Sentiment	Position Score
ICE Brent	66.57	21	16	28	-5	17	15
NYMEX WTI	62.64	-36	4	8	-27	-2	-11
NYMEX RBOB	197.82	15	-5	25	-30	-29	1
NYMEX Heating Oil	229.23	29	39	0	-23	-2	9
ICE Gasoil	689.00	49	43	41	-2	33	34
NYMEX Henry Hub	2.81	25	4	-46	-32	31	-13



	Week on week change		Input Scores				
	Δ Last Price	Δ Funds	Δ CTA	Δ Slope	Δ Technical	Δ Sentiment	Δ Position Score
ICE Brent	1.40	-5	23	12	24	18	15
NYMEX WTI	0.81	-4	24	-3	12	-1	6
NYMEX RBOB	1.69	0	9	-2	14	19	5
NYMEX Heating Oil	6.43	-1	7	12	13	25	10
ICE Gasoil	11.25	-1	5	-9	16	8	2
NYMEX Henry Hub	0.46	-32	-5	-2	26	-1	0




Market positioning index, and week change – selected energy contracts




Positioning inputs (left) and positioning index vs price (right) charts


ICE Brent input scores




ICE Brent position vs front month futures




NYMEX WTI input scores




NYMEX WTI position vs front month futures



NYMEX RBOB input scores




NYMEX RBOB position vs front month futures




Positioning inputs (left) and positioning (right) charts


NYMEX Heating Oil input scores




NYMEX Heating Oil position vs front month futures




ICE gasoil input scores




ICE gasoil position vs front month futures



NYMEX Henry Hub input scores



NYMEX Henry Hub position vs front month futures



The positioning index on a given underlying represents, on a given day, the deviation of a position away from its average value over a rolling 3-year period.

The positioning index takes on values between -50 and 50. When positioning is at the mean of its distribution, its value is zero, and positioning is considered neutral. When the positioning index approaches either of the extremes of the range above, this implies that the market's positioning has moved well above or below its average, which could signal that the market is stretched.

The market positioning index for a given commodity futures contract is obtained by taking a weighted average (see weighting methodology below) of the positioning index of five inputs deemed to reflect market positioning. These inputs are as follows:

- 1) Funds (Money Manager net length):** This is the open interest held by money managers on the futures contract obtained from either the CFTC or ICE Exchange commitment of traders' reports. The net length held by money managers is expressed as a percentage of the total open interest to capture the direction and magnitude of the position.
- 2) CTA positioning:** This is Onyx Capital Group's proprietary estimate of net length (in barrels) held by CTA market participants and is used as a reinforcing input to the money managers' input.
- 3) Slope (price structure):** this is the shape of the prompt portion of the futures curve, with the nearby inter-month spreads weighted by open interest. As the front month nears expiration and open interest migrates to the next month, the price structure becomes reflective of the next time spread. The steeper the backwardation, the greater the incentive to hold a passive long position due to the implied positive roll yield. The greater the backwardation, the greater the relative supply scarcity in the commodity, and again, the greater the incentive to be long the commodity.
- 4) Technical analysis momentum indicators:** This is a weighted average of the Relative Strength Index (RSI) calculated over 14 days and the MACD (12,26,9) momentum indicators. We use the MACD's histogram, which is the difference between the Signal line (a 9-day exponential moving average) and the MACD line (the difference between the 26-day moving average and the 12-day exponential moving average).
- 5) Sentiment:** This is a variation of technical analysis aiming to capture whether long or short positions are being added in the underlying based on the joint testing of daily open interest and price changes (see methodology below). We count the number of long and short positions added over a lookback period as a percentage of the total number of position changes in that period. The sum of long and shorts added always makes up 100% of the position change count, so we only need to track the evolution of the percentage of long positions added.

Positioning begins with z-scores


For each of the five inputs, we begin by calculating the z-score of each daily observation. The z-score is measured in terms of standard deviations away from the mean of the observations and is obtained simply by subtracting the mean from the observation's value and then dividing by the standard deviation of our sample period, in the case of this publication, a the 3-year rolling period of observations.

If the z-score is 0, the data point's score equals the mean. A z-score of 1 shows that the data is one standard deviation from the mean. The z-score can be either positive (above the mean) or negative (below the mean). It helps identify whether an observation is typical or atypical within the distribution of observations.

From z-scores to percentile ranking

Once the z-scores are obtained, we need to 'percentile rank' them. This means finding the percentage of scores in the distribution of the z-scores that is equal or lower than a given value, or more simply, how a particular score compares to the rest. The percentile rank is expressed on a continuum from 0 to 100%. A high percentile indicates that a given score is above the rest of the scores.

Assuming a normal distribution, 100% of the percentile ranking matches z-scores between -4 and 4, with the 50% percentile equal to a zero z-score, or the mean of the distribution (see figure below).



Constructing the positioning index with percentile ranking

When looking at the distribution of z-scores for a given input, the overall profile roughly approximates a normal distribution. This allows us to equate a percentile score of 50 with a z-score of 0 (see figure above). However, the presence of skew in the distribution suggests that the 50% percentile is closer to the median rather than the mean.

If we were to subtract 50 from the percentile of a given z-score, we would obtain a value that captures how far the z-score is away from its mean. By definition, a z-score close to the 100th percentile would have a value of nearly 50, while a score close to the 0 percentile would have a value of nearly -50. A value of zero, as indicated above, would equate to the mean or the 50th percentile. The positioning index thus becomes a series taking on values between -50 and 50.

When the index shows extreme values, above 40 or below -40, we interpret this as the market being respectively overexposed or underexposed to the underlying futures contract compared to its average exposure in the recent rolling 3-year period. At these levels, assuming mean reversion, a change in price direction is possible over the subsequent period.


Understanding the weighting of the inputs to determine the position score

To compute the position score, we calculate a weighted average of the input scores. To this end, we employ a **Random Forest decision tree model** (a machine learning process), as discussed in *Advanced Positioning, Flow, and Sentiment Analysis in Commodity Markets* by Keenan, Mark J.S. (2020). Our framework assumes position mean reversion: when the position score is well above the mean, we expect deleveraging to occur and, thus, by implication, the price to fall in the subsequent period. Vice-versa, positioning will grow when the position score is too low, pushing the price higher.

The decision tree's training incorporates the above mean reverting behavior and price dynamics to establish the weight of the inputs. In other words, the weight captures the significance of a given input in determining a position score that best predicts implied price change in the future.

Machine learning techniques, at their core, seek to uncover predictive relationships within datasets. Random Forests are particularly well suited to this task as they construct multiple decision trees and aggregate their predictions to improve predictive accuracy and robustness. Unlike a single decision tree, they are less prone to overfitting. The process of a Random Forest is depicted in the figure on Page 8. To train the model, we structure the dataset and define a forward period price return as our binary target – we either achieve or not the target. Given the sequential nature of financial markets, we use an ordered train-test split, which means that the data on which the model is trained precedes data in the following period. Finally, other treatments are applied, such as hyperparameter tuning and time-series cross-validation, so that the Random Forest method maximises predictive accuracy in both training and test samples.

The key to establishing the input weights is computing the feature importance. This means trying to quantify the contribution of each input in reducing uncertainty around predicting the overall score. The importance figures highlight which inputs influence future price movements the most. Once we have determined importance, we apply weights accordingly to the inputs to compute the position score. This approach ensures that the more predictive of the five inputs in our approach are emphasised in the weighted average, while those with lower explanatory power are de-emphasised.



Understanding the sentiment input

When there was no commitment of traders data, technical analysts looked for a workaround to infer overall position changes in the market. The analysis tests joint changes in a futures contract's price and open interest to determine whether long or short positions were being added or whether long or short positions were covered. These four outcomes are illustrated in the table below.

To build the sentiment input series, we test the conditions in the table below and then qualify the change as one of the four outcomes in the table. We then count the number of occurrences of each outcome in a lookback period to give the percentage of each outcome. The outcomes always add up to 100%. We then calculate the z-scores and percentile rank of the percentage of 'longs added', the sum of 'new longs' and 'shorts covered'.

Implied position changes based on changes in price and open interest		
Implied position	Price Change	Open Interest Change
Long positions added	▲	▲
Long positions liquidated	▼	▼
Short positions added	▼	▲
Short positions covered	▲	▼

Any recommendation, prediction, or suggestion as to an investment strategy has been prepared by Onyx Capital Advisory Limited ("Onyx") in accordance with legal requirements designed to promote the independence of investment research ("Research"). This research is directed at, and therefore should only be relied upon by, clients who have professional experience in matters relating to investments. Onyx's Research is not directed at retail clients or those in a jurisdiction in which this distribution may be restricted by local regulation or law. Onyx's publications are prepared without taking into account your specific investment objectives and financial situation, therefore before acting on any information, you should consider its appropriateness. Onyx's Research should not be regarded as a substitute for obtaining independent professional advice, including investment, tax and legal advice.

Onyx's policy is to only publish Research that is impartial, independent, clear, fair, and not misleading. Any views expressed are those of Onyx's at the time the Research was prepared. No assurances or guarantees are given as to the reliability, accuracy, or completeness of any such information or any matter contained in Onyx's Research and such Research may contain statements which are matters of judgement and which are subject to change at any time without notice. Onyx accepts no duty or liability, whatsoever, to any party in respect of its Research. Under no circumstances will Onyx be responsible for any losses incurred (whatever their nature) by its clients resulting directly or indirectly from the use or interpretation of any information contained in its Research. Such Research is solely produced and published by employees of Onyx and based on publicly available information. Past performance is not indicative of future performance.

Analysts are required to ensure that they have a reasonable basis for their analysis, predictions, and recommendations. Onyx maintains strict regulatory controls to mitigate any conflicts of interest including information barriers and restrictions on the undertaking of personal transactions in financial instruments.

Onyx is registered in England & Wales (company number 11472304) with its registered address at 95 Cromwell Road, Second Floor, London, United Kingdom, SW7 4DL. Onyx is authorised and regulated by the Financial Conduct Authority (FCA no. 822509).