The Officials Quarterly Volumes Review

Jorge Montepeque, Will Cunliffe, Ed Hayden-Briffett and Spyridon Kokas

Q3 2025

Abstract

The Officials present the Quarterly Volumes Review for Q3 2025. We document a sharp decline in traded oil derivatives volumes to 232 billion barrels, from roughly 308.2 billion barrels in the previous quarter, despite broadly stable open interest, down just 0.055%, and argue that this reflects a shift in trading behaviour rather than a structural retreat in market participation. Using a high-frequency text-based impact-probability model of Trump's Truth Social posts combined with a daily structural VAR, we show that informational shocks which generated meaningful increases in trading activity in Q1–Q2 largely ceased to elicit a volume response by Q3, providing direct evidence of headline fatigue. In parallel, we analyse the evolution of the global total-liquids balance and show that the mainstreaming of sizeable forward surpluses has encouraged a migration from short-horizon, headline-driven trading toward longer-duration, fundamentals-anchored speculative positioning. A case study of Dubai swaps highlights how localised episodes of intense activity can coexist with a broader low-intensity regime as benchmark structures evolve.

Building on these findings, we develop an empirical forecasting framework that links traded volumes and open interest to macroeconomic conditions, yield-curve dynamics and structural growth in market depth. Under our baseline assumptions, we project that total traded volumes will recover and resume their upward trend, rising by 11.05% in 2026 to around 1.18 trillion barrels and by a further 16.8% in 2027 to approximately 1.38 trillion barrels. Open interest is also expected to grow, albeit at a lower rate, increasing by roughly 8.2% in 2026 and 7.1% in 2027. The overarching conclusion is that Q3 2025 marks a temporary episode of fatigue and lower trading intensity within an otherwise robust, secular expansion of global oil derivatives markets.

Contents

1 Introduction

The first edition of the Quarterly Volumes Review emphasised the record-breaking quarter of trading volumes, continuing the overall uptrend that we identified. We determined that, in Q2, the primary driver of traded volumes was geopolitical events, made more potent by increasing access to the derivatives market.

The impetus provided by such major geopolitical and trade events pushed traded volumes in Q2 to record levels at over 308 billion barrels traded. In Q3, by contrast, traded volumes dropped to 232,169,303,661 barrels. That was the lowest quarterly traded value since Q4 2024 and lower than Q3 2024 values. Indeed, on a y/y basis, we recorded a 10.3% drop.

In terms of open interest, there was not a commensurate fall alongside the decline in traded volumes. This is a point of particular interest, which we shall explore in detail in the subsequent Money Managers section, where we analyse the change in speculative behaviour with regard to headline fatigue and the transition to longer-term position taking, rather than short-term trading on heightened volatility. This drives higher open interest yet lower traded volumes, a key theme of Q3.

In terms of thematic and narrative changes between Q2 and Q3, there was a major shift in market focus. As discussed in the previous Review, the imposition of wide-reaching and extensive tariffs by US President Donald Trump in early April and the 12 Day War between Israel and Iran in June were the standout events of Q2. By contrast, Q3 saw no similarly significant geopolitical impulses but was instead characterised by the emergence of long-term supply and demand narratives and an increasing numbness to headlines that would previously have precipitated a significant rise in traded volumes.

The third quarter saw a marked decrease in volatility in oil prices. Prompt Brent futures traded between a low of under $60 and a high above $80 in Q2 but this range was much tighter in Q3, between $65 and $74. Rather than on flat price, the market's attention has centred on the futures curve structure and the activity of Chinese buyers. Continued paper increases to production quotas by OPEC+ failed to provide meaningful impetus to traded volumes, as the market disregarded these as insignificant.

Through Q3, there were several benchmark moves worthy of note. The Dubai spot market, in particular, was under scrutiny, as a major 'play' unfolded. This forms the basis for our following case study on Dubai swap volumes.

Furthermore, since the publication of the previous Quarterly Review, there has been a marked change in monetary policy stance. In the US, the Federal Reserve has taken a more doveish position, having cut the federal funds policy rate by 25 basis points in its September meeting, lowering it from 4.25% to 4% by the end of the quarter. The change in monetary policy goes some way to explain the change in our open interest forecast, as cutting rates means that liquidity remains ample to enable speculative position taking. Yet, these positions taken have been of longer-duration than previously due to the overarching bearish narrative, low volatility, resulting in lower-than-expected volumes but higher than expected open interest.

In the Q2 Review, we forecast open interest to fall by 2% q/q. We actually observed a drop of 0.055% to 16.68 billion barrels, aided by the change in speculative behaviour of market participants. That constituted a surprise to the upside of 1.945%.

Volumes were distinctly low, despite this elevated level of open interest. Indeed, traded volumes disappointed our forecast published in the Q2 Review by 5% to the downside, realised at just over 232 billion barrels, against a forecast exceeding 245 billion barrels.

As volumes fell in Q3, the share prices of major exchanges also struggled. ICE stock fell by 8.1% from the open of July trading to the final close of September, while CME slid 2.7%. As a reference point, the S&P 500 returned 8.1% over the quarter.

The focus of this, second edition of the Quarterly Volumes Review, is on the headline fatigue that permeated the market, leading to lesser reactions to shocks, and the behaviour of speculative traders, as money manager positioning changed. We also provide an updated forecast to the end of 2027 for both traded volumes and open interest, on the basis of the additional data observed during Q3.

2 Volumes over time

As noted in the Quarterly Volumes Report for the Second Quarter, volumes from April to June this year were exceptional and indeed record-breaking. We observed a total of 308.784 billion barrels traded in Q2 and we had forecast Q3 traded volumes of only 245.280 billion barrels.

In fact, we observed 232.169 billion barrels, over 13 billion barrels lower than expected. As we shall discuss through the Review, we advance the view that this is the result of increasing levels of headline fatigue to impulses of geopolitical nature.

Indeed, traded volumes in Q3 were significantly lower across all three months than their level in each month in Q2. The lowest monthly volumes we assessed in Q2 were in May, with 87.675 billion barrels traded. That was higher than the highest monthly volumes we assessed in Q3, which were recorded in July at 86.42 billion barrels. August was a particularly weak month, with the lowest monthly traded volumes since December 2024. Furthermore, on a quarterly basis, traded volumes in Q3 were their lowest since Q4 2024. The cumulative effect of this was an aggregate drop of nearly 25% in traded volumes q/q and a fall of 10.3% y/y, as shown in Figure 1.

In this paper, we posit that this weakness in traded volumes in Q3 is largely the consequence of significant market fatigue to previously inciting geopolitical and wartime events. We endeavour to illustrate this through Trump Index Factors, as an indication of the market's sensitivity and responsiveness to political incitement. Price discovery is of course driven by the realisation of new information (Tetlock, 2007), even through social media channels (Bollen et al., 2011), and indeed those run by President Trump (Chen et al., 2023). This information can originate from many sources, and throughout the first two quarters one key source has been the Truth Social account of President Donald Trump.

One of the key features of Q3 was that volumes, as demonstrated above, were weak on a relative basis to recent quarters, but open interest remained robust, only 0.055% lower than in Q2, as shown in Figure 2. This indicates that the quantity of derivatives contracts held open remains elevated but that the frequency of trading/exchange has fallen, in line with market fatigue to impulses that previously instigated an increase in traded volumes.

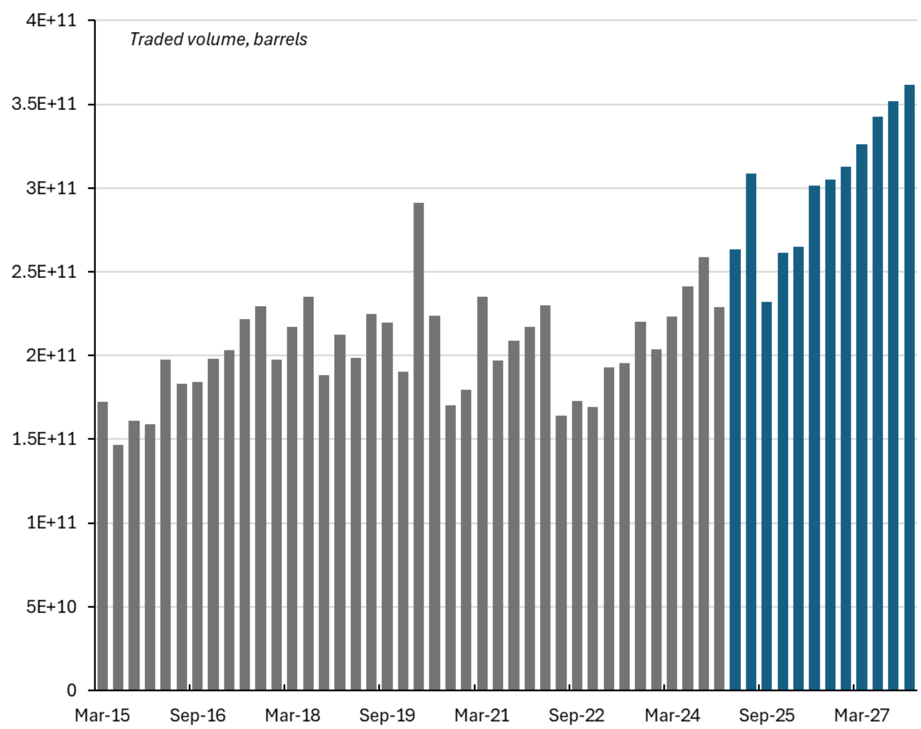

Figure 1: Total oil derivatives traded volumes as assessed and forecast by The Officials

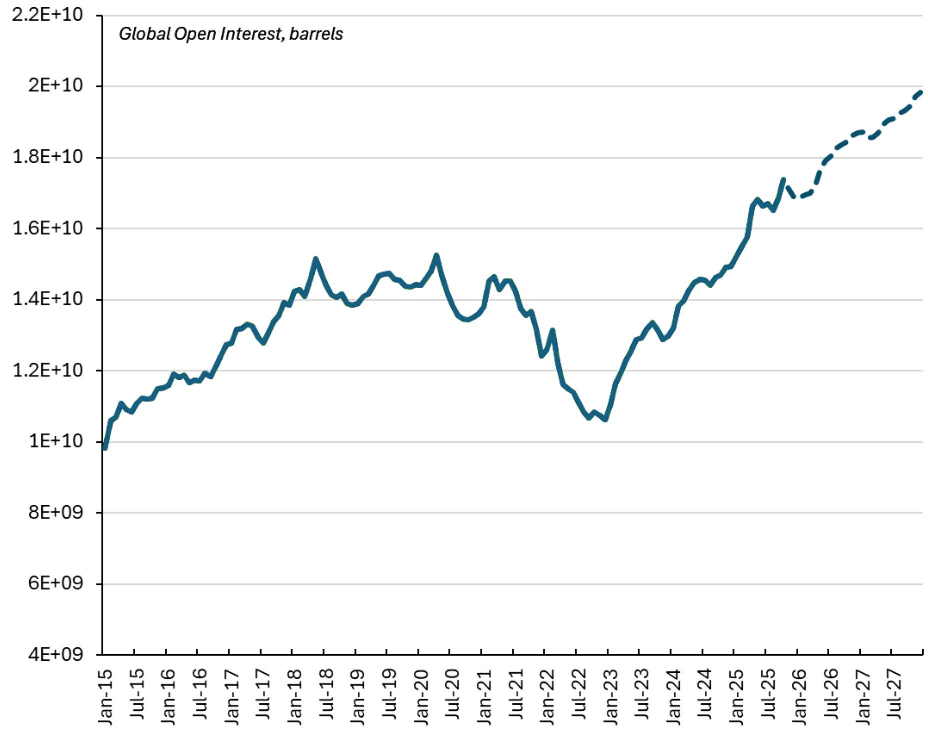

Figure 2: Total oil derivatives open interest as assessed and forecast by The Officials

Additionally, by contrast to Q2, when we saw large scale opportunistic hedging in large part due to the impact of the Israel-Iran war on prices, Q3 was relatively subdued in terms of hedging activity. The 12-Day War strengthened flat price and bolstered product crack spreads significantly, incentivising widespread hedging by both producers and refiners seeking to lock in attractive margins. The subdued pricing environment in Q3, as prompt Brent futures declined towards the low-$60 range again, did not inspire a strong hedging reaction from commercial players.

Another key feature of the third quarter was the emergence of the ‘super glut’ narrative to the mainstream market conversation and attention. Balances , as described by major institutions and banks had been developing significant surpluses for some time but this came to a head in Q3. The International Energy Agency’s August Oil Market Report showed a surplus of 2.96 million barrels per day in its projection for the global oil balance for 2026 – subsequently revised up to exceed 4 million barrels per day in Q1 2026 in following reports (International Energy Agency, 2025).

Morgan Stanley referenced the weakness of the Brent curve structure as early as April but this was a slow-burn development that came to widespread attention in Q3. The forum of conversation at APPEC in early September enabled discussions of excess supply to spread like wildfire and to grab the market’s attention. The Officials even detected mentions of long balances from major National Oil Companies that have a vested interest in balances showing shortfalls.

This led to an increase in money managers taking long-term positions on the basis of a surplus in the coming quarters and years, rather than short-term positioning. Such a change in trading dynamic by speculative players results in higher open interest but lower trading volumes; the volume of transactions is lower but the positions remain open for a longer period of time. As discussed above, this dynamic proved accurate, as open interest remained steady while traded volumes fell significantly q/q.

One of the other consequences of the change in speculative trading activity is the counterbalancing effect it had on hedging activity. Higher open interest from money managers helped to offset lower open interest arising from reduced hedging activity, while failing to compensate for lower traded volumes.

We shall analyse the role of money managers in greater granularity in the Dynamics in Open Interest section.

3 Drivers

3.1 The Trump effect

One theme has consistently permeated market commentary through the third quarter, namely the idea of "headline fatigue". By headline fatigue, we refer to a regime in which market participants are exposed to such a high volume and frequency of news, official statements and social media interventions that their marginal propensity to trade on any individual item of information diminishes. Information is no longer scarce; instead, the binding constraint becomes attention and risk appetite. In this environment, even objectively material developments can elicit only muted trading responses if they arrive in the context of a continuous stream of provocative or contradictory signals (Tetlock, 2011).

The Trump administration has provided an almost textbook case of this phenomenon. The rapid alternation between escalation and de-escalation on trade, sanctions and diplomatic issues has generated a sequence of sharp but often short-lived headlines, particularly via Truth Social. Markets have been repeatedly forced to reprice the outlook for tariffs, growth and geopolitical risk, only for the narrative to be reversed or diluted within days. Over time, this whiplash effect plausibly encourages a shift in market behaviour, away from reacting mechanically to each new post and towards discounting much of the flow as background noise unless accompanied by clear policy follow-through.

In the first two quarters, rhetoric emerging from Trump truth social posts had a demonstrable effect on market condition. One such example would be President Trumps Truth Social post on the 7 March this year, as displayed in Figure 3. The intraday evolution of Brent flat price on 7 March 2025 indicates a clear and sizeable market reaction to President Trump's post regarding renewed Russian attacks on Ukraine (text shown in the inset). In the hour preceding the post (13:06–14:16), Brent trades in a narrow and slightly declining range around \$70.3–70.5/bbl, briefly touching \$70.2/bbl, consistent with relatively orderly conditions. The release of the post, marked by the vertical line at approximately 14:16, coincides with an immediate discrete increase of roughly \$0.3–0.4/bbl and is followed by a sustained upward drift, with prices exceeding \$71.3/bbl by 15:06. The pattern is characteristic of an information shock that is interpreted as raising geopolitical risk and, by extension, the crude oil risk premium: an instantaneous jump in the level of prices, followed by continued buying as market participants reassess the implications of heightened conflict intensity for future supply and security of flows. In the third quarter, for the same impulse, the market has been notably devoid of such violent response. The markets have fatigued.

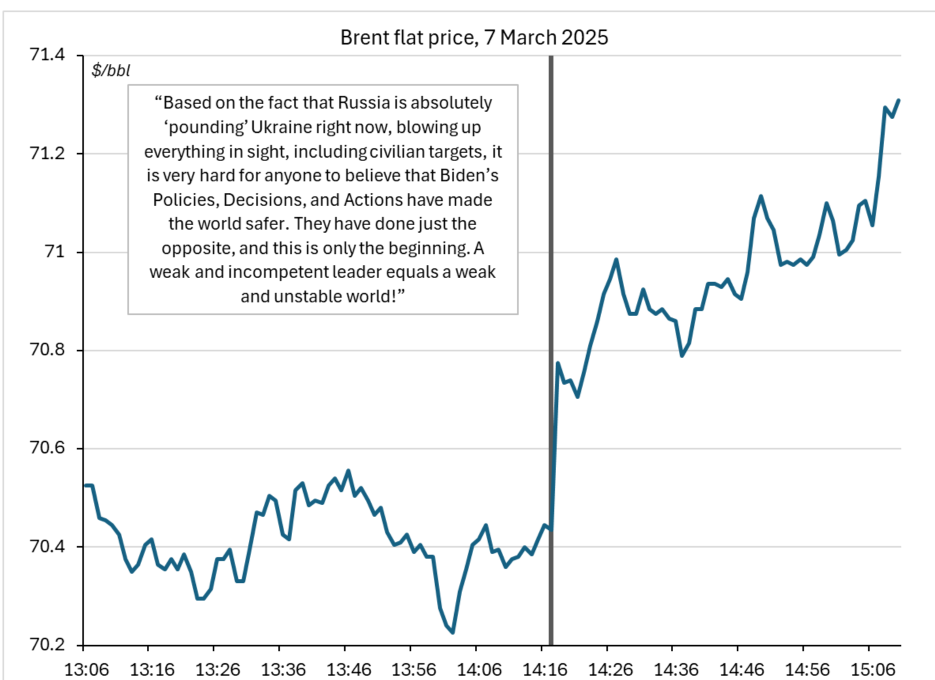

Figure 3: Brent Futures front month price response to President Donald Trump Truth Social post on 7 March

Given the centrality of this dynamic to trading conditions in 2025, it is natural that the second edition of the Quarterly Volumes Review should go beyond anecdote and quantify the extent to which headline fatigue has affected market activity. In what follows, we treat Trump-related social-media output as a laboratory for studying the interaction between political communication and trading behaviour. By combining a high-frequency text-based impact measure with a daily VAR framework for residualised trading volumes, we are able to trace how the market's responsiveness to these shocks has evolved over time, and to assess whether the Q3 collapse in traded volumes is consistent with a genuine attenuation in the information-to-volume transmission mechanism.

3.2 Fundamental narratives

Another clear driver of trading activity this month has been the evolving rhetoric surrounding the global balance of oil supply and demand, or rather the growing imbalance. Termed the ”super glut”, analysts and commentators around the world have inspected their spreadsheets to find something rather concerning: according to their calculations, next year the surplus is expected to be monumental.

Over the third quarter, the fundamental balance for total liquids has shifted in a more clearly bearish direction, and market behaviour has gradually converged towards that assessment. Downward revisions to demand, particularly in OECD economies, have coincided with resilient non-OPEC supply and a slower pace of inventory draws than anticipated earlier in the year. The IEA now estimates that the global oil market ran an average surplus of around 1.9 mb/d over January–September 2025, with forward balances for 2026 revised up to a surplus on the order of 3–4 mb/d in recent publications (International Energy Agency, 2025).

Several banks have reached similar conclusions. For example, Morgan Stanley now projects a surplus of roughly 0.8 mb/d in Q4 2025 and around 1.1 mb/d in the first half of 2026, while other sell-side houses place late-2025/2026 surpluses in the 1.8–2.4 mb/d range.

The cumulative effect of these revisions has been to erode the perception of imminent tightness and to normalise the idea of a forward surplus in the total-liquids balance. In contrast to the episodic dislocations observed in 2022–23, the Q3 data point towards a more structural loosening, driven by conventional supply–demand dynamics rather than idiosyncratic shocks.

This re-evaluation of fundamentals has been reflected primarily in the more deferred term structure, rather than in persistent stress at the front of the curve. Excess supply ultimately ends up in inventories, which softens time spread structures (Kilian and Murphy, 2014). Other than at the very prompt, Brent and Dubai time spreads have compressed, signalling a reduced scarcity premium, while outright prices have traded in a relatively narrow range despite a steady drumbeat of geopolitical headlines. In other words, the market has increasingly internalised the bearish total-liquids outlook via the shape of the curve and medium-horizon expectations, rather than through repeated repricing of prompt flat price in response to news flow. This is consistent with a transition from a headline-driven regime to one in which balance-sheet projections and the inventory trajectory dominate price discovery.

Speculative positioning has adjusted accordingly. The Q3 data show that open interest has been broadly maintained, even as traded volumes and trading intensity have declined. This pattern suggests that speculative risk has not been withdrawn from the market; instead, it has migrated into longer-duration, fundamentals-anchored positions that are turned over less frequently. Money managers increasingly appear to be expressing views on the anticipated surplus in total liquids over coming quarters, rather than seeking to monetise short-term volatility around political or geopolitical events. The weakening of the impulse response of volumes to high-frequency informational shocks is therefore consistent with a deeper shift in trading style: market participants remain engaged, but increasingly through slow-moving, balance-driven positioning rather than reactive headline-driven trading.

This shift in speculative engagement also manifested in the term structure of activity. While front-month volumes and day-to-day turnover softened, trading in the more deferred Brent contracts – particularly around the Dec26 tenor – picked up notably into the end of the quarter. This reinforces the idea that speculative participation was not withdrawn, but reallocated into longer-dated, fundamentals-anchored risk, consistent with investors positioning for the expected liquids surplus rather than chasing short-horizon volatility. In that sense, the behaviour documented in the trading-intensity analysis below is entirely coherent: open interest held steady, but the cadence of turnover slowed, and the market transitioned into a low-frequency regime where positions were held longer and adjusted less often.

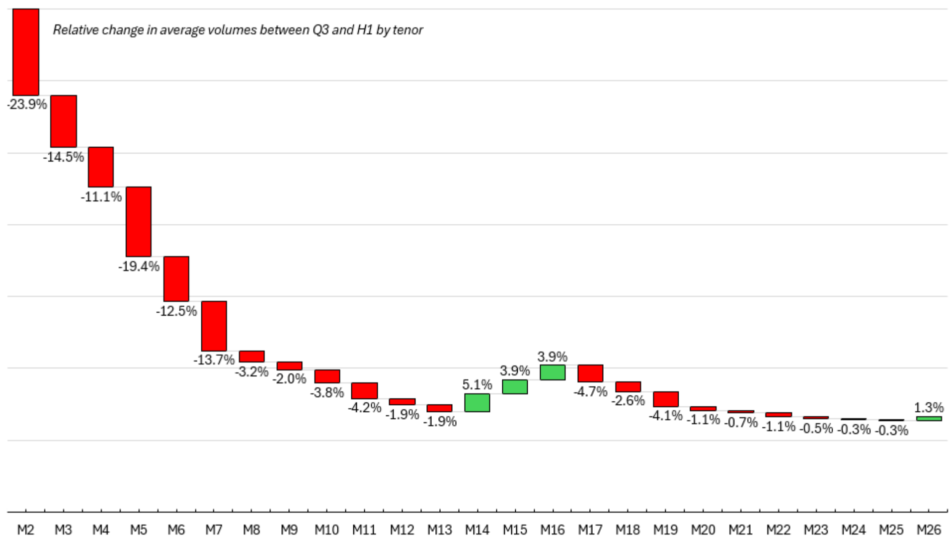

Figure 4: Relative Percentage Change in Average Tenor-Specific Trading Volumes vs. Average Change in Total Brent Futures Volumes in Q3 vs H1

4 Does anyone listen to Trump anymore?

This section develops a rigorous framework for constructing high-frequency identifiers that map political communication to contemporaneous Brent crude oil price dynamics. The objective is to quantify the immediate informational content embedded in political posts by observing the short-run market response to each post within a precisely defined event-time window. This identifier system establishes the empirical foundation for the sentiment and time-series modelling that follow, as it produces the labels necessary to train a text-based impact model and, simultaneously, characterises the short-run economic relevance of individual posts.

4.1 Data and Time Normalisation

4.1.1 Truth Social Archive

This study uses an archive of all public Truth Social posts ("Truths") published by President Donald Trump between **1 January and 30 September 2025**. The dataset contains 4,654 entries, of which 3,633 include non-empty textual content suitable for classification and analysis. For each post, the archive provides timestamps, text, repost counts, reply counts, and "favourites" (likes).

Table 1 summarises posting activity and average engagement across quarters. Posting frequency remains relatively stable across Q1–Q3. However, engagement patterns show a clear decline: average replies, re-shares, and favourites fall sharply after Q1 and remain subdued throughout Q2 and Q3. This indicates a cooling audience response despite consistent output.

Table 1: Quarterly Posting Activity and Engagement (Q1–Q3 2025)

Quarter	Posts	Avg. Replies	Avg. Re-shares	Avg. Likes
2025 Q1	1,584	1,132	4,078	18,278
2025 Q2	1,480	733	2,017	8,478
2025 Q3	1,590	715	1,866	7,482

To identify thematic patterns, posts were categorised using keyword-based classification into nine topical clusters. Table 2 reports the distribution and corresponding average engagement metrics. High-engagement categories include media criticism, US politics/elections, and immigration/border topics, each attracting above-average likes and re-shares. Legal-case-related posts and foreign-policy commentary also feature prominently. Campaign-related promotional posts constitute a very small share of overall activity.

Table 2: Topic Distribution and Engagement Metrics

Topic Category	Share (%)	Avg. Replies	Avg. Re-shares	Avg. Likes
Other / mixed	56.0	698	2,264	10,072
US politics / elections	15.9	1,390	3,593	14,726
Economy / business	13.1	1,111	3,249	14,167
Foreign policy / war	6.4	1,187	3,237	14,588
Immigration / border	2.9	1,161	3,355	14,634
Media criticism	2.9	1,238	3,640	16,135
Legal cases / investigations	2.4	927	2,960	12,231
Campaign events	0.4	1,250	3,354	15,365

Taken together, the data indicate that although Trump's posting frequency remains high, user engagement declines materially through 2025. Thematically, attention concentrates on politics, media criticism, the economy, and foreign policy, with these topics generating disproportionate interaction. This provides a structured empirical basis for evaluating the content, tone, and political communication strategies employed during the period.

4.1.2 Brent minute-level price data

The crude oil price data consist of a dense time series of Brent futures prices recorded at one-minute frequency. All timestamps are converted to Coordinated Universal Time (UTC), which is the reference timezone used throughout the entire methodology, in order to ensure temporal consistency. The resulting minute-level price series is cleaned and transformed into logarithmic price levels. Log transformation ensures that all subsequent return calculations are time-additive and symmetric for price increases and decreases.

4.1.3 Political post timestamps

The Truth Social communication data contain a timestamp field marking the time a post is published on the platform. All post timestamps converted to UTC and are then rounded up (ceiled) to the nearest minute. This reflects the earliest point in clock-time at which the post could be absorbed by market participants, and prevents the misalignment that would occur if posts with second-level timestamps were mapped to the preceding minute.

All posts are assigned a unique event identifier and an auxiliary "minute-before" timestamp defined as exactly one minute prior to the post's ceiled timestamp. This ensures the existence of a pre-event reference price.

For each post, a symmetric ±5-minute window is constructed. Let t_i denote the UTC minute-timestamp of post i. Then define:

Pre-event window:

$$\{\, t_i - 5,\ t_i - 4,\ \ldots,\ t_i - 1 \,\}$$

Post-event window:

$$\{\, t_i + 0,\ t_i + 1,\ \ldots,\ t_i + 4 \,\}$$

5 Impact-Probability Modelling

The second component of the empirical strategy develops a text-based measure of sentiment derived from the content and stylistic features of Truth Social posts, similar to (Tetlock, 2007). The aim is to map each post into a continuous index reflecting the ex ante probability that the post is capable of generating an economically significant reaction in crude oil markets.

This mapping converts unstructured text into a structured numerical series that is interpretable, temporally aligned with market data, and suitable for econometric modelling. Central to this framework is the construction of the "impact probability," a model-based estimate of the likelihood that a given post triggers a high-frequency market movement. This probability serves as the core impact-probability model: periods of frequent, highly predictable, or repetitively themed posts correspond to elevated fatigue, whereas posts that are novel, informationally rich, or stylistically distinct are assigned higher impact probabilities.

5.1 Text Normalisation Processing

All posts undergo a rigorous normalisation pipeline to ensure consistency in tokenisation, embedding, and model input representation. URLs, emoji, and hastily repeated punctuation are preserved unless they distort semantic information, because these features can contain impact-probability or emphasis relevant for classification. Exact duplicates- posts with identical character sequences - are removed. Posts that are empty after cleaning or that fall outside trading-relevant hours are excluded.

To train a supervised learning model capable of predicting market-moving content, each post must be labelled with an empirical outcome measuring its realised impact. The identifier dataset provides a five-minute return metric quantifying how Brent prices responded to each post as described by (5.) in section 4.3. Let r_i denote the five-minute delta return following post i, adjusted for ex-ante price action. We compute the Post-minus-pre delta returns as described in Appendix A.

To define "impactful" posts, we examine the distribution of r_i across the first two quarters of the sample period. The upper quartile breakpoint is

$$\tau = Q_{0.75}(r_i). \tag{1}$$

This threshold distinguishes background posts from those associated with economically relevant market reactions.

Thus, the binary label is:

$$y_i = \mathbf{1}\{\, |r_i| \geq \tau \,\}. \tag{2}$$

This procedure yields a conservative, data-driven classification problem: the median post has negligible effect, while the top quartile is treated as materially impactful. The classifier's task is to estimate the probability $P(y_i = 1 \mid \text{text}_i)$.

5.2 RoBERTa-Based Supervised Learning Architecture

The impact-probability model consists of a fine-tuned RoBERTa encoder (Liu et al., 2019). RoBERTa's contextual embeddings capture syntax, pragmatics, sequencing, and rhetorical structure, making it well suited for modelling political language, which often derives emphasis from sequential buildup, contrast, and rhetorical framing. Crucially the model is trained exclusively on Q1 and Q2 data, during the pre-fatigue regime.

The model architecture is two-stage:

(i) **Stage 1: Binary classification** A linear classification head is attached to the pooled RoBERTa embedding. The model is trained to minimise cross-entropy loss over the binary impact label.

(ii) **Stage 2: Robust magnitude regression** For posts classified as impactful, a second model estimates the conditional magnitude of the five-minute return. A regression head is trained using a Huber loss, chosen for robustness to outliers common in high-frequency financial data.

To prepare the shock series for time-series modelling, two transformations are applied:

(i) **Winsorisation** S_d is winsorised at the 0.5th and 99.5th percentiles, producing a cleaned series S'_d that suppresses extreme tail behaviour caused by rare event clusters.

(ii) **Rolling Z-Score Standardisation** To capture fluctuations in fatigue relative to local historical norms, a 60-day rolling mean μ_d and standard deviation σ_d are computed:

$$Z_d = \frac{S'_d - \mu_d}{\sigma_d}. \tag{3}$$

The Z-score Z_d is the processed shock used in the VAR. It represents information pressure relative to the prevailing environment and isolates unexpected surges or collapses in informational relevance.

6 Structural VAR (SVAR) Framework

6.1 Overview and Motivation

The final stage of the empirical analysis examines how political impact-probability shocks propagate through crude oil market activity using a structural vector autoregression (SVAR) in the spirit of Shapiro et al. (2022). Whereas earlier sections analyse event-level and daily contemporaneous outcomes, the SVAR captures multi-day dynamics, feedback effects, and the full transmission mechanism linking informational shocks to trading volume. The informational shocks used in this section are the counterfactual impact-probability innovations generated from the Q1-Q2 impact-probability model. These shocks measure the intensity of information flow under the pre-fatigue regime. Headline fatigue is therefore inferred not from the shocks themselves but from the degree to which their dynamic effect on trading behaviour weakens over time, in line with approaches that identify exogenous shocks externally and study their evolving impulse responses and counterfactual implications (e.g. Ramey, 2011).

A crucial feature of the design is the temporal structure of information arrival. Political posts are timestamped intraday and always precede the daily realisation of trading volume. The impact-probability shock for day d is thus predetermined with respect to all market variables on day d, ensuring weak exogeneity. This timing restriction yields a natural recursive structure: the informational shock may influence trading volume contemporaneously, but trading volume cannot influence the construction of the shock on the same day.

Empirically, a reduced-form VAR(3) is estimated for each quarter using two variables: the standardised daily informational shock and the residual component of daily log-volume growth after removing deterministic controls (day-of-week effects and smoothed open interest). Structural identification is then imposed using a recursive (Cholesky) scheme in which the informational shock is ordered first. This ordering implies that structural innovations to the shock equation correspond to exogenous information shocks, while trading volume is allowed to react contemporaneously but cannot feed back into the shock within the same day.

Under this recursive identification, the Cholesky factor of the reduced-form covariance matrix recovers the structural impact matrix, and the orthogonalised impulse response functions (IRFs) correspond to the structural responses implied by a recursive SVAR. Thus, although the system is estimated in reduced form, the dynamic responses reported in this section are structurally interpretable and equivalent to those obtained from a classical SVAR.

The remainder of this section describes the construction and residualisation of daily volume growth, the aggregation and standardisation of the counterfactual informational shocks, the estimation and stability assessment of the reduced-form VAR(3), and the recovery of structural shocks and structural IRFs using the recursive SVAR identification. This approach is well established in the financial econometrics literature, particularly in studies examining high-frequency news shocks and their transmission through market-microstructure variables.

6.2 Data Treatment

The outcome variable is derived from daily Brent crude oil trading volumes, measured in barrels. Raw volume typically exhibits strong periodicity, including day-of-week effects, monthly rollover patterns, and holiday-induced discontinuities. Using raw volume in a VAR would produce spurious dynamics and contaminate innovation-based identification with deterministic seasonal structure. Therefore, a series of transformations is applied to isolate the economically meaningful component - unexpected changes in trading activity.

First, raw daily volume V_d is smoothed using a five-day moving average:

$$\tilde{V}_t = \frac{1}{5} \sum_{i=t-4}^{t} V_i \tag{4}$$

This transformation stabilises the series against mechanically low or high observations associated with idiosyncratic trading conditions. The smoothed volume is log-transformed using:

$$\log V_d = \log(1 + \tilde{V}_d), \tag{5}$$

where the +1 term ensures well-defined behaviour during low-volume periods.

Then the primary outcome variable is the one-day log volume growth:

$$\Delta \log V_d = \log V_d - \log V_{d-1}. \tag{6}$$

This transformation removes persistent level effects and ensures stationarity under broad conditions. However, minute-level political shocks may not map linearly into log volume growth due to predictable components such as weekday seasonality and open-interest variation. To address this, we apply an OLS-based residualisation procedure.

For each day d, the following regression is estimated:

$$\Delta \log V_d = \alpha + \beta \cdot OIm_d + \sum_{k=1}^{4} \gamma_k \cdot DOW_{k,d} + \varepsilon_d. \tag{7}$$

where:

- OIm_d is a 5-day rolling average of open interest,
- $DOW_{k,d}$ are day-of-week dummies (Monday omitted).

The residual ε_d represents the unexpected component of volume growth, purged of deterministic weekday patterns and structural shifts in open interest. This residual serves as the primary market-response variable in the VAR.

6.3 Trump Impact Factors

The raw Trump Impact Factor scores are derived from the impact-probability values described in Section 5, displayed in Figure 5. Section 5 describes the construction of the daily impact-probability shock, defined as the 60-day rolling-Z standardised sum of predicted impact probabilities. The standardisation ensures that the measure captures deviations from local historical norms rather than mechanical increases associated with periods of higher posting volume. Crucially, these impact probabilities are obtained from a classifier trained exclusively on Q1-Q2 data, when the market was known to react strongly to political Truth Social posts. As a result, the shock series reflects the counterfactual impact that posts would be expected to have under pre-fatigue market sensitivity. The shock therefore represents informational pressure, not fatigue itself; fatigue is identified through the attenuation of market responses to these shocks.

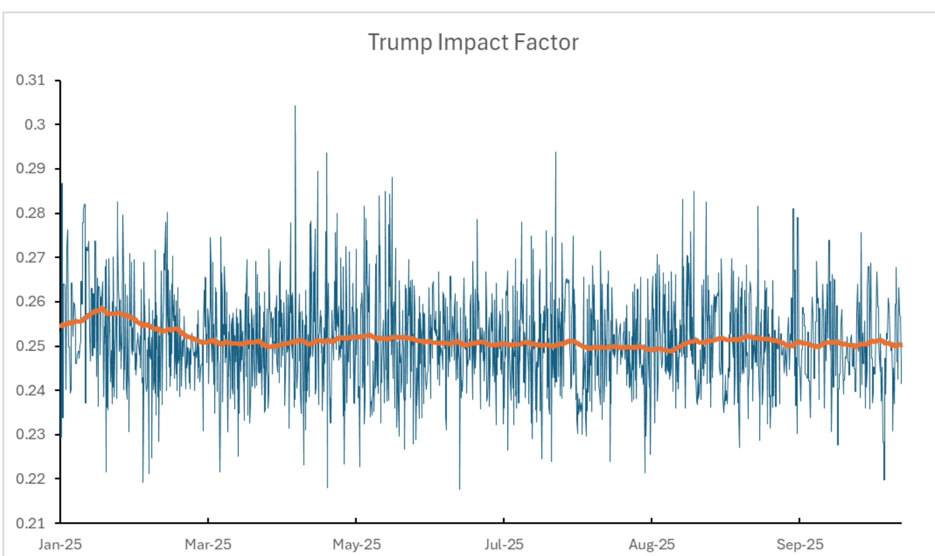

Figure 5: Trump Impact Factor scores

6.4 Identification

6.4.1 Recursive ordering and structural interpretation

We adopt a recursive identification scheme, ordering headline fatigue first and volume growth second:

$$Z_d \to \varepsilon_d^V.$$

Under this ordering, a structural innovation to Z_d is permitted to affect trading volume contemporaneously on day d.

The reverse contemporaneous channel—from volume to the shock variable—is ruled out. This structure reflects the informational timing inherent in the data: the full textual content of a political post (and the corresponding impact-probability assigned to it by the classifier) is fixed at the moment of publication, whereas market activity may respond immediately.

Formally, ordering Z_d first implies that the contemporaneous impact matrix is lower triangular, allowing the structural shocks to be obtained from the Cholesky factorisation of the reduced-form covariance matrix Σ_u.

6.4.2 Counterfactuals

An essential feature of the identification strategy is that the impact-probability model underlying is trained exclusively on Q1-Q2, a period in which the crude-oil market demonstrably responded to political Truth Social posts. The estimated parameters therefore encode the linguistic and rhetorical features that were genuinely price-relevant during the pre-fatigue regime. When these parameters are applied to Q3 posts, the resulting impact probabilities represent a counterfactual structural shock series:

$$\widehat{HF}_t = \Pr\left(\text{impact} \mid \text{text}_t; \theta_{\text{Q1+Q2}}\right). \tag{8}$$

This shock series reflects the level of impact that Q3 posts would have been expected to generate if the market had continued reacting with Q1-Q2 sensitivity, even though actual market sensitivity may have declined due to fatigue or crowding-out effects. Aggregating post-level estimates into daily shocks yields a predetermined exogenous series that is unaffected by contemporaneous or future market behaviour. Importantly, because the shock is constructed using pre-fatigue parameters, any weakening in the VAR-estimated transmission channel from Figure 7 to Figure 8 is interpreted as a reduction in the market's behavioural response, not a change in textual informativeness. This provides a transparent and empirically grounded identification of headline fatigue as attenuation in the information-to-volume transmission mechanism. This provides a transparent and empirically grounded identification of headline fatigue as attenuation in the information-to-volume transmission mechanism, in line with recent work that constructs text-based structural shocks outside the VAR using machine-learning methods and then studies their dynamic effects (e.g. Aruoba and Drechsel, 2022; Shapiro et al., 2022).

6.5 Impulse-Responses and Headline Fatigue

A positive structural shock to Z_d corresponds to an unexpected rise in predicted market-moving content from political communication, as inferred from the pre-fatigue (Q1–Q2) regime. The resulting IRF traces how this informational surprise propagates through trading activity over subsequent days. A significant and persistent response indicates that impact probability has systematic predictive power for crude-oil market participation. The sign and magnitude of the IRF reveal whether political noise crowds out market attention or stimulates speculative trading, while attenuation of the response across quarters provides evidence of headline fatigue.

Figures 6–8 report quarter-specific structural impulse response functions from the recursively identified SVAR(3), tracing the effect of a one–standard deviation informational shock on residualised daily trading volume. The preferred specification uses the model-implied impact probability as the shock variable. As discussed in the methodology, these impact probabilities are obtained from a classifier estimated on Q1–Q2 and then applied out of sample, so that the series $\widehat{HF}_t = \Pr(\text{impact} \mid \text{text}_t; \theta_{\text{Q1+Q2}})$ measures how impactful each post would have been under the pre-fatigue regime. Holding the textual mapping fixed means that cross-quarter changes in the impulse responses naturally reflect changes in the market's behavioural reaction to a given level of textual informativeness.

The impact-probability IRFs, as computed in Appendix B, show a clear progression across the year. In Q1, a structural impact-probability shock generates a small but positive contemporaneous increase in $\Delta \log V$, followed by a smooth decline (Figure 6). The point estimates remain modest, peaking between 0.003 and 0.006 before dissipating within a week. Confidence intervals are wide and overlap zero at all horizons, indicating that the effect is economically small and statistically imprecise early in the sample. In Q2, the market response strengthens markedly (Figure 7). The contemporaneous effect rises to roughly 0.008–0.009, and the IRF displays a mild hump shape during the first few days before converging back toward baseline. Short-horizon confidence intervals lie mostly above zero, suggesting that high-impact posts still generated incremental liquidity in Q2. By Q3, however, the response has largely disappeared (Figure 8). The IRF is essentially flat at all horizons, and the associated confidence intervals comfortably encompass zero. Even though the classifier's scoring rule is unchanged, structural impact shocks in Q3 no longer translate into systematic adjustments in trading volume.

Re-estimating the SVAR with alternative classifier-based shock measures, such as the expected price move or the predicted magnitude of the impact, produces a consistent pattern: a weak and noisy response in Q1, a clearer short-lived effect in Q2, and negligible responses in Q3. While the magnitudes differ slightly across shock definitions, the cross-quarter ranking is unchanged. This robustness indicates that the attenuation is not an artefact of the chosen sentiment proxy but reflects a genuine weakening of the information-to-volume transmission mechanism.

The SVAR results provide coherent evidence of headline fatigue in the crude-oil market. Under a fixed text-to-impact mapping learned from the pre-fatigue regime, both high-impact posts and repetitive clusters of posts are associated with meaningful increases in trading volume early in the year. As time passes, the market's responsiveness becomes more selective and ultimately fades out. In Q2, only posts that the classifier scores as highly impactful still generate a noticeable but transient rise in liquidity, while the effect of repetition is already muted. By Q3, even counterfactually high impact-probability shocks, $\widehat{HF}_t$, fail to move volume in a systematic way. Because the underlying shock processes are based on the fixed parameter vector $\theta_{\mathrm{Q1+Q2}}$, the observed attenuation reflects a decline in traders' behavioural responsiveness rather than changes in textual informativeness. This pattern constitutes direct empirical evidence of headline fatigue.

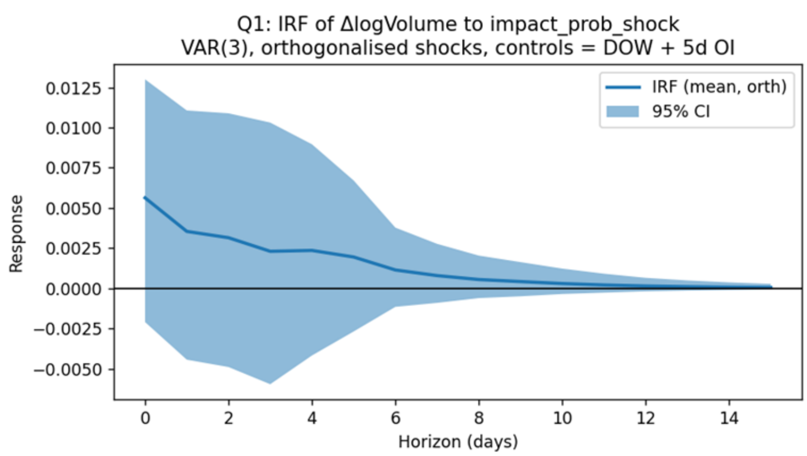

Figure 6: Impulse response function residualised volumes following a unit standard deviation shock to the Trump Impact Factor, during Q1

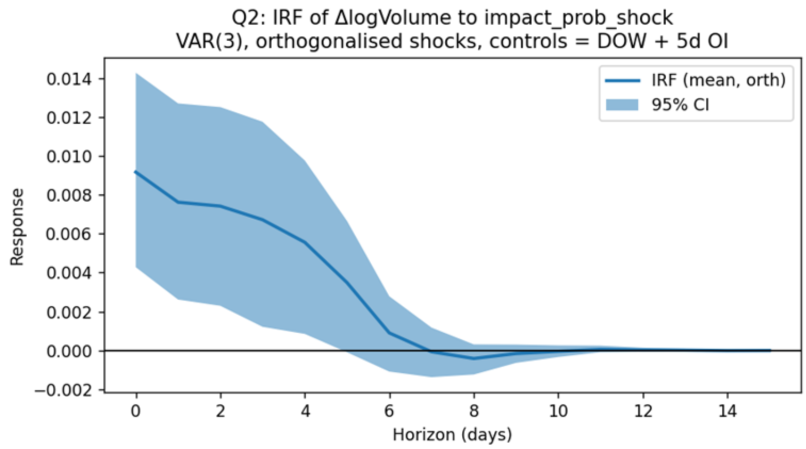

Figure 7: Impulse response function residualised volumes following a unit standard deviation shock to the Trump Impact Factor, during Q2

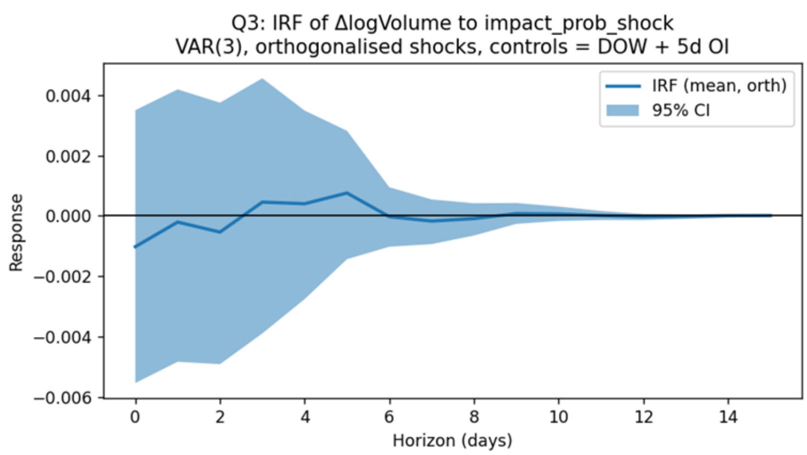

Figure 8: Impulse response function residualised volumes following a unit standard deviation shock to the Trump Impact Factor, during Q3

7 Dynamics in open interest

7.1 Trading Intensity

In this section we examine why trading activity in Brent futures slowed in Q3 2025 despite open interest remaining broadly unchanged, at -0.055% quarter-on-quarter, and to what extent money managers' behaviour – a proxy for speculative interest – contributed to that shift Commodity Futures Trading Commission (2025). The focus is placed on how trading intensity evolved over Q3, how aggregate (both long and short) managed-money positioning changed, the weekly volatility trend and how these patterns relate to the underlying crude environment. For the study we utilise the Brent front-month futures contract as a good representation of the overall state of the market: more specifically, front-month Brent futures volumes are on average around 10% of the global exchanged traded oil derivatives market. In Q3 2025 Brent front month volumes averaged 322.3 million barrels per day, while the global exchanged traded oil derivatives volumes were approximately 3.5 billion barrels per day.

To do this, we provide an indicator of trading intensity, which is equivalent to global exchange traded volumes over global open interest. This ratio is more informative than raw traded volumes in isolation as it allows us to capture the share of contracts traded with respect to the total market risk. Illustrating the trading intensity over the year indicates a clear regime change: in Q1 and Q2, trading intensity fluctuates around a relatively higher mean with more frequent swings, consistent with an environment of elevated trading activity and more active short-term position taking. On the other hand, in Q3 the series settles into a visibly lower range, with less frequent and less material spikes. In other words, the market maintained broadly unchanged levels of open risk, but turnover in those positions became less frequent.

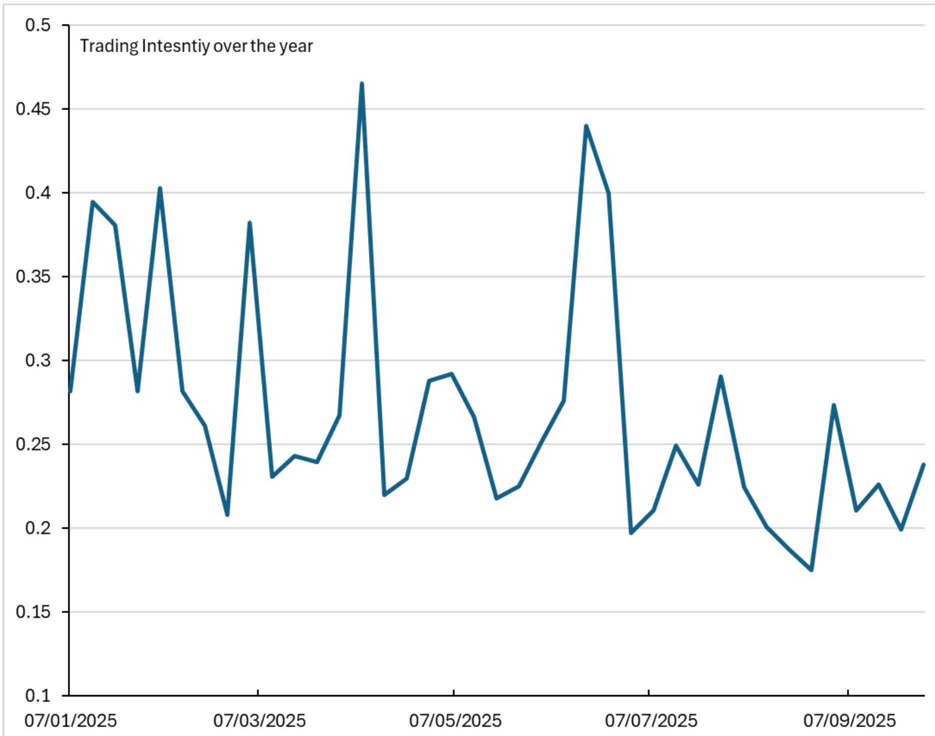

Figure 9: Trading Intensity in 2025

In addition, we examined money managers' aggregate positioning over the same period. Aggregate positioning is used instead of net positioning as the purpose of the study is to understand the volume of activity undertaken by funds as opposed to the direction of positioning. In order to approximate the weekly rebalancing in speculative positioning, we take the weekly change in the aggregate money manager positioning. Weekly change is defined as MM intensity:

$$MM_t^{\text{intensity}} = \left(\frac{MM_t - MM_{t-1}}{MM_{t-1}} \right) \times 100. \tag{9}$$

When we compare the speculative intensity across the year, a clear pattern emerges: both the sample mean and standard deviation of our variable are higher in the first half of the year and noticeably lower in Q3. In H1, funds were liquidating or adjusting their positions at relatively far higher rates, whereas in Q3 weekly changes showed weaker and less frequent pattern movements – further suggesting the assumption that Q3 was a low trading regime when compared to the first two quarters of the year.

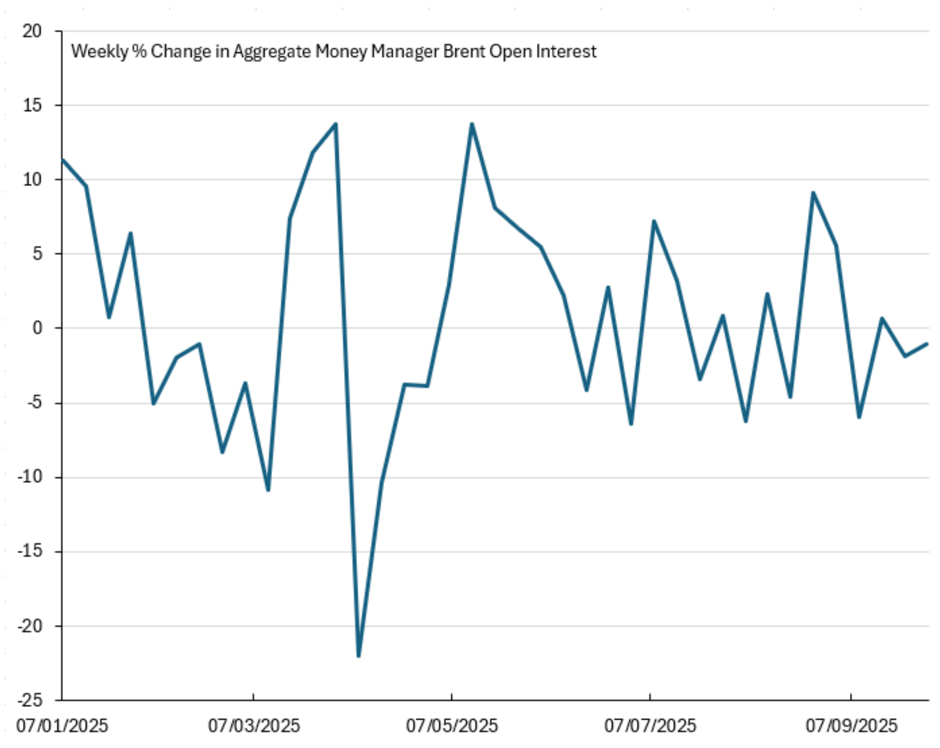

Figure 10: Money Managers Aggregate Brent positioning

The behavioural shift coincides with a backdrop of calmer Brent flat-price dynamics. Weekly front-month Brent futures returns in Q3 show smaller absolute moves than in H1 2025 and flat price traded within a significantly narrower range. Mechanically, larger price swings tend to be associated with elevated trading activity, and, conversely, stronger or more frequent adjustments in traders' positions would normally imply greater realised price volatility. Any attempt to attribute the decline in trading intensity to changes in speculative behaviour must therefore control explicitly for the volatility environment.

This motivates the empirical exercise that follows. The objective is not merely to show that lower volatility coincided with lower trading activity, but to assess whether money-manager behaviour itself contributed to the Q3 low-intensity regime, once volatility and curve structure are taken into account. We therefore estimate the following first-stage regression:

$$TI_t = \alpha + \beta_1 \left|\Delta \log(\text{Brent}_t)\right| + \beta_2 \Delta \log(\text{relCZ25–CZ26}_t) + u_t. \tag{10}$$

Where $|\Delta \log(\text{Brent}_t)|$ denotes the absolute weekly log return of front month Brent futures, as we are interested to regress against the magnitude of the move rather than the direction. While $\Delta \log(\text{relCZ25–CZ26}_t)$ represents the curve dynamics, more specifically, the weekly change of the log price of December 2025 minus the log price of the December 2026 Brent futures contract. By taking this spread we endeavour to incorporate the market signal of fundamental tightness or looseness a year apart in the Brent curve. Because money-manager behaviour is strongly shaped by curve structure – via expectations, roll yield and storage economics – omitting it would risk confusing structural shifts with genuine changes in trading intensity. Adding the spread, therefore, helps distinguish investor behaviour from broader fundamentals.

Despite the usefulness of adding the Dec/Dec spread, a common question one would raise is multi-collinearity. This concern is particularly relevant in our setting as one could argue that weekly moves in the prompt Brent futures contract would also affect the Dec/Dec spread.

To formally test for this, variance inflation factors (VIFs) are computed for the two regressors. To quantify multicollinearity between these two regressors, VIF is computed for the Dec/Dec term. By definition, the VIF for a regressor X_j is

$$VIF_j = \frac{1}{1 - R_j^2}. \tag{11}$$

Where R_j^2 is the coefficient of determination from the auxiliary regression of X_j on all the other regressors. In this case, the relevant auxiliary regression is:

$$\Delta \log(\text{relDecDec}_t) = \gamma_0 + \gamma_1 \, |\Delta \log(\text{Brent}_t)| + v_t. \tag{12}$$

and the resulting VIF for the Dec/Dec term is approximately

$$VIF_{\Delta \log(\text{relDecDec})} = 1.03. \tag{13}$$

$$R^2 = 1 - \frac{1}{VIF} = 1 - \frac{1}{1.03} \approx 0.03. \tag{14}$$

The results imply that only about 3% of the variance of the Dec25/Dec26 spread can be explained by the Brent returns and that 97% of its variance is independent. Econometrically, a VIF so close to one indicates that multicollinearity between the flat price volatility term and the Dec/Dec spread change is negligible: the curve term is not a near-linear function of the Brent term in this sample and does not inflate the standard errors in a way that would undermine the estimates.

The estimates of the first regression show that – as expected – outright Brent volatility has a strong, positive effect on trading intensity, while it is also the dominant driver, as weeks with larger absolute price swings exhibit significantly higher trading activity. By contrast, the Dec25/Dec26 spread change is not statistically significant in this sample, suggesting that, at least over this period, curve reshaping in that part of the strip did not systematically drive flows.

7.2 Money Manager behaviour

However, the objective of the empirical exercise is not only to document that lower volatility coincided with lower trading intensity, but also to assess whether there is more to the story and whether changes in money-manager behaviour contributed to a distinct low-activity regime in Q3. To address this, we then regress the residual of the first regression. The residual is computed as:

$$TI_t^{\text{resid}} = TI_t - \widehat{TI}_t. \tag{15}$$

Where TI is the trading intensity at week t and $\widehat{TI}$ is the fitted value from the Stage-1.

The residual captures the component of trading intensity that cannot be explained by the contemporaneous flat-price and curve changes included in the first specification. This residual is interpreted as a measure of excess or behavioural activity: it reflects variation in turnover that is more likely to be associated with changes in participation, trading style and other factors beyond the immediate magnitude of price moves. Money manager variables are introduced only in the second stage. The rationale is to avoid confounding the strong and well-documented effect of outright Brent futures volatility on activity with the more subtle behavioural effects the study seeks to isolate.

The second-stage regression takes the form:

$$TI_t^{\text{resid}} = \delta_0 + \delta_1 \, MM_t^{\text{intensity}} + \delta_2 \, D_t^{Q3} + \delta_3 \left(MM_t^{\text{intensity}} \cdot D_t^{Q3} \right) + \varepsilon_t. \tag{16}$$

Where $MM_{\text{intensity}}$ is the weekly percentage change aggregate managed-money length (the speculative activity proxy), D^{Q3} is a dummy that equals one for Q3 observations and zero otherwise, and the interaction term allows the sensitivity of residual trading intensity to money-manager activity to differ in Q3 relative to the rest of the sample.

The estimated coefficients from this second-stage regression indicate that the Q3 dummy enters with a negative and statistically significant coefficient, implying that Q3 weeks exhibit lower-than-expected trading activity even after controlling for the size of price moves and curve dynamics. By contrast, the money-manager intensity term is small and imprecisely estimated, and the interaction with Q3 does not materially alter this conclusion. In other words, once the mechanical effect of volatility has been taken into account, week-to-week speculative churn does not provide a strong additional explanation for the decline in trading intensity, although Q3 still stands out as a distinct low-activity regime.

Additional robustness checks reach similar conclusions. An amplifier regression that allows the effect of Brent volatility on trading intensity to differ between high- and low-intensity money-manager weeks finds that the interaction between volatility and a high-intensity dummy is statistically insignificant. Weeks with elevated money-manager churn do not exhibit a systematically stronger sensitivity of trading intensity to flat-price moves. Simple splits of the sample into high- and low-intensity weeks outside Q3 likewise show no significant differences in the mean or variance of trading intensity between the two groups.

Taken together, the evidence points to a clear change in money-manager behaviour in Q3, but not in the form of an abrupt exit from the market or a spike in speculative activity. Aggregate money-manager open interest remained high, yet the average size and variability of weekly rebalancing declined, consistent with a shift from a more frequent position to a more set-and-hold, less short-term focused stance. At the same time, the econometric results indicate that the dominant drivers of the Q3 decline in trading intensity are the calmer flat-price environment and the emergence of a distinct low-activity regime, rather than a strong, independent causal effect of week-to-week money-manager churn. The role of money managers in this episode is therefore best characterised as a change in trading behaviour – towards lower repositioning and greater willingness to sit on existing risk – rather than as the primary quantitative driver of the fall in observed turnover.

8 Case Study: Dubai

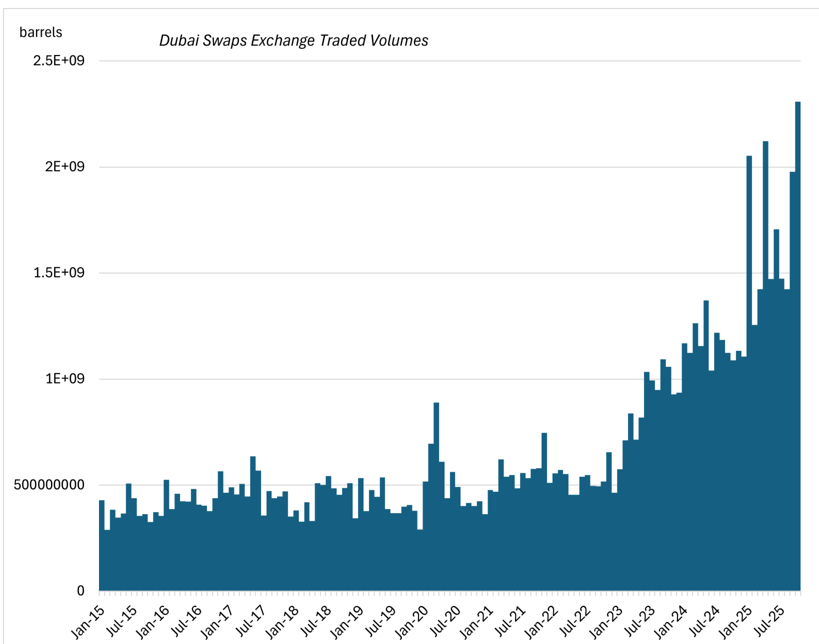

Figure 11: Monthly Exchange Traded Volumes of Dubai Swaps

The Dubai market in Q3 was particularly characterised by Mercuria's activity in the window, culminating in the busiest month of spot trading on record. In August, trading sources told The Officials Mercuria had accumulated near to 30 cargoes of Upper Zakum it intended to sell in the Dubai window. This was vindicated in September, when the trade house took a major short position in partials trading, offering and hitting bids extensively. Through the month's trading, they sold 45 cargoes, of which 38 were Upper Zakum nominations. To demonstrate Mercuria's dominance on the sellside, there were only 4 convergences declared by any other seller throughout the entire month, bringing the month's total tally to 49, an outright record, with 42 of those being Upper Zakums. According to calculations by The Officials based on market sources' information about Upper Zakum output, this equated to 75% of the field's production programme for the month: 45 cargoes of 500,000 barrels each is 750,000 barrels per day. Sources told The Officials Upper Zakum production was approximately 1 million barrels per day, meaning the volume of Upper Zakum crude declared in the window was three quarters that of total field production.

Typically this kind of window play will be accompanied by extremely sizable paper positions, as traders seek to make the most of a move in the physical market in the respective derivatives, often by placing a much larger notional value into paper positions than the physical. However, these derivative positions would likely be entered significantly before the physical trading of the relevant month began, though it is unlikely a single play of this kind would make a material and observable difference in a full month's traded volumes.

Indeed, in July and August, traded volumes of Dubai swaps showed no such anomalous signs, remaining relatively depressed at 1.473 billion and 1.423 billion barrels, respectively. The impact, instead, was felt in the month of September, as the physical window activity reached its crescendo.

Naturally, it is difficult to isolate this effect. However, the specific coincidence of the most declared cargoes in history in the Dubai window and extremely high Dubai swap trading volumes, absent a broader market increase in traded volumes, suggests at least a strong correlation.

The contrast of Dubai swaps to the overall change in traded volumes across instruments and products between Q2 and Q3 is striking. Dubai swaps traded volumes dropped 8.1% q/q, while total traded volumes fell much more, down almost 25%. The much lesser fall in Dubai swaps traded volumes indicates a dislocation of the instrument from the broader market, aided in particular by the extremely high volumes seen in September. Indeed, September volumes were 36.5% above the average of July and August volumes in Dubai swaps, while in the overall market, September volumes were 1.7% lower than the average of July and August. That is to say the strength in Dubai trading volumes was concentrated exclusively in extremely hectic September trading.

In terms of the development of liquidity in Dubai swaps, there is a distinct acceleration from the beginning of 2023, where volumes begin to pick up markedly. From January 2015 until December 2022, volumes had been consistently below 500 million barrels per month, with a few notable exceptions such as March 2020, in line with the overall uptick in volumes across the market due to the Covid-19 pandemic and global lockdowns. The overall increase was steady but strong, as the growth in traded volumes from 2015 to 2022 was 39%, far in excess of the overall market's growth of just over 15%. As a percentage of total traded volumes, in the same time period, Dubai swaps have grown from 0.071% to 0.857%.

Since the start of January 2023, however, the growth has been stratospheric. From 2022 to 2023, volumes traded in Dubai swaps surged by 68.9%, compounded by a further jump of 31.2% in 2024.

The outsized growth in Dubai swap volumes illustrates the growing significance of Dubai as a global benchmark, rising towards competing with the established Western benchmarks like Brent and WTI, although the volumes of the latter two (futures) contracts are far in excess of Dubai swaps.

This trend looks primed to continue, however, as Dubai's significance as one of the three largest global benchmarks is consolidated; as has been the case for some time, Brent is undermined by concerns over physical North Sea liquidity and WTI is exposed to US policy decisions. As a benchmark, Dubai also has the advantage of being a medium sour grade, in contrast to the light sweet Brent and WTI, an increasingly important proposition, providing greater granularity in trading and hedging, as liquidity deepens.

This development gathered steam in October, as traded volumes of Dubai swaps reached a record high of 2.308 billion barrels, surpassing the previous high in April 2025. That value was more than 185 million barrels higher than the April figure, or 8.7%. Unlike in April, when Trump's Liberation Day tariffs sparked the second highest month of traded volumes on record, there was no commensurate overall increase in traded volumes in October. Indeed, global traded volumes were actually 9.6% lower in October than in April.

9 Forecast

9.1 Methodology

Our forecast methodology is largely unchanged since the last iteration of this report, so for our diligent readers, feel free to skip to the assumptions. For new readers however, we now present our framework guiding our view of the future evolution of exchange-traded volumes in oil derivatives, both exchange-traded futures and over-the-counter swaps. We collate volumes data across an exhaustive set of exchanges, products, and tenors. From this, we derive an estimate for total globally traded volumes.

We then take an empirical approach, which captures both speculative and hedging components. Hedging open interest is fairly straightforward to capture, being a function of physically traded barrels and, therefore, a function of aggregate economic activity. Speculative interest is more difficult. To measure broad financial market liquidity, we turn to the US 2s10s. The 2s10s captures the steepness of the Treasury curve in the US. There are other metrics, but for our purposes, we found this to be most performant. The inclusion of the 2s10s yield curve slope in a model forecasting traded volumes of oil derivatives is justified by its ability to capture market expectations of macroeconomic conditions and financial cycles.

The slope of the yield curve, particularly the spread between the 2-year and 10-year Treasury yields, encapsulates market expectations of future economic growth, inflation, and monetary policy adjustments. These macroeconomic expectations are critically linked to risk-taking behaviour of financial market participants, including commodity speculators. Periods of curve steepening, often 16 associated with anticipated monetary easing or improved growth prospects, tend to coincide with enhanced risk appetite and increased speculative trading in higher-beta assets such as oil derivatives.

Conversely, yield curve inversions have been shown to precede economic slowdowns and are typically accompanied by tighter financial conditions and reduced speculative activity. Furthermore, the yield curve indirectly reflects shifts in liquidity conditions and credit availability, both of which are foundational to leveraged speculative positioning in derivatives markets. Incorporating the 2s10s measure thus provides a macro-financial dimension to the model, capturing forward-looking signals that are not readily observable through spot liquidity indicators alone. This is consistent with the broader literature linking term structure metrics to financial conditions and asset market dynamics (Estrella and Trubin, 2006).

We employ a time-series framework, specifically an autoregressive distributed lag model, to capture time-varying dynamics. We also employ dummy variables to capture seasonality.

9.2 Forecast Assumptions

We present our forecast in Figure **??**. The upward trend in traded volumes observed in recent quarters is projected to persist over the forecast horizon. Although the current tariff regime is expected to weigh on aggregate activity in the near term, we view the scope for a sustained escalation as limited. During Q3 the effective tariff on US imports continued to drift higher, from roughly 16% in Q2 to around 17% at the time of writing. However, the political economy constraints facing the administration make a prolonged policy of aggressive protectionism unlikely.

In particular, President Trump has strong incentives to avoid both equity market devaluation and a disorderly sell-off in sovereign debt. Equity price declines erode household wealth and, in the post-pandemic era of expanded retail participation, a growing share of the electorate is directly exposed to equity markets through brokerage accounts and retirement vehicles. Any substantial drawdown would therefore carry a high political cost. The fiscal arithmetic of tariffs is also unfavourable. While customs duties are directly revenue-generating, the associated demand destruction offsets much of this gain. Despite the elevated effective tariff rate, customs duties accounted for only about 3.7% of total federal tax revenues in September 2025, implying limited capacity to stabilise the fiscal position.

Macro–fiscal dynamics reinforce these constraints. Labour-market conditions have softened appreciably through Q3 2025, with rising job losses counteracting any consolidation achieved on the revenue side. The earlier "Big Beautiful Bill" introduced substantial tax cuts, further widening the structural deficit. In addition, the federal government shutdown that followed Q3 is likely to depress near-term growth. These developments have prompted a rebalancing of priorities at the Federal Reserve: downside risks to employment now dominate upside risks to inflation. Consistent with this shift, the Federal Funds Rate has been reduced to 3.75–4.00%. Term yields have moved lower, although they remain high relative to the pre-COVID period. Taken together, these factors suggest that tariff policy will remain noisy, with periods of sharp tightening and partial reversal, but is ultimately likely to retrace toward lower levels by the middle of next year. Episodes of heightened policy noise have historically coincided with evidence of market fatigue, which we expect to exert a modest drag on traded volumes throughout the profile.

The United States is not the only driver of our global growth outlook. China remains an important swing factor for the global growth outlook. Growth momentum softened through the year, with the economy expanding by 5.4% in Q1, 5.2% in Q2 and 4.8% in Q3, even though full-year guidance still implies a 5% growth rate for 2025. Again, as we have seen, weaker data prints are policy-positive for China, which has shown readiness to counter both producer and consumer softness with additional fiscal support. For example, producer prices have been in deflation for 37 consecutive months and policymakers have stepped up targeted measures in response.

Despite that, debt concerns are mounting globally, and so are in China, which has announced to support a backlog of unpaid bills by its local governments estimated to be over $1 trillion according to reports. With that being said the yield on Chinese 10-year bonds remains amongst the lowest in the world – trading for majority of the quarter below 1.90%, keeping borrowing costs cheap and supporting credit conditions. At the same time, China has been opening its bond market further to overseas participation, offering an alternative to dollar exposure at a time when trade tensions and tariff uncertainty remain elevated. Concluding, our view on China remains positive as the authorities have demonstrated a consistent willingness to cushion downside surprises in order to meet growth targets, accompanied by growing overseas investor appetite locally.

On this basis we project a material acceleration in global economic output next year, with growth strengthening further into 2027. Higher realised and expected physical oil demand should translate into increased hedging needs from producers and consumers, thereby supporting the expansion of futures and options volumes. Stronger aggregate growth also tends to relax financial constraints, allowing for additional monetary easing, improved market liquidity and a more constructive risk appetite, all of which are conducive to higher trading activity.

Geopolitical developments constitute an important source of uncertainty around this central scenario. While recent diplomatic efforts have raised hopes of de-escalation in several theatres, our baseline view is that the progress achieved to date will prove fragile. The war in Ukraine continues with limited signs of durable movement toward a ceasefire; although the United States has stepped up its efforts to broker an agreement, the terms currently under discussion still present substantial roadblocks to a lasting settlement. At the same time, tensions between China and Japan have intensified following comments by Takaichi regarding Taiwan, contributing to heightened trade frictions within Asia. These geopolitical frictions pose downside risks to global trade and confidence, even if a large, discrete shock is not incorporated into the baseline.

Despite these risks, we expect overall financial conditions to remain broadly accommodative. While the future path of US monetary policy is uncertain, we take the view that the Federal Reserve will remain more attentive to labour-market outcomes than to a rapid return of inflation to target, and will therefore tolerate modestly above-target inflation for some time. Under this assumption we anticipate a continued steepening of the US Treasury yield curve as policy rates are gradually reduced, while longer-dated breakeven inflation rates remain elevated.

In combination, these considerations underpin our central projection that traded crude-oil volumes will continue to trend higher over the next two years. However, the interaction of tariff-induced noise, labour-market softening and lingering geopolitical risks warrants a moderate downward revision relative to our previous forecast profile.

9.3 Conclusion and volumes in the years ahead

We conclude that the sharp decline in traded volumes in Q3 2025 is best viewed as a cyclical adjustment within an ongoing structural expansion of the oil derivatives market, rather than a turning point. Volumes fell to just over 232 billion barrels, their lowest quarterly level since late 2024, yet open interest remained broadly stable, only 0.055% lower than in Q2. This combination of weaker turnover and resilient open positions, together with the decline in trading intensity, points to a regime shift in trading behaviour: market participants have maintained risk but are turning over that risk less frequently.

The analysis of Trump-related informational shocks reinforces this interpretation. Using a fixed impact–probability model estimated on the pre-fatigue regime, we show that politically salient posts generated economically meaningful increases in residualised volumes in Q1 and, especially, Q2, but their effect largely vanished by Q3. The transmission channel from "impactful" content to trading activity weakens over time, even though the textual mapping is held constant. Similar patterns emerge for shocks based on expected price impact and novelty, while the response to information clusters (max similarity) is strongest in Q1 and decays thereafter. Taken together with the evidence on money-manager behaviour and the Dubai case study, this provides a coherent picture of headline fatigue: the market remains deep and liquid, but has become more selective about which events warrant incremental trading.

Against this backdrop, our baseline forecast envisages a resumption of volume growth, but from a lower base and with more weight on macro-fundamental drivers than on episodic geopolitical shocks. The forecast is conditioned on a noisy but ultimately constrained US tariff path, a more labour-market-focused Federal Reserve willing to tolerate above-target inflation, and a Chinese policy mix that remains firmly growth-supportive. We also assume that the Russia–Ukraine conflict persists without a decisive resolution and that tensions in Asia remain an intermittent, rather than catastrophic, source of risk premia. In such an environment, volumes are expected to be driven primarily by the evolution of global balances, the term structure of prices, and the availability of speculative capital, rather than by single-headline events.

On a quarterly basis, we project that traded volumes recover to around 261 billion barrels in Q4 2025, partially reversing the Q3 slump. For 2026, we forecast quarterly traded volumes of approximately 259, 295, 298 and 306 billion barrels in Q1–Q4, respectively, implying a full-year total of about 1.18 trillion barrels and a 11.05% increase relative to 2025. The expansion is expected to accelerate in 2027, with quarterly volumes of roughly 331, 346, 355 and 365 billion barrels, bringing the annual total to about 1.38 trillion barrels and a year-on-year growth rate of around 16.8%. Open interest is projected to grow broadly in line with this trajectory, rising by a little over 8% in 2026 and a further 7% in 2027, consistent with a gradual deepening of the market's risk-bearing capacity.

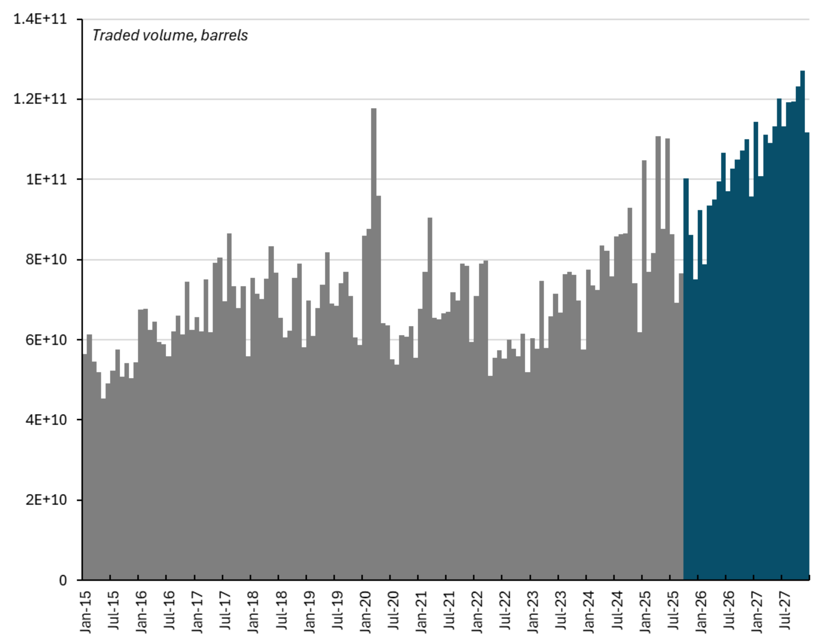

Figure 12: Total oil derivatives traded volumes as assessed and forecast by The Officials

Overall, we therefore judge that Q3 2025 marks a period of temporary fatigue and low trading intensity, rather than a structural break in the development of global oil derivatives markets. Our central view is that traded volumes will continue to trend higher over the next two years, supported by recovering global growth, accommodative financial conditions and a steady rebuild in risk appetite, albeit from a weaker starting point and with a modest downward revision relative to the profile envisaged in the previous Quarterly Volumes Review.

Quarterly Table

The Officials global volumes and open interest historical assessments and forecast (millions of barrels)

	Sum of Volumes	Average of Open Interest
2021		
Qtr1	235,010	14,330
Qtr2	197,080	14,450
Qtr3	208,680	13,850
Qtr4	216,980	13,090
2022		
Qtr1	229,790	12,660
Qtr2	163,960	11,510
Qtr3	173,030	10,870
Qtr4	169,350	10,740
2023		
Qtr1	192,890	11,530
Qtr2	195,470	12,580
Qtr3	220,030	13,150
Qtr4	203,630	13,000
2024		
Qtr1	223,520	13,660
Qtr2	241,440	14,440
Qtr3	258,650	14,530
Qtr4	229,100	14,850
2025		
Qtr1	263,470	15,480
Qtr2	308,780	16,690
Qtr3	232,170	16,680
Qtr4	***261,450***	***17,110***
2026		
Qtr1	***259,390***	***16,940***
Qtr2	***295,000***	***17,610***
Qtr3	***298,550***	***18,210***
Qtr4	***306,390***	***18,580***
2027		
Qtr1	***331,060***	***18,620***
Qtr2	***346,080***	***18,890***
Qtr3	***355,390***	***19,230***
Qtr4	***365,530***	***19,670***

Note: Forecast figures shown in bold italics.

Monthly Table

The Officials global volumes and open interest historical assessments and forecast (millions of barrels)

	Sum of Volumes	Average of Open Interest
2025		
Jan	104,840	15,210
Feb	76,950	15,460
Mar	81,690	15,770
Apr	110,820	16,630
May	87,680	16,820
Jun	110,290	16,630
Jul	86,420	16,700
Aug	69,200	16,500
Sep	76,550	16,840
Oct	100,230	17,370
Nov	***86,130***	***17,110***
Dec	***75,090***	***16,850***
2026		
Jan	***92,430***	***16,880***
Feb	***78,900***	***16,960***
Mar	***93,500***	***16,990***
Apr	***95,050***	***17,220***
May	***99,490***	***17,690***
Jun	***106,650***	***17,920***
Jul	***97,070***	***18,030***
Aug	***102,790***	***18,250***
Sep	***104,950***	***18,360***
Oct	***107,140***	***18,440***
Nov	***110,010***	***18,610***
Dec	***95,680***	***18,690***
2027		
Jan	***114,310***	***18,710***
Feb	***100,860***	***18,560***
Mar	***111,200***	***18,580***
Apr	***109,090***	***18,690***
May	***113,290***	***18,930***
Jun	***120,240***	***19,050***
Jul	***113,140***	***19,110***
Aug	***119,200***	***19,250***
Sep	***119,490***	***19,320***
Oct	***123,140***	***19,450***
Nov	***127,090***	***19,710***
Dec	***111,650***	***19,840***

Note: Forecast figures shown in bold italics.

References

Aruoba, S. B. and Drechsel, T. (2022). Identifying monetary policy shocks: A machine learning approach. Preprint.

Bollen, J., Mao, H., and Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8.

Chen, R. R., Xiao, X., and Chen, K. (2023). Presidential attack risk: The impact of Trump's tweets on financial markets. *Journal of Global Information Management*, 31(1):1–26.

Commodity Futures Trading Commission (2025). Commitments of traders (COT) reports. `https://www.cftc.gov/MarketReports/CommitmentsofTraders/index.htm`. Accessed 27 November 2025.

Estrella, A. and Trubin, M. R. (2006). The yield curve as a leading indicator: Some practical issues. *Current Issues in Economics and Finance*, 12(5).

International Energy Agency (2025). Oil market report – august 2025. Technical report, IEA, Paris.

Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics*, 126(1):1–50.

Shapiro, A. H., Sudhof, M., and Wilson, D. J. (2022). Measuring news sentiment. *Journal of Econometrics*, 228(2):221–243.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62(3):1139–1168.

Tetlock, P. C. (2011). All the news that's fit to reprint: Do investors react to stale information? *Review of Financial Studies*, 24(5):1481–1512.

A Return Metrics

This appendix outlines the return-based measures used to quantify short-horizon market reactions to President Donald Trump's Truth Social posts.

A.1 Definition of Metrics

The identifier system computes three related measures of intra-minute return dynamics:

1. **Net five-minute return**
   The principal outcome variable is the log return from the minute prior to the post to the end of the five-minute post-event window:

$$r_i(0-5) = \ln(p_{i,t_i+4}) - \ln(p_{i,t_i-1}). \tag{17}$$

   This captures the immediate directional price response commonly analysed in high-frequency event-study applications.

2. **Pre-trend control**
   A pre-event log return is computed over the five minutes preceding the post:

$$\text{pre}_i(5-0) = \ln(p_{i,t_i-1}) - \ln(p_{i,t_i-5}). \tag{18}$$

   This measure captures recent price drift within the event window and serves as a conditioning variable or placebo outcome in supervised learning settings.

3. **Post-minus-pre delta return**
   The net change in event returns is defined as:

$$\Delta_i = \big(r_i(0-5) - \text{pre}_i(5-0)\big) \times 100, \tag{19}$$

   expressing the differential return in basis-point units and isolating genuine post-information movements from preceding drift.

B Computing Impulse Responses

Impulse responses are derived from the structural moving-average representation of the recursively identified SVAR.

B.1 Reduced-Form VAR

Let the estimated reduced-form VAR(3) be:

$$Y_d = A_1 Y_{d-1} + A_2 Y_{d-2} + A_3 Y_{d-3} + u_d, \tag{20}$$

where u_d denotes reduced-form innovations. Structural shocks ε_d are obtained via the Cholesky decomposition of the covariance matrix, with:

$$u_d = P\varepsilon_d, \tag{21}$$

where P is lower triangular and orders the informational shock first.

B.2 Structural Impulse Responses

The structural moving-average representation is:

$$Y_{d+h} = \Psi_h \varepsilon_d, \tag{22}$$

with the recursion:

$$\Psi_h = A_1\Psi_{h-1} + A_2\Psi_{h-2} + A_3\Psi_{h-3}, \qquad h \geq 1, \quad \Psi_0 = P. \tag{23}$$

The impulse response at horizon h corresponds to the reaction of volume growth to a one-standard-deviation structural innovation in the informational-shock equation.