

Petroleum Inventory Scarcity: One Negative Macro-Commodity Identity

Arhan Canli

Founder, System Architect, and Quantitative Researcher
Canli Capital / AlphaC Algorithms

VERSION	0.1.0
RELEASED	23 August 2026
REGISTRY KEY	energy_inventory
STATUS	Working paper preprint / not peer reviewed

Research simulation and, where explicitly identified, Alpaca paper evidence. No funded performance, peer review, independent replication, external acceptance, or future return is claimed.

Short title: Petroleum inventory scarcity

Author: Arhan Canli, Founder and Quantitative Researcher, Canli Capital

Family key: energy_inventory · **System:** ALPHAC / AlphaForge

Status: public research record; not peer reviewed; not an investment solicitation

Evidence date: 2026-08-22

Abstract

This study tests one preregistered implementation of petroleum-inventory scarcity using EIA first-release data and liquid exchange-traded energy proxies. The signal enters on the next market open after release, normalizes inventory surprises against five years of seasonal history, and hedges trailing DBC beta. Over 2,669 immutable ledger observations, annualized Sharpe was -0.5893, skew was -1.5157, and kurtosis was 30.2245. The fuller probe measured net Sharpe -0.5892, Newey-West t-statistic -1.8450, DSR effectively zero, and maximum drawdown of 41.05%. It passed exposure and correlation checks but failed the return, robustness, product-consistency, portfolio, and capacity gates. The verdict is KILL and the family contributes zero sleeves.

Economic hypothesis

Petroleum inventories connect physical scarcity to spot and futures prices. Low inventories can raise convenience yield because holding the physical commodity protects production and delivery needs. Gorton, Hayashi, and Rouwenhorst document broad relationships between commodity inventories, the basis, and risk premia ([Review of Finance](#)).

The ALPHAC hypothesis is narrower. It asks whether the first public EIA petroleum release contains incremental information that can be traded after publication, once normal seasonality and broad commodity beta are controlled. This timing distinction is essential. A relation visible in revised inventory history or contemporaneous prices is not necessarily available to a trader at the next session open.

Preregistered implementation

The single charged probe uses USO and UGA as energy return proxies and DBC as the hedge instrument. The inventory score is normalized using five seasonal years, scaled over 52 weeks, and clipped to the interval $[-3, 3]$. Positions enter at the next session open after the EIA release. DBC hedge beta is estimated over 252 trailing sessions and clamped to $[-3, 3]$. One-way costs are 6 basis points for USO, 10 for UGA, and 3 for DBC. The out-of-sample period begins in 2016.

This is one configuration. PBO is not defined because no selection surface exists inside the probe. The identity is still charged to the 228-trial global union for selection-adjusted inference.

Result

Measure	Value
Net annualized Sharpe	-0.5892
Newey-West t-statistic	-1.8450
Probabilistic Sharpe ratio	0.0248
Deflated Sharpe ratio, 228-trial union	7.37e-11
Maximum drawdown	41.05%
Annualized turnover	52.30x
Net Sharpe at twice costs	-0.9977
Realized DBC beta	-0.0137

Both standalone products were negative: USO Sharpe was -0.4603 and UGA Sharpe was -0.7426. The twice-cost result moved farther below zero. The negative left tail and high turnover make the failure economically stronger than a marginally insignificant mean.

Diversification did not rescue the return

The candidate had low ordinary correlation with the four comparison sleeves. Average correlation was 0.0012 and the largest pair was 0.0323. The largest stressed correlation was 0.0790. Those figures clear the probe's overlap thresholds, but low correlation is not a sufficient admission criterion.

At a 10% test weight, the candidate increased common-window book Sharpe by 0.0132 before the mean-zero control. Under the mean-zero control the change was -0.0139, and the leave-one-year-out change was negative for 2023. The observed standalone loss therefore cannot be defended as a stable diversification premium.

The proxy capacity calculation was also poor. At 1% of ADV, the fifth-percentile estimate was only \$14,919. This model is not a market-capacity certification, but it decisively fails the preregistered \$5 million floor.

Gate decision and provenance

Five of thirteen preregistered research checks passed. Exposure and correlation limits passed; minimum Sharpe, DSR, Newey-West significance, twice-cost Sharpe, product consistency, mean-zero book contribution, leave-one-year-out stability, and capacity did not. The machine verdict is **KILL**.

The ETF input source is locally mapped to Yahoo adjusted market data. A sanitized execution receipt records 5,119 USO bars, 4,645 UGA bars, and 5,163 DBC bars, with persisted partition hashes and an ingestion timestamp inside the execution window. This establishes local source identity, not redistribution rights, exact historical loader bytes, or independent replication. Raw rows remain withheld.

Reproduction and decision

The governed probe command is:



The result is bound by its preregistration, input manifest, runner hash, environment lock, and admission-contract hash. The family packet is [energy_inventory_family.json](#).

Decision: FAIL / zero sleeves. First-release discipline, low beta, and low correlation do not offset a negative, left-tailed, cost-sensitive return. Research and implementation were directed by **Arhan Canli**; no historical result is relabeled as Alpaca or live performance.

References

1. Gary B. Gorton, Fumio Hayashi, K. Geert Rouwenhorst (2013). *The Fundamentals of Commodity Futures Returns*. Review of Finance. <https://doi.org/10.1093/rof/rfs019>