The Hormuz–Bab el-Mandeb Compound_Reading Gulf Energy Chokepoints as One System

Middle East theater · Tier 4 thesis post · anchor for shipping-chokepoint-stress, oil-wti, oil-brent, lng-qatar-asia, persian-gulf-insurance-risk

---

I. Dual-chokepoint geometry

Gulf hydrocarbon outflow is not a set of independent shipping lanes. It is a sequential transit problem governed by two straits whose geographic relationship is fixed and whose operational behavior is linked.

Every cargo originating from the Persian Gulf — crude from Saudi Arabia, the UAE, Iraq, Iran, Kuwait; LNG from Qatar; condensate and refined product across the basin — must exit through the Strait of Hormuz. This is obligate. No pipeline bypass handles more than a fraction of total Gulf liftings: the East–West Pipeline (Petroline) across Saudi Arabia terminates at Yanbu on the Red Sea at roughly 5 million barrels per day of nameplate capacity; the Habshan–Fujairah line in the UAE terminates outside Hormuz at roughly 1.5 million barrels per day. Combined bypass capacity is a meaningful fraction of Saudi and Emirati liftings but a small fraction of total Gulf outflow. Iranian, Iraqi, Kuwaiti, and Qatari cargoes have no non-Hormuz exit.

Once past Hormuz, cargoes split. European-bound cargoes and a portion of Asian-bound cargoes traverse the Gulf of Aden, enter the Bab el-Mandeb strait, transit the Red Sea, and exit through the Suez Canal. Asian-bound cargoes that do not take the Suez-Mandeb corridor route east across the Indian Ocean directly to South Asian, Southeast Asian, and Northeast Asian destinations — no second chokepoint transit required.

This produces the geometric asymmetry that drives the compound dynamic: Hormuz is obligate for all Gulf outflow; Bab el-Mandeb is obligate only for cargoes destined for or transiting toward the Mediterranean and Atlantic basin. The European market is structurally exposed to both straits. The East Asian market is structurally exposed to one.

II. Hormuz as the obligate passage

The Strait of Hormuz handles a volume of seaborne oil trade that has remained structurally stable across the past decade, in the range of 17–21 million barrels per day of crude and condensate plus roughly 20% of global LNG trade, depending on demand cycle and OPEC+ production posture. The specific number matters less than the structural fact: there is no redirection path for the majority of this volume on any operationally relevant timescale.

The pipeline bypass capacity described above exists but has two operational constraints. First, it is calibrated to steady-state bypass assumptions, not surge. Second, it terminates at Yanbu and Fujairah respectively — Yanbu cargoes re-enter the Red Sea flow and remain exposed to Bab el-Mandeb; only Fujairah cargoes exit into the Gulf of Oman directly. For practical purposes, only the Habshan–Fujairah line constitutes a genuine Hormuz-and-Mandeb bypass, and its volume is small relative to total Gulf liftings.

This is the structural basis for the priced-in premium on Hormuz risk: any credible threat of sustained interdiction at Hormuz has no release valve on timescales shorter than refinery and LNG-contract cycles. The 2019 Gulf of Oman tanker incidents, the 2019 Abqaiq strikes, and the recurrent Iranian seizure cycle (Stena Impero 2019, Advantage Sweet 2023, MSC Aries 2024) each tested this dynamic without triggering full closure. The persistent finding across these episodes: insurance premiums and freight rates respond within days; physical flows respond within weeks; demand-side substitution responds within months. Hormuz remains un-replaceable at each timescale.

III. Bab el-Mandeb as the elective passage

Bab el-Mandeb handles roughly 6–9 million barrels per day of crude and product in typical conditions, plus LNG and container traffic. The volume is smaller than Hormuz but the strait's strategic function is distinct: it is the entry point to the Suez-Mandeb corridor, which is the shortest sea route between Gulf production and European demand.

The alternate route — Cape of Good Hope — adds approximately 10–14 days to a VLCC voyage from the Gulf to Northwest European destinations, depending on vessel class, weather, and fuel-economy routing. This delay is the system's pressure-release mechanism.

The pressure-release function has concrete operational costs that rarely surface in headline chokepoint coverage:

LNG cargo scheduling. LNG carriers operate on tight turnaround windows governed by destination regasification slot contracts and boil-off-gas economics. A 10–14-day delay on a Qatar–Europe voyage does not simply push back delivery — it can miss the contracted regas window, triggering either cargo diversion or penalty clauses. For Qatari LNG specifically, Cape rerouting for European-contracted cargoes degrades fleet utilization and forces charter-market adjustments that ripple across Asian spot pricing.

Refinery turnaround windows. European refineries operate planned turnaround schedules calibrated to expected crude-delivery cadence. Cape rerouting adds 10–14 days of in-transit inventory that either pushes turnaround forward at operational cost or forces refiners onto spot-market purchases at the delta price. The effect compounds across the refining fleet.

Marine insurance. War-risk insurance markets, coordinated through the Lloyd's Joint War Committee and the Lloyd's Market Association, reprice rapidly in response to listed-area designations. Additional war-risk premium for Red Sea transits has, at peak cycle points, added meaningful per-voyage cost on top of hull and P&I baseline rates; operators with sufficient tonnage calculate whether the premium delta exceeds the Cape-route fuel and time cost.

Cape rerouting is therefore not a clean substitute. It is a release valve with quantifiable pressure costs that accumulate across the system.

IV. The compound mechanic

The core analytical claim of the compound frame: interdiction pressure at one strait does not stay at that strait. It redistributes.

When Bab el-Mandeb transit becomes operationally unattractive — through Houthi interdiction, through insurance repricing, or through direct security incidents — cargoes redirect around the Cape. Physical Hormuz transit volume is unchanged. But the *strategic value* of Hormuz leverage shifts: a Hormuz disruption layered on top of Red Sea disruption removes both the primary and alternate routing for European-bound cargoes, forcing reliance on Cape routing from an already stressed Gulf-exit system. The second shock compounds because the first has already consumed the system's slack.

The reverse also holds. A Hormuz incident that constrains Gulf outflow — even briefly — tightens the volume of cargoes available to transit Bab el-Mandeb, which interacts with whatever ongoing pressure exists in the Red Sea. If Houthi interdiction is active when a Hormuz incident occurs, the pricing delta on Brent widens more sharply than either event modeled independently would predict.

This compound behavior is observable in pricing data. Brent responds more sharply to combined-chokepoint events than to single-chokepoint events of equivalent individual magnitude; WTI, insulated from Bab el-Mandeb by geography of demand, responds principally to the Hormuz variable. The Brent–WTI spread is therefore a readable indicator of compound chokepoint stress, and widens predictably when both chokepoints are under simultaneous pressure.

The compound frame yields three operational predictions. First, any interdiction campaign that targets only one strait has a ceiling on its economic effect, because the alternate routing absorbs pressure. Second, coordinated or overlapping pressure on both straits produces superlinear effects in Brent pricing and in European LNG scheduling. Third, the strategic value of controlling *either* strait rises when the other is stressed — meaning the actor with leverage at Hormuz gains, without acting, when Bab el-Mandeb is contested, and vice versa.

V. 2023–2025 as a live test

The Houthi Red Sea interdiction campaign, active from late 2023 through 2025 with varying tempo, is the clearest available test of the compound frame. Several observations from the period:

Suez traffic collapse was rapid and deep. Suez Canal Authority transit and revenue data tracked a significant reduction in transits during peak campaign tempo, with revenue losses of a scale that materially affected Egyptian balance-of-payments calculations. Tanker and container operators rerouted around the Cape in large numbers; major carriers (Maersk, MSC, Hapag-Lloyd, CMA CGM) shifted posture.

War-risk premium repricing was sharp and persistent. Lloyd's Joint War Committee listed-areas designations for the southern Red Sea and Bab el-Mandeb triggered war-risk premium increases measured in multiples of baseline, with the premium remaining elevated through sustained periods rather than reverting quickly.

Hormuz activity during the same window was calibrated. Iranian direct action at Hormuz during the active Houthi campaign was notable for what it was *not*: there was no sustained Iranian interdiction of Hormuz transit during the peak Red Sea cycle. Discrete seizures occurred (Advantage Sweet, MSC Aries, and the April 2024 incident targeting the Aries specifically); an actor pursuing complementary escalation would have produced a sustained pattern.

Brent–WTI spread behavior was consistent with the compound model. During peak Red Sea cycle, Brent carried a persistent premium reflecting the combined Suez-Mandeb detour cost and the war-risk component; WTI, geographically insulated from the Bab el-Mandeb variable, tracked its own fundamentals more closely. The spread widened during the campaign and compressed as Red Sea pressure moderated.

The period therefore provides empirical support for the compound frame and, separately, evidence about Iranian-Houthi operational coordination that the next section addresses.

VI. Iran–Houthi coordination: complementary or substitutive?

The Houthi Red Sea campaign is typically described as coordinated Iranian proxy action within a broader interdiction strategy. Under the compound frame, this description is operationally ambiguous, because two distinct coordination models produce different predictions:

Complementary model. Iran and the Houthis coordinate escalation across both straits simultaneously, producing the superlinear compound effect described in Section IV. This model predicts that sustained Houthi Red Sea activity would be accompanied by sustained Iranian Hormuz pressure.

Substitutive model. Houthi Red Sea activity *substitutes* for Iranian Hormuz action, allowing Iran to extract economic and political pressure on the Gulf-Europe corridor without assuming the escalation costs of direct Hormuz interdiction. This model predicts that active Houthi Red

Sea pressure would be accompanied by *restrained* Iranian Hormuz behavior, because simultaneous pressure at both straits would draw direct US-Gulf military response that the substitutive posture is designed to avoid.

The 2023–2025 observable pattern — sustained Houthi campaign, calibrated and discrete rather than sustained Iranian Hormuz interdiction — is consistent with the substitutive model, not the complementary one.

The operational implications differ sharply. Under the complementary model, deterrence posture must simultaneously address both straits as one coordinated threat. Under the substitutive model, deterrence posture must address Hormuz and Bab el-Mandeb as strategically linked but operationally decoupled — degrading Houthi capability does not automatically reduce Iranian leverage, because Iranian leverage was being extracted *through* Houthi capability, not alongside it. Degrading the Houthi campaign may therefore produce an Iranian posture shift toward direct Hormuz action, not reduced overall pressure.

The substitutive reading also clarifies why post-2024 degradation of the broader proxy arc — particularly Hezbollah's 2024 losses and the Iranian direct-strike reprisal cycles — coincides with increased Iranian interest in direct posture options rather than continued reliance on arc-mediated pressure. This is the structural prediction the substitutive model makes, and it is the prediction the 2024–2025 operational record is tracking toward.

VII. The geographic asymmetry: Brent, WTI, and the European binding constraint

The compound dynamic does not affect all oil benchmarks equally. Bab el-Mandeb is a binding constraint for European crude supply in a way that it is not for North American or East Asian markets.

WTI pricing reflects North American supply-demand fundamentals, with Gulf cargo exposure mediated through crude imports and refined-product flows that do not predominantly transit Bab el-Mandeb. Hormuz events affect WTI through global pricing transmission; Red Sea events affect WTI indirectly, through the Brent–WTI spread and through product-flow arbitrage.

Brent, by contrast, is structurally exposed to both chokepoints. European refining runs on a crude slate in which Gulf cargoes are a meaningful input, and Gulf-European delivery is principally Suez-Mandeb-routed. Red Sea disruption imposes direct physical cost on the Brent-relevant supply chain.

This asymmetry has strategic consequences beyond pricing. US energy-security calculations inherited from the post-1973 posture treat Hormuz as the master variable and Bab el-Mandeb as a secondary concern. This calibration reflects the period when US import dependence was the organizing interest. Under current conditions — US as a net hydrocarbon exporter,

European market as the principal Gulf-volume demand sink for Suez-Mandeb-routed cargo, East Asian market taking direct-route deliveries — the strategic weight of Bab el-Mandeb has shifted. It is now more binding on allied European supply security than on US supply security.

US deterrence posture at Bab el-Mandeb therefore operates on behalf of an allied interest more than a direct national-supply interest, which is a different strategic configuration than the one that shaped pre-2015 Gulf force posture. This distinction is under-reflected in both public discourse and in some current budget-justification documents.

VIII. Policy implications: why deterrence posture cannot pick one strait

The compound frame produces a specific policy conclusion: US force posture in the Middle East cannot be redesigned to prioritize one strait at the expense of the other without accepting a predictable degradation in the deterrent effect at the de-prioritized strait, which under the compound model will be exploited — substitutively if not complementarily.

Three design constraints follow.

Posture cannot be geographically reduced without loss of compound deterrence. Drawdown at either the 5th Fleet's Bahrain footprint (Hormuz-adjacent) or the CTF-153 Red Sea footprint (Mandeb-adjacent) produces asymmetric signal under the substitutive coordination model: pressure at the de-prioritized strait becomes more attractive because the release valve becomes more reliable for the pressuring actor.

Coalition composition matters more than hull count. The CTF-153 Operation Prosperity Guardian coalition demonstrated both the feasibility and the limits of allied co-deployment at Bab el-Mandeb. European participation was partial; allied burden-sharing patterns revealed which states treated Bab el-Mandeb as a direct interest and which treated it as a coalition obligation. Future posture design must reflect the actual distribution of supply-security interests, not an inherited assumption that Gulf energy security is a uniformly shared allied priority.

Escalation-management tools must address both straits as one signaling space. Deterrence communications to Tehran regarding Hormuz cannot credibly decouple from posture at Bab el-Mandeb under the substitutive model, because substitutive leverage implies that action at one strait is read in Tehran as a policy signal about the other. This has implications for how sanctions, military posture, and diplomatic signaling are sequenced.

What the compound frame provides

The dashboard tracks five pressure cells whose behavior is linked by the geometry this piece describes: shipping chokepoint stress, WTI, Brent, Qatar-Asia LNG pricing, and Persian Gulf insurance risk. Movement in any one cell is an input to movement in the others through the compound mechanic. Reading the cells individually produces surprise when they move together; reading them as outputs of a single two-strait system produces calibrated expectation.

Subsequent briefs on Iran's strategic posture, on Houthi operational tempo, on Gulf-state hedging behavior, and on European energy-security adjustment all operate inside the geometry this piece establishes. The Hormuz–Mandeb compound is the Middle East theater's structural base variable. Everything else in the theater is priced against it.

Citation base: IEA Oil Market Report; Suez Canal Authority transit data; Clarksons and Platts tanker-rate data; Lloyd's Joint War Committee listed-area designations; CSIS AMTI and ACLED event tracking; IAEA Iran focus page; CENTCOM and 5th Fleet operational disclosures; Tanker War (1984–88) and 2019 Strait incident historical baselines.