Ships and Storage

The Global Gas Market

On a January evening in 2024, UK power prices hit £500/MWh while Norwegian hydro was generating at £40/MWh. The two countries are connected by a 1,400 MW subsea cable, yet the price spread persisted. Why didn't the interconnector arbitrage it away? Because energy markets aren't simple supply-demand curves—they're global systems where ships move molecules across oceans, storage facilities buffer seasonal swings, and physical constraints create friction that pure financial models miss.

Understanding where UK gas comes from—and how traders manage seasonal price swings—is essential before diving into spark spreads and power trading. This guide answers two fundamental questions: *How does US shale gas reach UK power plants?* and *Why do summer and winter gas prices diverge?*

What you'll learn:

You'll understand why UK gas prices track Henry Hub in Louisiana with a 6–8 week lag, how LNG shipping costs determine whether US producers export to Europe, and why UK's tiny storage capacity (<2% of annual demand) makes winter prices volatile. You'll also learn why LNG is the marginal price-setter while pipelines provide baseload, why NBP can decouple from TTF despite interconnection, and how storage rate constraints affect trading strategies. By the end, you'll see how these global gas dynamics feed into the spark spread calculations you'll learn in the next guide.

Prerequisites:

This guide assumes you've completed the Physical Foundations guide and understand merit order dispatch and why gas-fired power plants are marginal price-setters. You don't need to know spark spreads yet—this guide explains the gas price input before you learn how it combines with power and carbon prices.

A note on currencies and units:

LNG is a global commodity traded in US dollars ($) because most contracts reference Henry Hub. European gas trades in euros (€) at TTF or pounds (£) at NBP. We'll use the currency native to each market and convert when necessary. One useful conversion: US gas is quoted in dollars per million British thermal units ($/MMBtu), while UK gas is quoted in pence per therm. For reference, 1 MMBtu = 10 therms, and typical energy content is roughly 10 kWh per cubic metre of gas.

You might wonder: if the energy conversion is that simple (1 MMBtu = 10 therms), why do traders still make unit conversion errors? Two reasons. First, currency moves complicate things—when USD/GBP swings 5% in a week, your netback calculation from yesterday is wrong today. Second, timing conventions create subtle mismatches. Gas markets use a "Gas Day" running 05:00–05:00, while power markets use midnight-to-midnight. If you're calculating spark spreads across both markets, you're comparing prices from different delivery periods, which creates small but persistent P&L errors. The energy conversion is simple; the trading reality is messier.

1. Why UK Gas Prices Track Henry Hub

If you pull up a chart of Henry Hub (the US gas benchmark in Louisiana) and NBP (the UK gas benchmark), you'll notice something surprising: they move together, but NBP lags Henry Hub by 6–8 weeks. When US shale production surges and Henry Hub falls, UK prices don't react immediately. But two months later, NBP drops too. Why?

The answer is ships. LNG tankers take 13–16 days to cross the Atlantic from the US Gulf Coast to Milford Haven (UK's largest LNG import terminal). But the lag is longer than the voyage time because the economics work like this: when Henry Hub falls, it becomes more profitable to export US gas to Europe. US producers increase LNG shipments. Those cargoes arrive in Europe over the next 4–8 weeks. European supply increases. TTF and NBP fall. The delay between US production changes and UK price adjustments is the time it takes for physical molecules to move across oceans.

Why 6–8 weeks instead of the 2-week voyage time? Because **price follows molecules, and molecules follow contracts**. When Henry Hub drops today, US producers don't instantly load ships—they negotiate new export contracts, charter vessels, schedule terminal slots, and load cargoes over the next 2–4 weeks. Those ships then take 2 weeks to cross the Atlantic. Once they arrive, they unload over several days and inject into European storage or pipelines. Only when enough cargoes accumulate (typically 3–5 additional deliveries) does European supply increase enough to pressure prices down. The full cycle—from Henry Hub move to European price response—takes 6–8 weeks. The voyage is 2 weeks; the supply response is 6–8 weeks.

For traders, this creates opportunity. If you see Henry Hub drop $1/MMBtu due to a Permian Basin production increase, you can position for lower UK gas prices 6–8 weeks out. You're trading the physical lag between markets—buying time to react before the UK market catches up.

Where UK Gas Comes From

The UK gets gas from three main sources in 2026. First, **Norwegian pipeline gas**—subsea pipes like Langeled deliver roughly 40% of UK supply directly from Norwegian fields. Second, **European interconnectors**—the IUK pipeline from Belgium allows bidirectional gas flows between continental Europe and the UK. Third, **LNG imports**—tankers arriving at Milford Haven and other terminals, mostly from the US Gulf Coast but also Qatar and other export hubs.

That third source—LNG—is what physically connects UK prices to Henry Hub. When a cargo leaves Sabine Pass (Louisiana) bound for Milford Haven, it's effectively arbitraging the Henry Hub-NBP spread. If NBP is high enough relative to Henry Hub plus shipping costs, the cargo flows. If not, it redirects to Asia or the US producer sells domestically instead. This arbitrage mechanism is what ties UK gas prices to US production trends.

Pipelines vs LNG: Baseload vs Flexible Supply

Here's a critical distinction that shapes how gas markets work: Norwegian pipeline gas and European interconnector flows are relatively **inflexible**—they provide baseload supply that changes slowly in response to seasonal contracts and long-term field production profiles. LNG, by contrast, is highly **flexible**—ships can redirect mid-voyage, producers can choose to export or sell domestically based on daily spreads, and import terminals can ramp up or down based on storage levels and demand.

This means LNG is the **marginal source** that sets UK gas prices during most market conditions. Norwegian pipeline flows are relatively constant throughout the year (though slightly higher in winter when Norwegian fields ramp up production). When UK demand spikes during a cold snap, it's LNG cargoes that respond—terminals ramp up regasification, charter rates spike to pull more ships to UK waters, and producers accelerate export schedules. When demand falls during mild weather, LNG cargoes redirect to Asia or US exports slow.

Why does this matter for traders? Because the Henry Hub-NBP spread (adjusted for shipping costs) is the equilibrium mechanism that determines how much gas flows to the UK. Norwegian pipelines don't respond to short-term price signals—their output is locked into long-term contracts and field production schedules. But LNG does. When NBP rises $2/MMBtu above its normal relationship to Henry Hub, more cargoes arrive within 2–3 weeks, pressuring NBP back down. When NBP falls relative to Henry Hub, cargoes divert away and NBP support returns.

This is why charter rates are so important: they directly affect the marginal cost of delivering LNG, which determines how responsive LNG flows are to price signals. When charter rates spike from $29,000/day to $145,000/day (as they did from early 2025 to late 2025), the shipping cost increases by $0.20/MMBtu. That might not sound like much, but it's enough to delay or redirect cargoes, tightening European supply and supporting higher prices. Traders who understand this dynamic watch charter rates as a leading indicator for European gas prices.

Why LNG Ships Exist

Natural gas at atmospheric pressure is 600 times less dense than when liquefied. This is why pipelines work for regional gas trade—you can push gas through a pipe under pressure, maintaining energy density. But for intercontinental distances, pipelines are impractical (too expensive, geopolitical risk). Instead, you cool the gas to $-162°$C, turning it into LNG. At this temperature, gas condenses into liquid form, shrinking its volume by 600x. Only then can you fit enough energy into a ship to make the voyage economical.

A typical LNG carrier holds 174,000 cubic metres of LNG. When regasified, that's enough energy to supply the entire UK for roughly 12 hours. A busy import terminal might receive 2–4 cargoes per week during peak winter months, meaning LNG arrivals are a material portion of UK supply—any disruption (terminal congestion, ships delayed by weather, cargoes redirected to Asia) can spike UK prices within days.

Gas Quality and Pipeline Specifications

When you buy gas, you're buying *energy*, not just volume. Not all gas is created equal. A cubic metre of gas from a US shale well contains slightly different energy than a cubic metre from a Qatari gas field. This variation matters for pricing, terminal acceptance, and pipeline compatibility.

The key metric is **heating value** (also called calorific value)—the amount of energy released when you burn the gas, measured in megajoules per cubic metre (MJ/m^3) or British thermal units per standard cubic foot (Btu/scf). UK NBP gas typically has a gross heating value around 39.0–40.0 MJ/m^3. US Henry Hub gas ranges 38.0–39.5 MJ/m^3. Qatari LNG tends slightly higher. These differences sound small, but on a 174,000 m^3 cargo carrying 3.5 million MMBtu, a 2% heating value difference shifts the effective price by $0.20/MMBtu—that's $700,000 on a single cargo.

Why does heating value vary? Because natural gas isn't pure methane (CH_4). It contains ethane, propane, butane, and other hydrocarbons (collectively called *natural gas liquids* or NGLs), plus small amounts of nitrogen, carbon dioxide, and trace impurities. Higher NGL content means higher energy density. US shale gas from the Permian Basin is "wet" gas with high NGL content, giving it higher heating value. Norwegian gas from offshore fields is "dry" with less NGL content and lower heating value. When an LNG cargo arrives, the terminal measures heating value at regasification and adjusts the delivered energy accordingly.

For **pipeline gas**, quality is relatively consistent. Norwegian pipeline flows to the UK have stable heating values because they come from long-producing fields with predictable composition. Traders don't worry much about quality adjustments for pipeline deliveries—the gas that enters the pipe in Norway is the gas that arrives at UK terminals. Contracts specify a narrow heating value range, and field operators blend production to meet it.

For **LNG imports**, quality is critical. Cargoes from different export terminals have different compositions. US Gulf Coast LNG (from shale gas) differs from Qatari LNG (from conventional fields). When a cargo arrives at Milford Haven, terminal operators check several specifications:

- **Gross heating value:** Must fall within 1030–1130 Btu/scf (equivalent to 38.3–42.0 MJ/m^3). Cargoes outside this range may require blending with other supplies before injection into UK pipelines.
- **Wobbe Index:** A measure of interchangeability—can this gas be safely mixed with existing pipeline gas without changing burner behavior in end-user equipment? If Wobbe Index is too far off spec, the gas can't be injected directly.
- **Sulfur content:** Maximum 5 mg/Nm^3 (milligrams per normal cubic metre). Higher sulfur corrodes pipelines and creates emissions compliance issues.
- **Moisture and impurities:** Water content, CO_2, nitrogen (N_2), and hydrogen sulfide (H_2S) must be within tight limits to prevent pipeline corrosion and meet combustion standards.

If a cargo arrives out of spec, the terminal has two options. First, **blend it** with other cargoes or pipeline gas to bring the mix within specification. This is why large import terminals like

Milford Haven have blending facilities—they can combine a high-heating-value US cargo with lower-heating-value Norwegian pipeline gas to hit the UK network's target range. Second, **reprice it**. LNG contracts include quality adjustment clauses. If the delivered gas has 2% lower heating value than contracted, the buyer pays 2% less per cubic metre to get the same energy cost. This protects both parties from quality risk.

Why does this matter for traders? Two reasons. First, **netback calculations**. When a US producer evaluates whether to export LNG to Europe or Asia, they compare delivered energy value, not just volumetric price. If European gas trades at €10.00/MWh but has 5% lower heating value than contracted, the effective price is €9.50/MWh. That shifts the netback by \$0.18/MMBtu—enough to make Asian markets more attractive even if headline European prices look higher.

Second, **delivery disputes**. Imagine you've bought a cargo at \$10.60/MMBtu based on a contracted heating value of 40.0 MJ/m³. The cargo arrives with 38.5 MJ/m³. You're entitled to a price adjustment:

$$\text{Adjusted price} = \$10.60 \times \frac{38.5}{40.0} = \$10.21/\text{MMBtu}$$

That's \$0.39/MMBtu less, or \$1.37 million on a 3.5 million MMBtu cargo. Without understanding quality specs, you'd miss this adjustment and overpay.

Quick Check

A cargo arrives with heating value 39.0 MJ/m³ instead of the contracted 40.5 MJ/m³. The contract price is £2.80/therm. What should the buyer pay per therm to maintain the same energy cost?
Answer: Adjust by the ratio of heating values: £2.80 × (39.0 / 40.5) = £2.80 × 0.963 = £2.70/therm. The buyer should pay 10p less per therm because the cargo delivers 3.7% less energy than contracted.

The broader lesson: energy markets trade *energy*, not molecules. When you see NBP quoted at 70p/therm or TTF at €28/MWh, those prices assume a standard heating value. Actual delivered energy can differ by 2–5%, and sophisticated traders adjust their positions accordingly. For pipeline gas this is background detail—Norwegian flows are stable. But for LNG, where cargoes come from diverse sources with varying quality, it's a material pricing factor. Always confirm what heating value your contract references, and always verify delivered quality matches before settling payments.

2. LNG Shipping Economics

The cost to deliver an LNG cargo from Louisiana to the UK comes in three pieces. First, **liquefaction**—cooling the gas at the US export terminal costs around \$2.50/MMBtu (this is the benchmark cost for modern facilities like Sabine Pass). Second, **shipping**—the charter rate for the vessel itself, which varies wildly with market conditions. Third, **regasification**—warming the LNG back into gas at Milford Haven, typically \$0.40/MMBtu.

It’s the middle piece—shipping—that creates volatility and trading opportunity. LNG charter rates are spectacularly unstable. In early 2025, average charter rates hit \$29,000 per day, a 48% drop from 2024 levels as new ships entered service and demand softened. But by late November 2025, winter demand in both Europe and Asia drove Atlantic basin rates to \$145,000 per day, a 400%+ spike. That’s a 5x swing in a single year.

Why such volatility? When terminals are busy and ships are in high demand (typically November through February), rates spike. Traders sometimes charter ships just to store LNG at sea, waiting for prices to rise before delivering—this “floating storage” absorbs available tonnage and pushes charter rates higher. Ships moving empty between basins (from Asia back to the US to pick up another cargo) add non-revenue days that charterers must recover through higher rates on revenue voyages. And repositioning delays compound during peak periods when cargoes queue at congested terminals.

The good news for traders is that 170+ new LNG carriers entered service in 2026, creating a fleet surplus expected to dampen charter rate volatility over the next 2–3 years. But for anyone trading the 2025–26 winter, charter rate swings were the difference between profitable and unprofitable positions.

Calculating the Shipping Cost

Let’s work through a concrete example to understand what it costs to move gas from Henry Hub to NBP. We’ll ship a cargo from Sabine Pass (Louisiana) to Milford Haven (UK).

The voyage takes roughly 13 days one-way. Add 1–2 days for loading at the export terminal and 1–2 days for unloading at the import terminal, giving 15–16 days outbound. But the ship doesn’t teleport back to Louisiana—it needs to return empty for the next cargo. Total round-trip time: 28–30 days.

If we charter a ship at \$120,000 per day (a mid-range rate for late 2025), the total voyage cost is:

$$\$120{,}000 \times 28 = \$3{,}360{,}000$$

Our cargo holds 174,000 cubic metres of LNG. When regasified, that’s about 3.5 billion cubic feet (bcf) of natural gas. In the units US traders use, 1 bcf equals 1 million MMBtu, so the cargo is 3.5 million MMBtu. Divide the voyage cost by cargo size:

$$\frac{\$3{,}360{,}000}{3{,}500{,}000\text{ MMBtu}} = \$0.96/\text{MMBtu}$$

That’s just the charter cost. Add liquefaction (\$2.50/MMBtu) and regasification (\$0.40/MMBtu):

$$\boxed{\text{Shipping stack} = \$0.96 + \$2.50 + \$0.40 = \$3.86/\text{MMBtu}}$$

This “shipping stack” is the key number for US producers deciding whether to export. If European gas prices are high enough to cover Henry Hub price plus \$3.86, the cargo flows. If not, the producer sells domestically or looks for better-paying markets in Asia.

Quick Check

What happens to the shipping cost if charter rates spike to \$145,000/day? Assume the same 28-day round-trip and 3.5 million MMBtu cargo.
Answer: Charter cost rises to \$145,000 $\times$ 28 = \$4,060,000, giving \$1.16/MMBtu for shipping alone. Total stack becomes \$1.16 + \$2.50 + \$0.40 = \$4.06/MMBtu. That's \$0.20/MMBtu more than baseline—which sounds small, but on a full cargo that's \$700,000 in extra cost. A trade that was profitable at \$3.86 shipping might become unprofitable at \$4.06, causing producers to cancel or delay exports.

The Netback Calculation: Should US Producers Export?

US producers face a daily decision: sell gas domestically at Henry Hub prices, or export it to Europe (or Asia) at TTF/NBP prices minus shipping costs? The **netback price** answers this question. It's what the producer can "net back" to Henry Hub after subtracting all costs to reach Europe.

The formula:

$$\boxed{\text{Netback} = \text{TTF Price} - \text{Shipping Stack}}$$

If the netback exceeds Henry Hub plus production costs, export is profitable. If not, sell domestically.

Let's use late 2025 prices. TTF (the European benchmark) traded around \$10.60/MMBtu equivalent. Henry Hub sat at \$4.30/MMBtu. Subtracting our \$3.86 shipping stack:

$$\text{Netback} = \$10.60 - \$3.86 = \$6.74/\text{MMBtu}$$

Now compare to US production costs. Shale gas production from the Permian Basin costs around \$5.00–6.00/MMBtu depending on well productivity. If your cost is \$5.50/MMBtu and your netback is \$6.74, you're making \$1.24/MMBtu profit by exporting. On a full cargo (3.5 million MMBtu), that's \$4.3 million.

Here's the trading insight: this netback calculation runs continuously in US producers' heads. When TTF rises or Henry Hub falls, the netback improves and more cargoes flow to Europe. Over 4–8 weeks, those cargoes arrive, increasing European supply and pushing TTF and NBP down. Eventually, the netback shrinks and export volumes stabilize. The netback is the equilibrium mechanism—it's what prevents Henry Hub and TTF from diverging indefinitely.

For UK traders, watching the Henry Hub-TTF spread tells you whether more or fewer LNG cargoes are heading your way. A widening spread signals increasing exports to Europe, which will pressure UK prices down in 6–8 weeks. A narrowing spread signals slowing exports, which might tighten UK supply and support prices.

But there's a third player in this calculation: Asia. US producers don't just choose between domestic sales and European exports—they also evaluate Asian markets. Let's compare the netbacks for a typical late-2025 scenario:

Market	Price	Shipping Stack	Netback to US
TTF/NBP (Europe)	$10.60/MMBtu	$3.86/MMBtu	$6.74/MMBtu
JKM (Asia)	$11.20/MMBtu	$4.90/MMBtu	$6.30/MMBtu

Notice the counter-intuitive result: Asia pays $0.60/MMBtu more than Europe ($11.20 vs $10.60), but the netback to a US producer is **lower** ($6.30 vs $6.74). Why? Because the voyage to Asia is roughly twice as long—25 days transit vs 13 days to Europe—and that extra distance costs $1.04/MMBtu in additional shipping. The "expensive" market isn't always the most profitable market once you factor in distance.

This is why charter rates and voyage times matter so much for global gas flows. If Asian prices spike to $14/MMBtu (say, due to a cold snap in Japan), the netback jumps to $9.10—suddenly $2.36/MMBtu better than Europe. Cargoes that were heading to Milford Haven will divert mid-voyage to Tokyo. But when Asian and European prices are close, the shorter voyage to Europe wins. US producers constantly run these three-way netback calculations (domestic, Europe, Asia) to decide where to send each cargo.

Asia Competes for the Same Cargoes

The table above shows a normal market state where European netbacks beat Asian netbacks, so cargoes prefer Europe. But markets shift. When Asian prices spike—say, a cold winter in Japan drives JKM to $14/MMBtu—the Asian netback becomes more attractive ($9.10 vs Europe's $6.74). Suddenly, ships that were heading to Milford Haven divert mid-voyage to Tokyo. European imports drop. TTF and NBP rise.

This flexibility is what makes LNG the global connector. Unlike pipeline gas (locked into fixed routes), LNG cargoes optimize in real-time. If Asia outbids Europe, the molecules go east. If Europe outbids Asia, they go west. UK traders who only watch Henry Hub and UK supply-demand miss half the picture. When Asian demand surges, it pulls cargoes away from Europe, tightening UK markets even though nothing changed domestically. This three-way arbitrage (Henry Hub, TTF, JKM) is why UK gas trading requires a global perspective—and why traders monitor Asian weather forecasts as closely as UK ones.

3. The UK Storage Problem

The UK has a structural vulnerability that amplifies price volatility: tiny storage capacity. UK storage holds less than 2% of annual gas demand. For comparison, continental Europe has roughly 20% storage capacity. This 10x difference has massive trading implications.

Why does storage matter? Because gas demand is wildly seasonal. Winter heating demand can be 2–3 times higher than summer. Without storage, you'd need enough pipeline and LNG import capacity to meet peak winter demand, which would sit idle all summer (economically wasteful). Or you'd face winter shortages and rationing (politically unacceptable). Storage solves this by letting you inject gas during low-demand summer months and withdraw during high-demand winter.

But the UK barely has any. Rough—a depleted gas field in the North Sea—used to provide

70% of UK storage capacity (about 3.2 bcm). It partially closed in 2017 due to structural issues, reopened with reduced capacity (1.5 bcm), and in 2025 announced it wouldn't refill for winter 2025/26 due to potential conversion to hydrogen storage. That left the UK entering winter with roughly half its normal storage buffer.

Why This Makes UK Prices Volatile

With minimal storage, the UK is heavily dependent on **continuous imports**—LNG cargoes arriving at Milford Haven and pipeline gas from Norway—to meet daily demand. There's no cushion. When a cold snap hits, demand spikes immediately and prices must rise to ration demand (households don't cut heating usage until prices get painfully high). When LNG cargoes are delayed by weather or redirected to Asia, UK supply tightens and prices spike.

In January 2024, a brief cold period combined with low wind generation (wind farms not generating, so more gas-fired CCGTs needed to run) drove NBP above £5.00/therm. At the same time, TTF on the continent stayed below €60/MWh equivalent (roughly £2.50/therm) because European storage provided a buffer. The UK-Europe **basis risk**—the risk that NBP decouples from TTF despite being connected by pipelines—can blow out to £1–2/therm during UK-specific tightness.

For traders, this means UK gas positions require different risk management than European positions. You can't assume NBP will track TTF just because the IUK pipeline connects them (20 bcm/year capacity). When UK storage is low and winter weather hits, NBP can decouple sharply upward, and traders short NBP against TTF positions discover painful losses.

What is Basis Risk?

Basis risk deserves explicit attention because it's one of the most dangerous traps for gas traders. The "basis" is the price difference between two related markets—in this case, NBP (UK) versus TTF (continental Europe). These two benchmarks are physically connected by the IUK pipeline, which has 20 bcm/year bidirectional capacity. In theory, gas should flow from the cheaper market to the more expensive one, keeping prices aligned.

In practice, basis can explode. During the January 2024 cold snap, NBP traded at a £2.50/therm premium to TTF. Why didn't arbitrage close the gap? Three reasons. First, the IUK pipeline was running at maximum capacity in the *Europe-to-UK direction*—yes, backwards. This is deeply counter-intuitive: Europe, which produces almost no natural gas domestically, was *exporting* to the UK, which has North Sea production. But Europe's storage buffer (20% of annual demand) was more valuable than UK's production flow during peak demand. European traders withdrew from storage and sold to the UK at £5.00/therm, arbitraging the massive basis. But the IUK pipeline's physical capacity (20 bcm/year) was maxed out—enough flow to profit from the spread, but not enough to equalize prices. Second, UK regasification terminals (Milford Haven, South Hook) were already running at maximum throughput—you couldn't quickly increase LNG imports to relieve UK tightness. Third, Norwegian pipeline flows were fixed by long-term contracts and field production limits—they couldn't ramp up on short notice.

The result: NBP spiked while TTF stayed subdued. Traders who were long TTF and short NBP (betting on convergence) lost money as the basis widened instead of narrowing. This is basis risk—the risk that two correlated markets decouple due to physical constraints, policy differences, or local supply-demand imbalances.

You might wonder: given the UK has minimal storage, why don't UK traders just lease German or Dutch storage and pipe gas back when needed? They do—UK utilities and trading houses routinely book continental storage for seasonal positions. But this strategy doesn't eliminate basis risk; it transforms it. If you inject 100 mcm into German storage in summer and plan to withdraw it during a UK cold snap in January, you're betting the IUK pipeline will have spare capacity when you need it. During the January 2024 event, the IUK was maxed out. Traders with gas in European storage couldn't move it to the UK fast enough. Their gas was "trapped" in the wrong market—worth €60/MWh in Germany while UK prices hit £130/MWh equivalent. The physical pipe is the bottleneck, and no amount of storage booking eliminates that constraint.

For traders, the lesson is clear: don't assume interconnected markets move in lockstep. Physical infrastructure has limits. Storage disparities create different volatility profiles. And during extreme events (cold snaps, supply disruptions, wind droughts), basis can move against you faster than you can hedge. Always size basis trades assuming the spread could widen beyond historical norms.

Storage Types: Depleted Fields vs Salt Caverns

The small amount of UK storage that exists comes in two types with very different characteristics. Understanding the difference helps explain why storage capacity numbers don't tell the whole story.

Depleted gas fields like Rough are old reservoirs that have been emptied and repurposed for storage. They're enormous—Rough at 1.5 bcm is by far the UK's largest facility, roughly half of total UK storage. But they're slow. Injection and withdrawal rates are constrained by reservoir geology, taking weeks to fill or empty. They're designed for seasonal storage: inject all summer (April–September), withdraw all winter (October–March). You can't respond to day-to-day price swings. If prices spike mid-November, Rough can't ramp up withdrawal fast enough to capitalize.

Salt caverns like Stublach and Holford are carved into underground salt formations. They're much smaller—Stublach holds 0.41 bcm, less than a third of Rough. But they're fast. You can fill or empty a salt cavern in days, cycling multiple times per year. This makes them valuable for short-term trading: if a cold snap drives prices up mid-November, salt caverns can respond immediately. If prices crash during a warm week in January, you can re-inject and wait for the next cold period.

The economic difference is crucial. Depleted fields have lower per-unit costs (fixed costs spread over huge capacity) but no flexibility. Salt caverns cost more per unit but generate higher returns through multiple cycles and volatility capture. Sophisticated storage traders use depleted fields for bulk seasonal trades and salt caverns for tactical short-term plays.

Storage Rate Constraints

Here's a critical detail that textbook models often miss: you can't inject or withdraw gas from storage instantaneously. Every facility has **rate limits**—maximum flow rates measured in millions of cubic metres per day (mcm/day) or gigawatt-hours per day (GWh/day). These physical constraints determine how quickly storage can respond to price signals and shape trading strategies.

Rough, for example, has maximum withdrawal rates around 42 mcm/day when full, declining to around 20 mcm/day when nearly empty (lower reservoir pressure reduces flow). Over a 120-day winter season (October–February), that means Rough can withdraw its full 1.5 bcm capacity at an average rate of 12.5 mcm/day. But if a 7-day cold snap hits in mid-January and UK demand spikes by 50 mcm/day, Rough can only supply a fraction of the shortfall. The rest must come from increased LNG imports or Norwegian pipeline flows—and if those can't ramp up fast enough, NBP spikes.

Salt caverns are faster. Stublach (0.41 bcm capacity) can withdraw at up to 20 mcm/day when full, meaning it can empty in roughly 20 days. This high withdrawal-to-capacity ratio makes salt caverns extremely valuable during price spikes—they can inject 100 mcm during a warm October week when prices are low, then withdraw it over 5 days during a November cold snap when prices triple. Depleted fields can't cycle this fast, so they miss short-duration arbitrage opportunities.

Why does this matter for traders? Two reasons. First, storage acts as a **supply buffer** that dampens price volatility, but only up to the rate limit. Once storage is withdrawing at maximum capacity, further demand increases must be met by more expensive marginal sources (redirecting LNG cargoes from Asia, displacing industrial gas users). This creates step-function price jumps. Second, traders with access to fast-cycle storage (salt caverns, or depleted fields with high withdrawal rates) can capture short-term volatility that longer-cycle storage misses.

The practical implication: when evaluating storage arbitrage opportunities, don't just calculate the seasonal spread (summer vs winter prices). Also consider whether you can actually execute the withdrawal during the high-price period. If winter prices spike for 3 days due to a cold snap, but your storage facility needs 10 days to empty, you'll miss most of the profit opportunity. Rate constraints are the hidden friction that separates theoretical storage value from realized profits.

4. Storage Arbitrage: The Seasonal Trade

Storage economics boil down to one question: is the summer-winter price spread wide enough to cover the costs of storing gas for 6 months? This is **seasonal arbitrage**—buying cheap in summer, holding through low-demand months, selling expensive in winter.

But it's not as simple as "buy at €28, sell at €36, pocket €8." Storage has costs at every step that eat into the spread.

The Break-Even Calculation

Let's work through a realistic example using a German salt cavern (we'll use Germany because UK storage is too small for most traders to access; UK traders often use European storage for seasonal trades and then transport gas back to UK when needed).

You have access to 100 million cubic metres (mcm) of salt cavern storage. TTF summer prices (Q2 2025) trade at €28/MWh. Winter prices (Q1 2026) trade at €36/MWh. That's an €8/MWh seasonal spread. Looks profitable, right?

Storage costs:

1. **Gas purchase:** €28/MWh × 1.055 million MWh (100 mcm converted using 10.55 kWh per cubic metre) = €29.54 million
2. **Injection cost:** €0.15/MWh to inject, totaling €158,000
3. **Capacity fee:** €2/MWh/month for 6 months (June–December), totaling €12.66 million
4. **Withdrawal cost:** €0.20/MWh to withdraw, totaling €211,000

Total costs: €29.54m + €0.16m + €12.66m + €0.21m = €42.57 million. Now sell in winter at €36/MWh: revenue is €37.98 million. Net result: **a loss of €4.59 million**.

The problem is clear. The €8/MWh spread looks attractive, but storage costs €12.35/MWh (€0.15 injection + €12.00 capacity + €0.20 withdrawal). You're underwater by €4.35/MWh. To break even, winter prices would need to hit:

$$\boxed{\text{Break-even} = \text{Summer price} + \text{Storage costs} = €28 + €12.35 = €40.35/\text{MWh}}$$

This is the key insight: forward curves that look like they reward storage often don't once you add costs. A €6/MWh move in winter prices (from €36 to €42) turns a €4.6 million loss into a €1.7 million profit. This sensitivity is why storage traders watch forward curves obsessively and adjust positions as winter expectations shift.

Quick Check

If winter prices rise to €45/MWh (everything else unchanged), what's the profit?
Answer: Revenue becomes €45 × 1.055m = €47.48m. Costs stay at €42.57m. Profit: €47.48m - €42.57m = €4.91m. The spread widened from €8 to €17, and the €9 improvement translated directly to profit (€9 × 1.055m = €9.5m additional revenue, minus rounding).

Forward Curves: Contango and Backwardation

The shape of the gas forward curve determines whether storage makes economic sense. When forward prices rise as you move out in time (summer €28, winter €36), the curve is in **contango**. Storage is potentially profitable—you buy cheap now, hold, sell expensive later. When forward prices fall as you move forward (summer €28, winter €24), the curve is in **backwardation**. Storage destroys value—you'd be buying at €28 and selling at €24 six months later.

Why do curves invert to backwardation? It signals current tightness and expected future looseness. Perhaps LNG imports are low right now (pushing spot prices up), but traders expect new US export capacity to come online in 3 months (keeping forward prices depressed). Or perhaps storage is already at 95% full (no demand for summer injections, so summer stays high), but expectations are for a mild winter (keeping winter forwards low).

The 2022 European gas crisis is the textbook example of extreme contango. Following Russia-Ukraine supply disruptions, summer 2022 gas traded around €80/MWh while winter 2022/23 hit €200/MWh—a €120 spread driven by fear of winter shortages. Traders with storage capacity made extraordinary returns. But by 2023, as European storage filled to record levels (governments mandated 90% fills by November 2022, later relaxed to 75% in 2025), the curve flattened and even inverted slightly. Storage operators who expected fat spreads to persist lost money.

The lesson: forward curves are policy-sensitive and regime-dependent. When the EU mandated storage fills, it forced massive injection volumes, pushing summer prices up and compressing contango. Storage traders must watch policy changes, weather forecasts, and supply disruptions—all reshape the curve and make or break seasonal trades.

Who Actually Profits from Storage?

If the break-even is €40.35 but winter only trades at €36, who's using storage? The arithmetic says you lose €4.59 million—so why do it?

The answer lies in a critical distinction between **capacity costs** (the €12.66 million fee to lease storage space for 6 months) and **variable costs** (injection + withdrawal, totaling €0.35/MWh). For traders renting storage short-term, capacity is a real cost—you pay it regardless of whether you inject. But for **long-term capacity holders**—facility owners or firms with 10-year leases—capacity is a **sunk cost**. They've already committed to the capacity fee through long-term contracts. Their marginal decision on each trade is just: do injection + withdrawal costs (€0.35/MWh) justify the spread?

With this framing, the economics flip. If you own capacity, your break-even is:

$$\text{Break-even} = £28 + £0.35 = £28.35/\text{MWh}$$

Winter at €36 now yields €7.65/MWh profit, not a €4.35 loss. On 1.055 million MWh, that's €8.07 million profit instead of a €4.59 million loss. **Owning capacity fundamentally changes storage economics.** This is why storage facilities are valuable assets—the capacity fee that crushes short-term traders is irrelevant to facility owners, who profit even when curves look unprofitable to outsiders.

Second, **fast-cycle traders**. Salt caverns can inject and withdraw 5–10 times per year. If you spread the 6-month capacity fee (€12/MWh) across 3 cycles, effective cost per cycle drops to €4/MWh. Now you need €32.35/MWh winter price, not €40.35. Multiple cycles make storage viable even with narrower spreads.

Third, **hedgers, not arbitrageurs**. A UK utility buying gas for winter delivery might inject storage in summer as insurance against price spikes, even if the forward curve doesn't justify it. They're paying €4–5/MWh for price protection (optionality), not chasing arbitrage profits. For them, storage is a hedge, not a spread trade.

5. Python Implementation: Storage Arbitrage Calculator

Here's a practical tool for evaluating whether a seasonal spread justifies storage costs:

```python
def storage_arbitrage_profit(summer_price, winter_price, volume_mwh,
                              injection_cost, withdrawal_cost,
                              capacity_fee_per_month, months):
    """
    Calculate storage arbitrage profit/loss.

    Args:
        summer_price: Summer gas price (EUR/MWh)
        winter_price: Winter gas price (EUR/MWh)
        volume_mwh: Storage volume (MWh)
        injection_cost: Cost to inject (EUR/MWh)
        withdrawal_cost: Cost to withdraw (EUR/MWh)
        capacity_fee_per_month: Monthly capacity fee (EUR/MWh/month)
        months: Number of months storage held

    Returns:
        dict: profit/loss, break_even_price, and cost breakdown
    """
    # Purchase gas in summer
    purchase_cost = volume_mwh * summer_price

    # Injection, capacity, withdrawal
    injection_total = volume_mwh * injection_cost
    capacity_total = volume_mwh * capacity_fee_per_month * months
    withdrawal_total = volume_mwh * withdrawal_cost

    # Revenue in winter
    revenue = volume_mwh * winter_price

    # Calculate profit/loss
    total_costs = (purchase_cost + injection_total +
                   capacity_total + withdrawal_total)
    profit_loss = revenue - total_costs

    # Break-even calculation
    cost_per_mwh = (injection_cost +
                    (capacity_fee_per_month * months) +
                    withdrawal_cost)
    break_even = summer_price + cost_per_mwh

    return {
        'profit_loss': profit_loss,
        'break_even_winter_price': break_even,
        'total_costs': total_costs,
        'revenue': revenue,
        'cost_per_mwh': cost_per_mwh
    }

# Example: German salt cavern, Q2 2025 -> Q1 2026
```

```python
result = storage_arbitrage_profit(
    summer_price=28.0,            # EUR/MWh Q2 2025
    winter_price=36.0,             # EUR/MWh Q1 2026
    volume_mwh=1_055_000,          # 100 mcm ~ 1.055m MWh
    injection_cost=0.15,
    withdrawal_cost=0.20,
    capacity_fee_per_month=2.0,
    months=6
)

print(f"Profit/Loss: EUR {result['profit_loss']:,.0f}")
print(f"Break-even winter price: EUR
   {result['break_even_winter_price']:.2f}/MWh")
print(f"Cost per MWh stored: EUR {result['cost_per_mwh']:.2f}/MWh")

# Output:
# Profit/Loss: EUR -4,589,250
# Break-even winter price: EUR 40.35/MWh
# Cost per MWh stored: EUR 12.35/MWh
```

This calculator helps traders evaluate whether forward curves justify storage positions. In practice, you'd extend this to model multiple scenarios (mild winter = €32, normal winter = €36, cold winter = €45) and calculate expected value across probabilistic outcomes.

The Data Freshness Challenge

The calculator above is clean Python—it runs in milliseconds and the logic is straightforward. But productionizing this for real trading exposes a data problem that would break a naive implementation: **data freshness and availability**.

To run this calculator effectively in production, you need real-time or near-real-time feeds for three categories of data, each with different freshness requirements and access challenges:

Charter rates: These are the most opaque. LNG charter rates trade over-the-counter (OTC), not on exchanges. Published rates from Spark Commodities or Argus Media lag by 1–2 days and represent averages, not the marginal rate for your specific route and timing. If you're building a system that needs to decide today whether to charter a ship for next week, you need direct broker relationships or proprietary market intelligence. Historical charter rate data won't capture the 5x spikes we discussed—you'll underestimate volatility and mis-price risk.

Currency pairs: USD/GBP and EUR/GBP exchange rates are easy to source (financial data providers like Bloomberg, Refinitiv, or free APIs). But these move intraday. If your netback calculation uses yesterday's FX rate and GBP strengthens 2% overnight, your \$6.74 netback becomes \$6.61 in GBP terms—potentially flipping a profitable cargo to marginal. Production systems need live FX feeds (sub-second latency for high-frequency traders, sub-minute for most use cases) and must recalculate positions whenever FX moves.

Forward curves: TTF and NBP forward prices are liquid on ICE and EEX exchanges, with real-time data feeds available. But there's a subtlety: the *quoted* price (what you see on ICE) and the *executable* price (what you actually trade at) can differ by €0.50–1.00/MWh during volatile periods due to bid-ask spreads and market impact. If you're backtesting strategies using settlement prices, you're assuming zero slippage—but in January 2024, trying to execute a 100 mcm trade would have moved the market 2–3%. Your calculator says "profitable," but execution says "you just lost €2 million to slippage."

The software engineering challenge is that these three data sources have different latencies (OTC charter rates lag days, FX updates sub-second, forward curves update every few seconds), different reliability profiles (exchange data is clean, OTC data is noisy), and different costs (FX is cheap/free, forward curve data costs $1,000–5,000/month, charter rate intelligence requires broker relationships). A production system needs data validation (sanity-check charter rates against historical ranges, reject stale FX quotes), fallback logic (if live feed drops, use cached rate with timestamp warning), and alerting (if spreads move 10% while your data is stale, halt trading).

For software engineers: this is the unglamorous reality of energy trading systems. The Python math is the easy part. Data freshness, feed reliability, and handling missing/stale/corrupt data is where most bugs happen—and where P&L gets destroyed. Always build with the assumption that data will be late, wrong, or missing, and design your system to fail safely when it is.

When Policy Overrides Economics

There's a critical assumption baked into the calculator above: it assumes all market participants optimize for profit. Every actor is comparing spreads to costs and only injecting when the math works. But real markets don't always behave this way. Government policy can force non-economic behavior that breaks the model.

The EU storage mandate is the perfect example. Following the 2022 Russia-Ukraine crisis, European governments mandated that storage facilities reach 90% fill by November 2022 (later relaxed to 75% for 2025 onwards). This created **forced buyers**—utilities and national importers who had to inject gas regardless of economics. Even when summer prices hit €80/MWh and winter forwards sat at €90/MWh, they kept injecting. The calculator would scream "unprofitable!" (€10 spread minus €12.35 costs = €2.35 loss per MWh). But policy compliance overrode profit optimization.

This forced buying distorts both sides of the market. On the injection side, it pushes summer prices UP—artificial demand from mandated fills absorbs supply that would otherwise keep summer prices low. On the withdrawal side, it pushes winter prices DOWN—guaranteed high storage levels remove winter scarcity risk, capping winter prices. Both effects compress the spread. Your calculator might predict an €8 spread is too narrow to profit, but in 2022, policy-driven flows crushed spreads even further. By summer 2023, the curve actually inverted (backwardation) in some months as storage filled to record levels while winter demand expectations normalized.

For software engineers building trading systems, this is a crucial lesson: your optimization model can be mathematically correct but operationally incomplete. The calculator assumes actors maximize profit subject to physical constraints (storage costs, rate limits). It doesn't account for policy constraints—mandated fill levels, strategic reserves, political pressure to ensure supply security. When policy intervenes, markets decouple from pure economics. Traders who ran the numbers in mid-2023 and saw "profitable storage spreads" discovered those spreads evaporated as forced injections compressed the curve.

The practical implication: always check for policy regimes that create forced flows. EU storage mandates, strategic petroleum reserve releases, renewable energy subsidies, carbon price floors—these all inject non-economic actors into markets. Your model needs either explicit policy variables or at minimum, alerts when policy-driven flows are likely. Otherwise, you'll build a system that's optimizing a model that no longer matches market reality.

6. How This Connects to Spark Spreads

You now understand the mechanisms that connect UK gas prices to global markets and shape seasonal price dynamics. LNG ships create spatial arbitrage between Henry Hub and NBP, establishing why UK gas prices track US production with a 6–8 week lag. The flexibility of LNG (versus inflexible pipeline baseload) makes it the marginal source that sets UK prices. Storage creates temporal arbitrage between summer and winter, but rate constraints limit how fast storage can respond to price spikes. And basis risk means NBP can decouple from TTF despite physical interconnection, creating dangerous spreads for traders who assume convergence.

These aren't abstract concepts—they directly feed into the spark spread calculations you'll learn next. In the Spark Spread guide, you'll see formulas like "Clean Spark Spread = Power - (Gas × Heat Rate) - (Carbon × Emissions Factor)." Now you understand where that gas price comes from. When you see NBP at 70p/therm, you'll know that price reflects:

1. Henry Hub fundamentals (US shale production, domestic demand)
2. LNG shipping costs ($3.86/MMBtu to cross the Atlantic, but volatile with charter rates)
3. LNG as the flexible marginal source (pipelines are baseload, LNG responds to spreads)
4. Seasonal storage dynamics (contango pushing summer down, winter up)
5. UK-specific tightness (minimal storage, dependence on continuous imports)
6. Basis risk (NBP can decouple from TTF when UK storage is low or imports disrupted)

When Henry Hub moves, you can anticipate NBP moving 6–8 weeks later. When the summer-winter spread compresses, you know storage injections are slowing and winter supply might tighten. When Asian LNG prices spike, you know cargoes might divert away from Europe, squeezing UK supply. When charter rates spike, you know the shipping stack is widening and fewer cargoes will flow economically. This global view transforms spark spread trading from "power minus gas minus carbon" arithmetic into strategic positioning based on leading indicators.

The spark spread is the bridge between gas markets (what you've just learned) and power markets (what comes next). Gas is the fuel input; power is the output. Understanding how gas prices form globally—and how quickly they can move when physical constraints bind—makes spark spread trading less mechanical and more informed by physical market structure. You're ready for the next guide.

About Jordan Dimov

Portfolio CTO and energy trading software specialist with over 20 years of software engineering experience and 7 years in the energy commodities sector. Based in London since 2011, working through A115 Ltd, a London-based contracting and consulting company.

Previous roles include trading platform development at Shell, Centrica Energy, and Limejump, delivering systems for front office trading, middle office risk management, and back office settlement across gas, power, and environmental markets.

Professional Services

Individual & Team Training

- Energy trading bootcamps for software engineers
- ETRM development training for technical teams

Technical Consulting

- Architectural review and validation sprints
- Code review and technical assessment
- Bespoke energy trading software development

Business Development

- Strategic advisory for energy trading technology firms
- Investor relations support in the energy trading sector

Contact

Email:	jdimov@a115.co.uk
LinkedIn:	linkedin.com/in/jdimov
Website:	a115.co.uk