Pipes, Wires, and the Merit Order

Physical Foundations of Energy Trading

Most people think of trading as numbers on screens—prices, volumes, P&L. But energy trading is different. Unlike equities or currencies, you can't just "settle in cash" if things go wrong. When you trade gas or power, someone, somewhere, has to physically deliver it through a pipe or a wire. And if they can't—because a pipeline is full, an interconnector is at capacity, or the grid is out of balance—the consequences are immediate and expensive.

This guide is about the **physical reality** beneath energy markets. The pipes that carry gas. The wires that move electricity. The system operators who keep everything balanced in real time. The merit order that determines which power plants run and which sit idle. These aren't abstractions—they're the constraints that create trading opportunities, shape prices, and define risk.

Understanding these physical foundations is what separates energy traders from other commodity traders. The reason power prices can swing from £80/MWh to negative in a few hours, or why Scottish wind farms get paid millions to turn *off*, isn't finance—it's physics.

We'll start with electricity, not because it's more important than gas, but because its constraint is the most extreme. Once you understand why instant balance is non-negotiable for power, everything else—system operators, merit order dispatch, real-time pricing—makes sense as *solutions* to that constraint. Gas is actually easier to grasp once you've seen the hard case first.

1. The 50Hz Problem

Electricity is the only commodity that must balance supply and demand *instantly*. Not within the hour, not within the minute—within *milliseconds*. The UK grid operates at a frequency of 50 Hz (50 cycles per second). If generation exceeds demand, frequency rises. If demand exceeds generation, frequency falls. There is no inventory, no buffer stock, no warehouse. Every electron consumed must be generated at that exact moment.

Why does imbalance change frequency? Frequency is the rotational speed of generators. When more power is generated than consumed, the turbines spin faster—frequency rises. When demand exceeds generation, turbines slow down—frequency falls. Grid frequency is the real-time speedometer of system balance.

This real-time balance is visible in the grid frequency itself. The National Energy System Operator (NESO)—an independent public body that took over from National Grid ESO in 2026—monitors frequency second by second, and keeps it within a tight operational band—typically 49.5 to 50.5 Hz under normal conditions. If frequency drops below 49 Hz, automatic systems start disconnecting large industrial users to prevent a cascading blackout. If it rises above 51 Hz, generators begin tripping offline to protect their turbines. In extreme cases, the grid can collapse entirely—a "blackout"—if the imbalance isn't corrected within seconds.

International note: Most of Europe operates at 50 Hz like the UK, while the US and most of North America operate at 60 Hz. The difference is historical (early manufacturers converged on different standards), not technical superiority. For traders, frequency choice doesn't change market economics, but it affects inertia dynamics and generator design—relevant when evaluating cross-border projects or equipment compatibility.

Why can't we just store electricity? Storage does exist—batteries are growing rapidly in the UK—but the scale is tiny compared to demand. The UK uses roughly 30–40 GW of power continuously. Even large battery installations can discharge for only a few hours before they're empty. There's no equivalent to the vast underground gas storage facilities or the millions of barrels of oil sitting in tanks. Electricity is produced, transmitted, and consumed in the same instant.

Gas, by contrast, is *flexible*. Pipelines hold gas under pressure—this is called **linepack**. You can push more gas into the system during low-demand periods and draw it out later. Underground storage facilities can hold billions of cubic metres for months. If a gas trader misjudges demand by an hour, the system absorbs it. If a power trader misjudges demand by a minute, frequency drops and alarms go off in the National Grid control room.

The Physical System: Generation to Consumption

But where does this instant balance happen? What system must be kept synchronised? To understand that, you need to see the physical path electricity takes from power station to home. The grid has four layers:

1. **Generation:** Power plants (wind farms, gas turbines, nuclear reactors) produce electricity.
2. **Transmission:** High-voltage wires (275 kV and 400 kV in the UK) carry power over long distances—from a Scottish wind farm to London, for example. High voltage minimizes energy losses over long distances. This is the "motorway" of electricity. Wholesale trading happens at this level.
3. **Distribution:** Lower-voltage networks (11 kV down to 230 V) deliver power from substations to homes and businesses. This is the "local road network."
4. **Consumption:** Homes, factories, data centres—anything that uses electricity.

Traditionally, this was a one-way flow: large central power stations fed the transmission grid, which fed the distribution network, which fed consumers. But that's changing. Rooftop solar panels and batteries mean power can now flow *upward*—from a home's solar array back into the distribution network, and sometimes into the transmission grid. This complicates balancing in two ways: the system operator must now coordinate not just a few hundred large generators but potentially millions of small distributed sources, and crucially, NESO cannot directly see most behind-the-meter generation, making demand forecasting significantly harder.

The key point for traders: wholesale energy markets operate at the **transmission level**. When you trade "UK power," you're trading the right to inject or withdraw power from the high-voltage grid. Distribution is a separate, regulated business. Gas works similarly: wholesale trading happens at the **National Balancing Point** (a virtual hub representing the high-pressure transmission system), not at individual homes or factories.

Across all four layers of this system, balance must hold every millisecond. Because electricity has no storage buffer, someone must orchestrate this in real time. That someone is the system operator.

2. System Operators: The Invisible Conductor

System operators don't trade. They don't take positions. They don't speculate on prices. But they are arguably the most important players in energy markets, because they ensure the physical system works. Think of them as air traffic controllers—they don't fly the planes, but they coordinate everyone else to prevent collisions.

In power markets, the system operator takes direct control: instructing generators to ramp up or down second by second. In gas markets, the model is more relaxed: commercial participants balance their own books, and the system operator steps in only when things go wrong. The difference comes back to time: electricity has no buffer, so control must be centralised. Gas has storage, so responsibility can be distributed.

NESO: The Power Balancer

NESO (National Energy System Operator) is the independent public body responsible for keeping the lights on in Great Britain. Created from the former National Grid ESO in 2026, NESO has an expanded remit covering both electricity system operation and strategic gas planning. Its core job remains pure coordination: match supply to demand, second by second, and keep frequency at 50 Hz.

NESO operates two main mechanisms:

The Balancing Mechanism (BM) is the real-time control room. Every large generator and some large consumers are connected to the BM. Throughout the day, NESO issues instructions: "increase output by 50 MW," "decrease by 100 MW." Generators bid prices to go up or down, and NESO accepts the cheapest bids that restore balance. If wind suddenly drops or a power plant trips offline, the BM responds within 5–10 minutes via dispatch instructions. Traders watch BM instructions closely—they reveal whether the system is **long** (supply exceeds demand, creating downward price pressure) or **short** (demand exceeds supply, creating upward price pressure) in real time.

But here's the critical detail most new traders miss: **you don't just pay the BM price—you pay the Imbalance Price.**

Every 30 minutes (the settlement period), NESO calculates two prices from all the BM actions it took:

- **System Buy Price (SBP):** If you're *short* (you sold more power than you generated or bought), you pay SBP to buy back your shortfall. This is derived from the most expensive *upward* BM actions NESO took in that period.
- **System Sell Price (SSP):** If you're *long* (you generated or bought more than you sold), you receive SSP for your surplus. This is derived from the cheapest *downward* BM actions NESO took.

GB-specific note: Since the P305 reform in November 2015, GB operates a "single imbalance pricing" regime where SSP and SBP effectively converge to a single system price for most volumes. While both prices are still calculated and published by ELEXON, participants are typically settled at one price rather than facing the dual pricing seen in some other markets. This simplifies imbalance settlement but doesn't change the fundamental risk: being out of balance during system stress is expensive.

NIV chasing: Some sophisticated participants—particularly batteries and flexible assets—deliberately run imbalanced positions, attempting to profit by predicting the system price. This strategy, known as "NIV chasing" (Net Imbalance Volume chasing), involves forecasting whether the system will be long or short and positioning accordingly. For example, a battery might discharge without fully hedging if it expects a high system price (system short), accepting the imbalance volume at a profitable rate. This shows that the GB market is *not* purely centrally controlled—even BM-registered assets have agency. If their bids and offers are extreme enough relative to system needs, they can choose their dispatch level and profit from imbalance rather than being forced to balance.

Why is imbalance exposure punitive? Because SBP and SSP are calculated from the *marginal* cost of balancing actions, not the average. During normal conditions, SBP might be close to the day-ahead price. But during system stress—low wind, a large generator trips offline, interconnector failure—NESO must call on expensive peaker plants or emergency reserves. SBP can spike to £200–300/MWh or higher.

Example of imbalance risk:

- You sell 100 MWh day-ahead at £80/MWh (your revenue: £8,000)
- Your wind farm generates only 80 MWh (you're short 20 MWh)
- Day-ahead cleared at £80/MWh, but during that settlement period, low wind forced NESO to dispatch expensive OCGTs
- SBP settles at £220/MWh (system stress)
- You pay: 20 MWh × £220/MWh = £4,400 for your imbalance
- Your profit: £8,000 (revenue) − £4,400 (imbalance) = £3,600

If you'd accurately forecast your generation and only sold 80 MWh day-ahead, you'd have earned £6,400 (£80 × 80 MWh) with no imbalance. Instead, you earned £3,600—a £2,800 penalty for being 20 MWh short during a stress event.

This is why accurate forecasting and position management are *existential* for power traders. Day-ahead prices are predictable. Imbalance prices during system stress are not. A single bad forecast on a high-wind day that doesn't materialize can wipe out a week's profits.

Ancillary services keep the system stable beyond just energy balance. **Frequency response** providers inject or absorb power automatically when frequency deviates, typically within 1 second. **Reserve** services provide backup capacity that can be called upon within minutes. **Inertia**—the spinning mass of large turbines—resists sudden frequency changes and gives NESO time to react.

As the grid shifts toward wind and solar (which provide no inertia), these services are becoming more valuable and more expensive.

Most energy trading happens in the **day-ahead market**, where participants buy and sell power for delivery tomorrow. But day-ahead is based on forecasts—and forecasts are never perfect. The Balancing Mechanism handles reality: when wind drops unexpectedly or a plant trips offline, NESO buys or sells power in real time to keep frequency at 50 Hz. As we've seen, getting your forecast wrong means paying imbalance prices that can be far more expensive than the day-ahead market.

Gas Shippers and the Gemini System

Gas works differently. There is no central dispatcher telling producers to turn on or off. Instead, commercial entities called **shippers** manage their own portfolios. A shipper might supply gas to homes, operate storage, or trade on behalf of a power plant. Each shipper is responsible for balancing its own inputs (gas entering the system) and outputs (gas leaving the system). If they take more gas than they nominated, they pay an imbalance charge. If they take less, they get paid (or charged, depending on system conditions).

The **Gemini system** is the IT platform where shippers submit nominations—forecasts of how much gas they'll inject or withdraw at each entry and exit point. Nominations are updated throughout the gas day (which runs 5am to 5am in the UK, not midnight to midnight). Shippers can trade with each other bilaterally to adjust their positions, or buy/sell on exchanges like ICE.

Why does this self-balancing model work for gas but not for power? *Storage*. If a shipper slightly over-nominates in the morning, the excess gas sits in the pipeline network as linepack or flows into an underground storage facility. The system can absorb imbalances for hours or even days. The system operator (National Gas Transmission) only intervenes when aggregate imbalances threaten network pressures or flows, issuing market-wide signals to correct the imbalance. For gas, coordination can be loose. For power, it must be tight.

> **Key Insight**
>
> Power requires a central conductor (NESO) because imbalances must be corrected in seconds. Gas allows distributed responsibility (shippers) because storage provides a buffer measured in hours or days.

3. The Merit Order & Dispatch

We've established that electricity must balance instantly and that NESO orchestrates this balance. But *how* does NESO decide which power plants to turn on and which to leave idle? The answer is the **merit order**—a ranking of every available generator from cheapest to most expensive. NESO dispatches plants starting from the bottom of this stack and works upward until demand is met. The last plant needed—the **marginal plant**—sets the price that everyone receives.

This isn't arbitrary. It's economically efficient: you meet demand at the lowest possible cost by using the cheapest generation first. But it also creates the fundamental dynamic of power trading: if you can predict which plant will be marginal, you can predict the clearing price.

Caveat for GB traders: The merit order stack described in this section is a *teaching tool* that explains the economic principles of price formation. In reality, GB operates a **self-dispatch** market: generators decide themselves what to run based on their contracts and price expectations, and NESO only adjusts these schedules via Balancing Mechanism bids and offers. When NESO does act, it doesn't simply choose by price—it considers network constraints, ramp rates, minimum stable levels, and reserve requirements. The simplified stack shown here helps you understand *why* prices form as they do, but actual GB dispatch is more nuanced. For algorithmic desks building GB strategies, you must model self-dispatch behavior, BM bid/offer curves, and constraint-driven actions, not just a pure cost stack.

To understand how this works in practice, we need to start with the demand side. What does UK electricity demand actually look like, hour by hour? How much does it vary between winter and summer, weekday and weekend? Once we understand the demand patterns traders must forecast, we'll see how NESO meets that demand by dispatching plants in merit order.

Understanding Demand: The 24-Hour Pattern

Before we build the merit order stack, let's understand what drives electricity demand in the UK. Unlike oil or wheat—where buyers can wait for a better price, or sellers can store inventory—electricity must be consumed the instant it's generated. When you flip a light switch at 6 PM, the lights come on regardless of whether power costs £60/MWh or £300/MWh. Demand is almost completely **inelastic** in the short term: it doesn't respond to price hour-to-hour.

Instead, demand follows predictable daily patterns driven by human behaviour. Figure 1 shows actual GB electricity demand over a typical 24-hour period, using recent National Grid ESO data. The shape is remarkably consistent day after day—but the absolute levels vary dramatically between winter and summer.

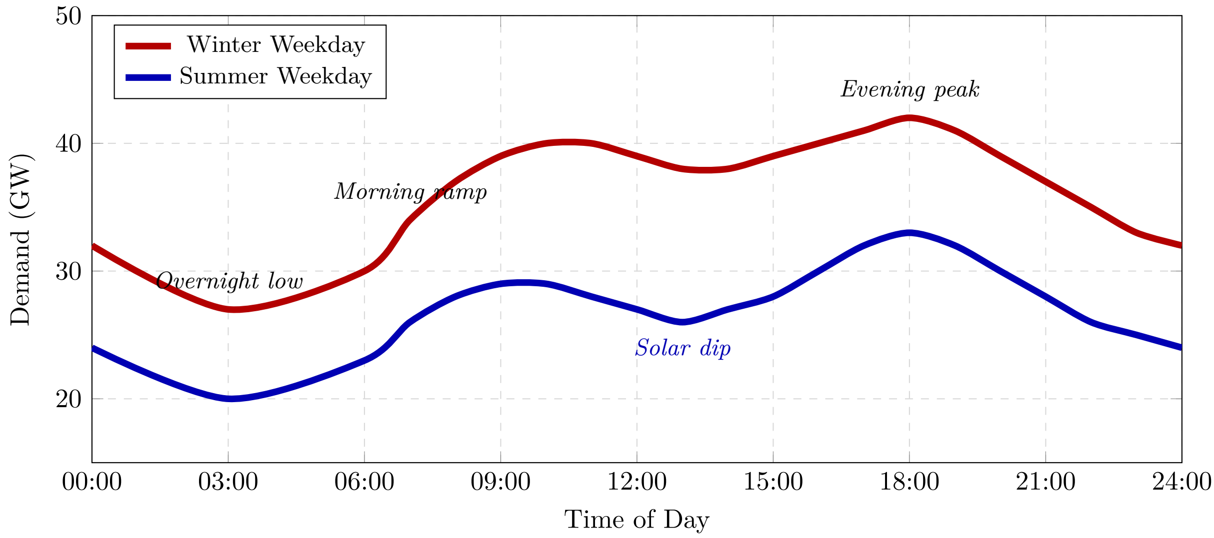

Figure 1: GB electricity demand over 24 hours: winter vs summer weekday. The shape is consistent—overnight trough, morning ramp, evening peak—but winter demand is 10–15 GW higher due to heating. Summer shows a midday dip as rooftop solar meets part of demand behind the meter. Data: National Grid ESO historic demand (2021–2023 typical ranges).

What drives this shape? Three components:

- **Residential:** Heating (biggest driver in UK winter), cooking, lighting, appliances. Peaks when people wake (6–9 AM) and return home (5–9 PM).
- **Commercial:** Offices, shops, restaurants. Drives weekday demand during business hours (8 AM–6 PM). Weekends are 5–8 GW lower all day.
- **Industrial:** Factories, data centers, manufacturing. More constant, running 24/7, but many large industrials (steel smelters, aluminum refiners) shift load when price signals justify it.

The daily cycle:

- **3 AM (overnight low):** 20–27 GW. Most people asleep, minimal commercial activity. Only essential services and some industrial processes.
- **6–9 AM (morning ramp):** Demand climbs sharply as people wake, heat homes, make breakfast, commute. Offices and shops open. Can rise by 10–15 GW in three hours.
- **12–2 PM (midday):** Demand plateaus at 38–40 GW (winter) or 26–29 GW (summer). In summer, rooftop solar meets part of residential and commercial load, creating a visible dip in grid demand.
- **5–9 PM (evening peak):** The "big one." People return home, cook dinner, turn on heating and lights, watch TV. Winter peaks can hit 42–45 GW; summer peaks typically 30–35 GW. This is the highest-priced period of the day.
- **10 PM–3 AM (wind-down):** Demand falls back toward the overnight trough as people go to bed and commercial premises close.

Seasonal variation: UK electricity demand is winter-peaking because heating dominates. Winter evening peaks can reach 45–50 GW on very cold days; summer peaks rarely exceed 35 GW. The UK has much lower summer peaks than the US because air conditioning penetration is far lower (only 5% of UK homes vs 90% of US homes). Week-to-week, demand varies with temperature: a cold snap can add 5–8 GW to evening peaks.

Weekly and event-driven patterns: Weekdays peak 5–8 GW higher than weekends due to commercial and industrial activity. Bank holidays behave like Sundays. Major events create sudden spikes: a World Cup final can add 2–3 GW as millions switch on TVs and kettles at half-time.

Why this matters for trading: Electricity demand is *predictable in shape* but *variable in magnitude*. You know there will be an evening peak—but will it be 38 GW or 45 GW? That 7 GW difference determines whether cheap CCGT plants are marginal (clearing price £60/MWh) or expensive old CCGT and peakers are needed (clearing price £80–100/MWh). Forecasting demand accurately—accounting for weather, day of week, holidays, and unexpected events—is the foundation of power trading.

The economic point: Unlike most commodities, electricity demand doesn't respond to price in real time. Consumers pay fixed tariffs; they don't see or react to wholesale prices hour-to-hour.

This means demand is a near-*vertical line* at any given moment—it must be met regardless of cost. When demand rises from 35 GW to 42 GW, NESO can't wait for price-sensitive consumers to reduce consumption. It must immediately dispatch more expensive plants. This inelastic demand, combined with the step-function supply of the merit order (which we'll explore next), creates the extreme price volatility characteristic of power markets: small demand changes cause large price swings.

Now let's see how NESO meets this varying demand by dispatching power plants in merit order.

Building the Dispatch Stack

Imagine GB demand is currently 35 GW. NESO looks at all available generation and ranks it by price. Before we see the numbers, let's quickly define the main plant types:

- **Nuclear:** Large steam turbines powered by nuclear fission. Very low fuel cost, runs constantly (baseload).
- **Wind and Solar:** Renewables with zero fuel cost (once built). Output depends on weather, not economics.
- **CCGT (Combined Cycle Gas Turbine):** Gas-fired plants that use both a gas turbine and a steam turbine, recovering waste heat for efficiency. Modern CCGTs are highly efficient (49–52%); older ones less so (45–48%). This efficiency gap explains why we distinguish "new" and "old" CCGT in the merit order.
- **OCGT (Open Cycle Gas Turbine):** Simpler gas turbines (like jet engines) with no heat recovery. Much less efficient (30–35%) but very fast to start. Used as "peakers"—they run only during the highest-demand, highest-price periods.

Here's a simplified snapshot using 2026 numbers:

Plant Type	Available Capacity (GW)	Price (£/MWh)
Nuclear	5	£8
Wind	12	£0
Solar	0	£0
CCGT (new)	20	£60
CCGT (old)	15	£75
OCGT (peaker)	5	£95

Note: Solar is zero because it's night-time in this example. Wind is available at full capacity (12 GW), but this varies hour by hour depending on weather.

NESO "stacks" these plants from cheapest to most expensive and dispatches them in order:

1. **Wind** (£0/MWh): Dispatch all 12 GW. *Cumulative: 12 GW.*
2. **Nuclear** (£8/MWh): Dispatch all 5 GW. *Cumulative: 17 GW.*
3. **CCGT (new)** (£60/MWh): Need 35 GW total, so dispatch 18 GW of the 20 GW available. *Cumulative: 35 GW.*

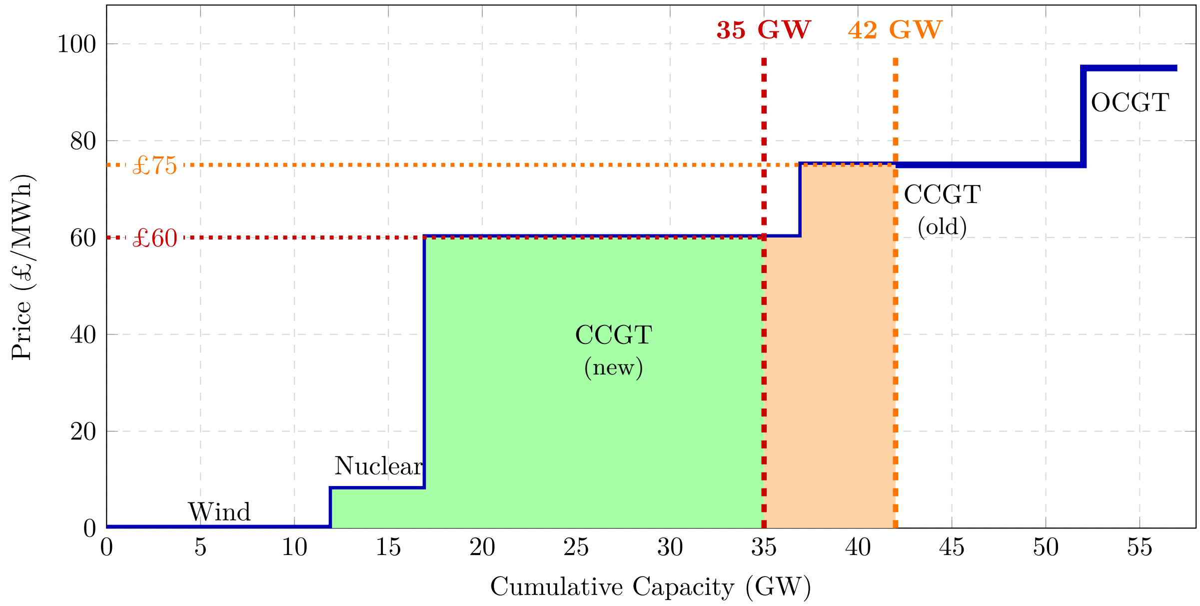

Figure 2: Merit order dispatch stack showing how demand determines the clearing price. At 35 GW demand (red), CCGT (new) is marginal and sets price at £60/MWh (green shaded area dispatched). At 42 GW demand (orange), CCGT (old) becomes marginal, pushing price to £75/MWh.

Demand is met. The **marginal plant** is the new CCGT running at £60/MWh. That becomes the **clearing price**—the price paid to *all* generators, including the wind farms that bid £0 and the nuclear plant that bid £8. Everyone receives £60/MWh.

Why uniform pricing? It seems wasteful to pay wind £60/MWh when they bid £0. But this is economically efficient for three reasons: (1) The marginal plant must earn exactly its cost, or it won't run when needed—and we *need* that plant to meet demand. (2) If you paid different generators different prices, they'd game the system by withholding capacity to drive prices up. (3) The "windfall" profits to cheap generators incentivize investment in low-cost generation—exactly what society wants. Uniform pricing aligns private incentives with system efficiency. Every plant lower in the stack earns a profit (their "inframarginal rent"), which funds future investment.

Now imagine demand rises to 42 GW. NESO must dispatch more capacity:

1. Wind: 12 GW (*cumulative: 12 GW*)
2. Nuclear: 5 GW (*cumulative: 17 GW*)
3. CCGT (new): All 20 GW now (*cumulative: 37 GW*)
4. CCGT (old): 5 GW of the 15 GW available (*cumulative: 42 GW*)

The marginal plant is now the old CCGT at £75/MWh. The clearing price jumps from £60 to £75—a 25% increase—because demand rose by just 7 GW (20%). This is why power prices are so volatile: small changes in demand can push the system up or down the merit order stack, crossing step-functions in price.

What determines a plant's price? For fossil fuel plants, it's the cost of fuel plus carbon. A modern CCGT burning gas at 70 p/therm with carbon at £50/tonne has a variable cost around £60/MWh (we'll derive this in the Spark Spread guide). An older, less efficient CCGT might be £75/MWh. Wind and solar have zero fuel cost—once built, the marginal cost of generation is effectively zero. Nuclear fuel is so cheap (£5–10/MWh) that nuclear plants run constantly unless they're offline for maintenance.

Carbon price volatility: Carbon prices (UKAs in the UK, EUAs in Europe) are set by policy—auction volumes, cap adjustments, regulatory changes—making them highly volatile. In recent years, UKAs have ranged from £30–80/tonne. This matters because carbon and gas prices often move independently. When gas is cheap but carbon is expensive (say, 50 p/therm and £80/tonne), gas plants become relatively expensive in the merit order. When gas is expensive but carbon is cheap (100 p/therm and £30/tonne), the opposite occurs. Traders constantly model this "two-dimensional" cost surface. We'll explore carbon markets in detail in the next guide.

UK-specific carbon costs: UK generators face a two-component carbon cost structure. They pay the **UK Allowance (UKA)** price—a traded carbon permit that varies daily based on market conditions—*plus* the **Carbon Price Support (CPS)**, a fixed tax set by government (£18/tonne as of 2025). Total carbon cost for a UK CCGT is therefore **UKA + CPS** (roughly £58–78/tonne when UKAs trade at £40–60/tonne). This dual structure means UK CCGT costs can diverge from EU plants even with similar gas prices, because EU plants only pay the EUA (no equivalent CPS tax). For traders, UKA risk is market risk (hedge with UKA futures), while CPS risk is political risk (changes in government budgets, typically announced annually).

International note: The UK operates a single GB-wide clearing price. The US uses **nodal pricing** (Locational Marginal Pricing, LMP)—thousands of location-specific prices that explicitly reflect congestion and losses at each grid node. Continental Europe uses **zonal pricing** via EUPHEMIA, where countries/regions are coupled zones with implicit cross-border capacity allocation. Post-Brexit, GB trades with Europe via explicit interconnector capacity auctions rather than implicit coupling. For algorithmic desks, US nodal markets require network topology modeling and constraint-based strategies, while UK/EU focus more on day-ahead/intraday auction optimization and balancing mechanism spreads.

The Must-Run Paradox: Why Expensive Plants Run Alongside Free Wind

Pure merit order assumes the cheapest generation always runs first. But in reality, NESO sometimes pays expensive gas plants to run *even when there's enough cheap wind and solar to meet demand*. This seems wasteful—why pay for £60/MWh CCGT when you have free wind available?

The answer: **the grid needs more than just energy—it needs frequency stability.** Recall that grid frequency (50 Hz) is the rotational speed of generators. When demand suddenly spikes or a plant trips offline, frequency starts to drop. What stops it from crashing? **Inertia**—the physical spinning mass of large turbines. Heavy generators (steam turbines in gas plants, nuclear, or the now-retired coal plants) act like flywheels: they resist changes in rotational speed, giving NESO precious seconds to dispatch balancing actions.

Wind and solar provide *zero* inertia. They're **inverter-based resources**—they convert DC to AC electronically, with no mechanical spinning mass.[a] A grid running purely on wind and solar would have almost no natural resistance to frequency changes. A sudden demand spike or generator trip could cause frequency to crash before NESO could react, triggering a blackout.

This is why NESO sometimes runs "must-run" plants:

- **Inertia:** Keep synchronous generators spinning to provide mechanical inertia and system strength.
- **Voltage support:** Maintain local voltage stability (wind/solar connected via long transmission lines can't always provide this).
- **Black start capability:** Ensure some plants can restart the grid if it collapses (most renewables can't).

The 2026 context: With the last UK coal plant closing in September 2024, the grid lost its largest source of inertia. Coal plants were massive: heavy steam turbines spinning at 3,000 RPM, providing enormous mechanical mass. Now, as wind and solar capacity grows (36 GW wind, 25 GW solar in 2026), the system has less natural inertia than ever. NESO must increasingly pay gas plants to run just for their spinning mass, even when their energy isn't needed.

This creates a trading opportunity: **ancillary services markets**. Plants providing inertia, frequency response, and voltage support earn revenue beyond energy sales. These services are becoming more valuable and more expensive as the renewable share grows. We'll explore this further in the next guide, but for now, understand that "must-run" plants can bid above their marginal cost in the Balancing Mechanism, knowing NESO must accept them for stability reasons. This is why BM prices occasionally spike far above what pure merit order would predict.

[a] Technically, emerging **grid-forming inverters** can provide *synthetic* inertia by electronically emulating the behavior of spinning mass. As of 2026, most UK wind and solar use **grid-following** inverters that provide zero inertia. Grid-forming technology is being deployed in some large battery installations and is expected to grow, but conventional synchronous generators (CCGT, nuclear) remain the primary source of system inertia.

Time Variation: A Day in the Life

The merit order is not static. Demand varies hour by hour, and so does the marginal plant. A typical winter weekday in the UK follows a predictable pattern:

- **3 AM (Overnight low):** Demand might be 28 GW. Wind (12 GW) plus nuclear (5 GW) plus some CCGT (11 GW) covers it. Clearing price: £60/MWh (new CCGT marginal).
- **8 AM (Morning peak):** Offices open, trains run, kettles boil. Demand hits 38 GW. More CCGT comes online, and perhaps some old CCGT starts. Clearing price: £75/MWh.
- **2 PM (Afternoon dip):** Demand drops to 32 GW as the morning rush ends. Back to mostly new CCGT. Clearing price: £60/MWh.
- **6 PM (Evening peak):** Peak demand. Everyone gets home, turns on heating, cooking, lights. Demand hits 42 GW. Old CCGT fully dispatched, maybe even some OCGT peakers fire up. Clearing price: £75–95/MWh.

This daily "shape" creates trading opportunities. A **baseload** contract commits to constant delivery 24/7 at a single fixed price agreed at signing, but the actual clearing price swings between £60 and £95 hour-to-hour. A **peak** contract covers only high-demand hours (typically 7 AM–7 PM), also at a fixed price. Traders who can take positions in peak versus off-peak periods can capture the spread between these products and the volatile hourly clearing prices.

Modern dispatch as optionality: While baseload contracts still exist, modern CCGT operators increasingly treat dispatch as a continuous optimization problem. Rather than "running baseload," they decide every settlement period (30 minutes) whether to generate based on the spark spread (power price minus gas cost). Economically, this is equivalent to holding a **call option** on the spark spread—the plant only "exercises" (runs) when the option is in the money. We'll explore this optionality framework and its trading implications in the Spark Spread guide.

Now add weather into the mix. On a calm night with no wind, that 3 AM demand of 28 GW might need to be met entirely by nuclear and CCGT—wind contributes zero. Suddenly, even at low demand, you're dispatching old CCGT. Clearing price: £75/MWh in the middle of the night. Conversely, on a very windy afternoon, wind might be generating 30 GW. Combined with solar (say, 10 GW) and nuclear (5 GW), you have 45 GW of zero-marginal-cost generation. If demand is only 32 GW, some wind farms must be *turned off*. The clearing price can drop to zero or even go negative.

Why Negative Prices?

Negative prices occur when supply far exceeds demand and inflexible generators (nuclear, some renewables with subsidy contracts) prefer to pay to stay online rather than shut down. Reasons: (1) Shutdown/restart costs exceed the negative price, (2) subsidy payments (like CfDs) mean they still profit even at negative prices, (3) wind farms avoiding turbine wear from cycling. Negative prices signal: "We have too much power—please consume more or store it." Batteries and flexible demand profit by getting paid to absorb excess generation. **Why don't retail consumers benefit?** Most households and businesses are on fixed tariffs set months in advance—they don't see real-time wholesale prices. Only large industrials with half-hourly metering and time-of-use contracts, or battery operators trading in wholesale markets, can capture negative price opportunities.

The Renewable Transformation: How We Got Here

Before 2010, the merit order was simpler. Nuclear provided baseload. Coal and gas competed in the middle. The marginal plant was almost always a fossil fuel—either coal or gas, depending on relative fuel prices. Power prices were relatively stable because fuel prices moved slowly and predictably.

Then wind and solar scaled up. By 2026, the UK has roughly 36 GW of wind capacity and 24 GW of solar. On a windy, sunny spring afternoon, renewables can cover nearly all demand. Because their marginal cost is zero, they sit at the bottom of the merit order and push fossil fuels out. Gas plants that would have run for 20 hours a day now run for 10. Some days, they don't run at all.

This has three major consequences for trading:

1. Price volatility has increased. When renewables are generating heavily, clearing prices drop toward zero (or below). When the wind stops, prices spike as expensive gas plants are called upon rapidly. The same hour of the day might clear at £20/MWh one week and £100/MWh the next, depending entirely on weather.

2. Forecasting became critical. In the old world, you could predict tomorrow's price reasonably well by looking at gas prices. Now you need wind forecasts, solar irradiance models, and interconnector flow predictions. A 5 GW error in your wind forecast translates directly into being wrong about which plant is marginal—and therefore wrong about price by £20–40/MWh.

3. Gas plants became "swing" assets. They no longer run 24/7. They turn on and off multiple times per day, filling the gaps when renewables dip. This is why traders running gas-fired plants obsess over the **Spark Spread**—the profit margin between power prices and gas costs. If the Spark Spread is positive, you bid your plant into the market. If it's negative, you sit idle and hope it improves. Predicting the Spark Spread means predicting whether your plant will be marginal, which means predicting demand, renewable output, and interconnector flows.

The merit order is no longer a stable stack—it's a dynamic, weather-driven system where the marginal plant can shift violently within minutes.

Quick Check

If UK demand is 35 GW and you have:

- Nuclear: 5 GW at £8/MWh
- Wind: 12 GW at £0/MWh
- CCGT (new): 20 GW at £60/MWh
- CCGT (old): 15 GW at £75/MWh

Which plants run? What is the clearing price? What if demand rises to 42 GW?

Hint: Stack from cheapest to most expensive. At 35 GW, you'll dispatch Wind (12 GW) + Nuclear (5 GW) + CCGT new (18 GW) = 35 GW. The marginal plant (CCGT new at £60/MWh) sets the clearing price.

Python Implementation: Dispatch Simulator

Here's how you'd automate the merit order dispatch logic. This function takes a list of available plants and a demand level, then returns which plants run and what the clearing price is.

```python
def dispatch_merit_order(plants, demand_mw):
    """
    Simple merit order dispatch simulator.

    Args:
        plants: List of tuples (name, capacity_mw, price_per_mwh)
        demand_mw: Total demand to meet in MW

    Returns:
        dict with:
            - dispatched: list of (name, dispatched_mw, price)
            - clearing_price: marginal plant price
            - total_cost: sum of (dispatched_mw * marginal_price)
            - unmet_demand: MW of demand not covered (if insufficient capacity)
    """
    # Sort plants by price (merit order)
    sorted_plants = sorted(plants, key=lambda p: p[2])

    dispatched = []
    remaining_demand = demand_mw
    clearing_price = 0

    for name, capacity_mw, price in sorted_plants:
        if remaining_demand <= 0:
            break

        # Dispatch as much as we can from this plant
        dispatched_mw = min(capacity_mw, remaining_demand)
        dispatched.append((name, dispatched_mw, price))

```

```python
        remaining_demand -= dispatched_mw
        clearing_price = price  # Marginal plant sets the price

    # Calculate total cost: all dispatched MW paid at clearing price
    total_mw_dispatched = sum(mw for _, mw, _ in dispatched)
    total_cost = total_mw_dispatched * clearing_price

    return {
        "dispatched": dispatched,
        "clearing_price": clearing_price,
        "total_cost": total_cost,
        "unmet_demand": max(0, remaining_demand),
    }


# Example: UK dispatch for 35 GW demand
uk_plants = [
    ("Nuclear", 5000, 8.0),
    ("Wind", 12000, 0.0),
    ("CCGT_new", 20000, 60.0),
    ("CCGT_old", 15000, 75.0),
    ("OCGT", 5000, 95.0),
]

result = dispatch_merit_order(uk_plants, demand_mw=35000)

print(f"Clearing price: GBP {result['clearing_price']:.2f}/MWh")
print(f"Total cost: GBP {result['total_cost']:,.0f}")
print("\nDispatched plants:")
for name, mw, price in result["dispatched"]:
    print(f"  {name}: {mw:,} MW at GBP {price}/MWh")

# Output:
# Clearing price: GBP 60.00/MWh
# Total cost: GBP 2,100,000
#
# Dispatched plants:
#   Wind: 12,000 MW at GBP 0/MWh
#   Nuclear: 5,000 MW at GBP 8/MWh
#   CCGT_new: 18,000 MW at GBP 60/MWh

# Now try with 42 GW demand:
result_high = dispatch_merit_order(uk_plants, demand_mw=42000)
print(f"\nHigh demand (42 GW):")
print(f"Clearing price: GBP
    {result_high['clearing_price']:.2f}/MWh")
for name, mw, price in result_high["dispatched"]:
    print(f"  {name}: {mw:,} MW at GBP {price}/MWh")

# Output:
```

```
# High demand (42 GW):
# Clearing price: GBP 75.00/MWh
#   Wind: 12,000 MW at GBP 0/MWh
#   Nuclear: 5,000 MW at GBP 8/MWh
#   CCGT_new: 20,000 MW at GBP 60/MWh
#   CCGT_old: 5,000 MW at GBP 75/MWh
```

A 7 GW increase in demand (20%) pushes the clearing price up by 25%. This simple function captures the core logic of power markets: cheapest first, marginal plant sets price, everyone gets paid the same.

Checkpoint: Before Moving On

Test your understanding:

1. Why does the marginal plant set the price for *all* generators, not just itself?
2. If wind output increases by 5 GW, what happens to the clearing price? (Hint: think about which plant drops out of the merit order.)
3. Why are power prices more volatile than gas prices?

4. Gas: The Flexible Counterpart

We've spent significant time on power because it's the hardest case. Gas is simpler to understand once you've grasped the power system—but it operates on fundamentally different principles. Where power requires instant, centralized orchestration, gas allows for distributed, slower-paced coordination. The key is storage.

The National Balancing Point: A Useful Fiction

When UK traders talk about "NBP gas," they're referring to the **National Balancing Point**. Where is it? Nowhere. You can't visit NBP on a map. It's not a physical location—it's an *accounting construct*, a virtual hub that exists only on paper.

In physical reality, gas enters the UK transmission system at specific terminals: Milford Haven on the Welsh coast (LNG imports), Bacton in Norfolk (pipeline imports from Europe), St Fergus in Scotland (North Sea production). It exits at power plants, industrial users, and homes across the country. But for trading purposes, all gas entering the system is deemed to enter "at NBP," and all gas leaving is deemed to exit "at NBP."

Why create this fiction? Because it dramatically simplifies trading. Instead of negotiating delivery at hundreds of specific points, traders can buy and sell a standardized product—"NBP gas"—knowing that the physical system will handle the complexity of actual delivery. The NBP creates a liquid, centralized market where prices are transparent and volumes are high. It's the UK's gas trading hub, analogous to Henry Hub in the US or TTF (Title Transfer Facility) in the Netherlands.

This fiction holds most of the time—but not always. When physical infrastructure hits capacity limits (pipelines full, regional demand spikes), local prices diverge from the NBP reference. You might hedge at NBP but pay a premium (or receive a discount) for physical delivery at your location. This is **basis risk**, and we'll explore it in detail in Section 5.

Key Insight

NBP is to UK gas what "par" is to bonds—a reference point for pricing, not a physical location. You can't visit NBP, but you can trade millions of therms there every day.

Units & Conversions: Gas to Power

Gas markets quote in **pence/therm** (UK NBP), while power markets quote in **£/MWh**. To compare them:

$\boxed{1\text{ MWh} \approx 34.12\text{ therms}}$ (the "magic number")

Quick conversion formula: p/therm × 0.3412 = £/MWh

Example: Gas trading at 70 p/therm is equivalent to £23.88/MWh in energy terms:

$$70 \times 0.3412 = 23.88$$

Why this matters: When calculating the **Spark Spread** (a CCGT plant's profit margin), you need to compare gas input costs (p/therm) with power output revenue (£/MWh). This conversion is the first step. We'll cover Spark Spread calculations in detail in the dedicated guide, but understanding the energy equivalence now builds intuition for how gas and power markets link.

Inverse conversion: £/MWh ÷ 0.3412 = p/therm

International note: Gas hubs use different pricing units. NBP quotes in **pence/therm** (UK convention). TTF (Netherlands/Europe) quotes in **EUR/MWh**. Henry Hub (US) quotes in **USD/MMBtu**. When trading cross-hub spreads, you must normalize both energy units and currency: 1 MWh ≈ 3.412 MMBtu ≈ 34.12 therms. These conversions are critical for LNG arbitrage and cross-border gas positions.

Day-rate conventions and clock-change day quirks: Gas trading uses day-rate units that can create subtle errors:

- **NBP (UK):** Traded in **therms per day**. Straightforward, but remember the gas day runs 05:00–05:00 UTC (Section 4.2).
- **Henry Hub (US):** Traded in **MMBtu per day**. Gas day runs 09:00–09:00 Central Time.
- **Continental Europe (TTF, etc.):** Often quoted in **MW**, which is *ambiguous*—it can mean MWh per day *or* MWh per hour depending on context and counterparty convention. Always confirm which is meant, especially in nominations and confirmations.

Clock-change day impact: On daylight saving transitions, the "day" has 23 or 25 hours instead of 24, affecting energy delivery:

- **Germany:** Gas day starts at 06:00 CET. On the October clock-change day (clocks go back), Germany gets an *extra* hour, so a contract for "X MWh per day" delivers 25 hours worth of gas at X/24 MW average rate (total: $X \times 25/24$).

- **France:** Interprets the same contract differently, delivering exactly X MWh regardless of day length, so the hourly flow rate adjusts but total energy stays flat.

This creates a small volume mismatch on clock-change days if your system doesn't handle day-length explicitly. For ETRM systems, store delivery periods with explicit start/end UTC timestamps and day-length as a dimension, not just "date." Test your P&L calculations across DST boundaries to catch these off-by-one errors. We cover this in detail in the *Position & P&L* guide (Pitfall 4).

Shippers, Linepack, and Self-Balancing

Unlike power, where NESO dispatches every plant, gas operates on a **self-balancing** model. Commercial entities called **shippers** are each responsible for balancing their own portfolios. A shipper might be a power plant buying gas to burn, a supplier delivering gas to homes, or a trader speculating on price movements. Each shipper forecasts how much gas they'll inject into the system (from terminals or storage) and how much they'll withdraw (to customers or storage). These forecasts are called **nominations**, and they're submitted via the **Gemini system**.

Throughout the gas day (which runs 5 AM to 5 AM in the UK—chosen because it's the period of lowest demand, minimizing market disruption at the rollover), shippers can update their nominations and trade with each other to adjust positions. If a shipper ends the day having taken more gas than they nominated, they're charged for the imbalance. If they took less, they may be paid or charged, depending on whether the system was short or long overall.

SOFTWARE ENGINEERING PITFALL: Gas Day Offset

A notorious source of systematic errors in energy trading systems.
The UK gas day runs **05:00 to 05:00**, not midnight to midnight. When someone requests "Wednesday gas trades," they mean the gas day running **Wednesday 05:00 to Thursday 04:59**, not the calendar day.

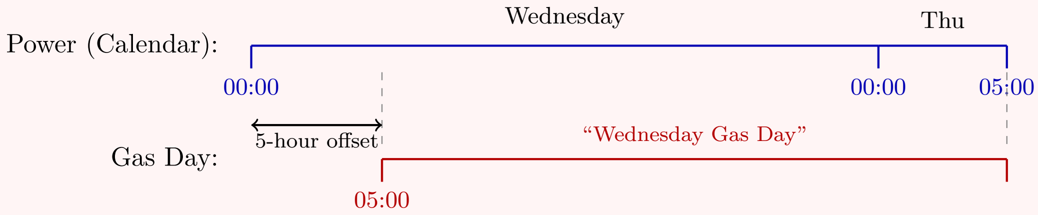

Why this breaks systems:

- Naively joining gas and power data by "day" creates a systematic 5-hour misalignment
- Daily P&L calculations will be wrong if you don't account for the offset
- Position reconciliations fail when aggregating across markets (gas vs power)
- Off-by-one errors when rolling up intraday data to daily summaries

The fix:

- Store delivery periods with explicit start/end timestamps in UTC
- Never assume "day" means calendar day—clarify gas day vs calendar day
- Test your join logic across gas day boundaries (especially around 05:00)
- Document your aggregation strategy when combining gas and power data

International note: The US gas day runs 09:00–09:00 Central Time. If you're aggregating North American and European positions, you have *two* different offset conventions to handle. See the *Position & P&L* guide (Pitfall 4) for detailed implementation guidance.

This model works because gas has flexibility that power does not. Pipelines themselves act as short-term storage—**linepack**. By varying the pressure in the pipes, you can hold more or less gas. If a shipper over-nominates slightly in the morning, that extra gas just increases the linepack. Later in the day, when demand rises, the system draws down that pressure. Underground storage facilities (like Rough, an offshore depleted gas field with roughly 1.5 bcm of working capacity) provide even longer time buffers—weeks or months.

The system operator (National Gas Transmission) monitors aggregate system balance and intervenes only when network pressures approach limits. Compared to power, where balance must be perfect every second, gas can tolerate imbalances for hours or even days.

Storage as a Time Buffer

Underground gas storage is the reason gas markets are less volatile than power markets. The UK has several storage facilities, including the Rough offshore field (1.5 bcm working capacity) and salt cavern sites like Stublach and Holford. These facilities inject gas during low-demand periods (typically summer) and withdraw during high-demand periods (winter). This seasonal arbitrage dampens price spikes: if winter gas prices soar, storage operators release gas, which increases supply and caps prices.

We won't dive deep into storage economics here—that's for the next guide on global energy flows. The key point for now: storage provides a time buffer. A gas trader who misjudges demand today can correct the error tomorrow by trading in the market or adjusting storage flows. A power trader has no such luxury.

5. Physical Constraints & Basis Risk

So far, we've talked about NBP as if it's a single, unified market where all gas trades at the same price, and the UK power grid as if it's one clearing price everywhere. In theory, that's true. In practice, physical infrastructure has limits, and those limits create price divergences. This is **basis risk**—the risk that the price at your physical location differs from the hub price you hedged against.

Pipeline Capacity Constraints

The UK gas transmission network is not infinitely flexible. Pipelines have maximum flow capacities. Gas enters the system at terminals like St Fergus in Scotland (North Sea production) and exits in the South (London, industrial users). If demand in the South exceeds the capacity of the pipelines running north-to-south, local prices in the South will rise above the NBP reference price. Conversely, if Scottish terminals are producing heavily but the pipes are full, Scottish prices might fall below NBP.

For a trader, this creates a problem. You've hedged your position at NBP—say, you've bought 10,000 therms at 75 p/therm. But your physical gas plant is in Kent, and on delivery day, the South is tight. You end up paying 78 p/therm to get gas delivered locally. Your hedge locked in NBP, but your physical cost is NBP plus a 3 p/therm "basis." On 10,000 therms, that's a £300 loss you didn't expect.

Basis risk is small most of the time in the UK gas market because the network is well-connected. But during extreme weather (cold snaps that max out demand) or infrastructure outages (a key pipeline offline), basis can spike.

Grid Congestion and Curtailment

Power has the same problem, but it's more severe because you can't store electricity. The UK's wind capacity is concentrated in Scotland—offshore farms in the North Sea and onshore farms in the Highlands. Demand is concentrated in England, particularly the South East. To move power from Scotland to England, you need transmission lines across the **B6 boundary**—the network bottleneck between Scotland and England.

When Scottish wind is generating heavily but demand in Scotland is low, power must flow south. If the B6 boundary hits its capacity limit (roughly 5–7 GW, depending on network conditions), the excess Scottish generation has nowhere to go. The grid becomes "constrained."

What happens? NESO pays Scottish wind farms to turn *off*. This is called **curtailment**. At the same time, NESO pays gas plants in England to turn *on* to meet southern demand. The result: Scotland has cheap (or negative) power prices, England has expensive power prices, and consumers pay for both the curtailment and the replacement generation.

The numbers are staggering. In 2025, Scottish wind farms were paid roughly £0.8–1.5 billion to curtail output—around 4 TWh of lost generation in the first half of the year alone. Some wind farms curtailed 37–50% of their potential output in peak months. On a single day in December 2025, the GB system paid £12.6 million in constraint costs: £4.2 million to turn off Scottish wind, and £8.5 million to buy replacement power in England. These costs are ultimately socialised—paid by consumers through network charges on their energy bills.

Why does this seemingly wasteful arrangement make sense? The B6 boundary is a *physical limit*. The transmission wires can only carry 5–7 GW. When Scottish wind generation exceeds Scottish demand plus B6 capacity, the excess power physically cannot be exported south. If generators kept producing, local frequency in Scotland would rise and the grid would become unstable. Curtailment isn't wasteful economics—it's necessary physics. The alternative (grid instability or collapse) is far more expensive. The long-term solution is building more transmission capacity, but that takes years and costs billions.

REMA and the zonal pricing debate: These massive constraint costs sparked a debate about whether the UK should move to **zonal pricing** (different prices in Scotland vs England, like the US nodal system or EU's multi-zonal structure). In 2025, the UK government's Review of Electricity Market Arrangements (REMA) ruled this out, choosing instead "Reformed National Pricing"—keeping a single GB-wide clearing price while using the Strategic Spatial Energy Plan (SSEP) and enhanced network planning to address constraints over time. For traders, this means curtailment and basis risk will remain features of the market, managed through constraint pricing and network charges rather than explicit zonal day-ahead prices.

For a wind farm owner in Scotland, this creates basis risk. The day-ahead market might clear at £80/MWh (the GB average). But if your farm is curtailed, you don't receive that. Instead, you receive the Balancing Mechanism bid-off price—possibly £60/MWh or lower, depending on your contract terms. Your hedge assumes you'll receive the GB clearing price. Your physical reality is £20/MWh less.

For traders, this matters enormously. If you're long Scottish wind, you're exposed to curtailment risk. If you're short English power, you're exposed to the cost of replacement generation when Scotland is constrained. Basis risk isn't a footnote—it's a structural feature of the market.

Try It Yourself

Your gas plant in Kent needs 10,000 therms for tomorrow. You've hedged at NBP for 75 p/therm. On delivery day:

- NBP spot price: 75 p/therm (your hedge is flat)
- Kent basis: +3 p/therm (pipeline from North is constrained)

What is your actual gas cost per therm? What is your total basis risk loss in pounds?

Answer: Physical cost = 78 p/therm. Total loss = 3 p/therm × 10,000 therms = 30,000 p = £300.

6. Market Coupling & Interconnectors

We've seen how physical constraints within the UK create basis risk—Scottish wind constrained by the B6 boundary, gas pipelines hitting capacity limits. But the UK doesn't operate as an isolated island. Subsea cables connect the GB power grid to France, Belgium, the Netherlands, Norway, and Denmark. Gas pipelines link NBP to European hubs. These **interconnectors** add another layer of physical constraint—and another source of trading opportunity.

These interconnectors allow energy to flow between markets, driven by price differentials. When power is cheap in France and expensive in the UK, electricity flows across the English Channel. When gas is cheap at TTF (the Dutch hub) and expensive at NBP, gas flows into the UK.

Interconnectors serve two purposes: they provide *physical flexibility* (the UK can import power when demand is tight or wind is calm) and they create *price convergence* (arbitrage flows push prices in connected markets toward each other).

Physical Links: Subsea Cables and Pipelines

As of 2026, the UK has roughly 11.7 GW of power interconnector capacity to Europe and Norway. The major links are:

Link	Destination	Capacity (MW)
IFA	France	2,000
IFA2	France	1,000
BritNed	Netherlands	1,000
Nemo Link	Belgium	1,000
NSL (North Sea Link)	Norway	1,400
Viking Link	Denmark	1,400

Each interconnector has a finite capacity—IFA can carry 2 GW, but no more. When that capacity is full, prices in the UK and France can diverge significantly, even though they're physically connected.

Gas interconnectors are fewer and simpler. The main pipeline link is BBL (Balgzand-Bacton Line), connecting the Dutch TTF hub to the UK's Bacton terminal. Gas can also flow via LNG tankers, which deliver liquefied gas to terminals like Milford Haven. We'll cover global gas flows in detail in the next guide.

Arbitrage: Flow Follows Price

The principle is simple: energy flows from the cheap market to the expensive market. If UK power is trading at £90/MWh and French power is at €80/MWh (roughly £68/MWh), traders will buy power in France and sell it in the UK, pocketing the £22/MWh spread. This is called **arbitrage**.

Power flows across the IFA interconnector from France to the UK. As French exports increase, French prices rise (less supply available domestically). As UK imports increase, UK prices fall (more supply arriving). The arbitrage continues until one of two things happens: either prices converge (the spread narrows to near zero), or the interconnector fills to capacity (2 GW for IFA).

Once capacity is full, the arbitrage stops—even if prices remain divergent. If the UK is very tight and France is very cheap, you might see UK power at £120/MWh and French power at £50/MWh, with IFA running flat-out at 2 GW. The physical constraint prevents further flows, so prices stay apart.

This dynamic matters for traders. If you're forecasting UK power prices, you need to know: what are French, Dutch, and Norwegian prices? Are the interconnectors likely to be full or empty? A 2 GW swing in imports can move UK clearing prices by £10–20/MWh, depending on where you are on the merit order stack.

A technical detail: loop flows and Kirchhoff's laws. Power doesn't flow in straight lines from source to destination. AC grids obey Kirchhoff's laws—current flows along all available paths inversely proportional to impedance. When the IFA interconnector flows from France to the UK at 2 GW, some power may physically "loop" through Belgium and the Netherlands due to grid topology, even though the commercial transaction is direct France-to-UK. For GB traders, this affects constraint costs when Scottish wind loops through European interconnectors back into England. For US nodal traders, loop flows are fundamental—congestion on one transmission line affects prices at distant nodes because power physically flows everywhere the grid is connected. In EU flow-based market coupling, these loop flow effects are explicitly modeled in the day-ahead auction algorithm (EUPHEMIA).

We won't go deep into cross-border hub trading here—that's for the next guide, where we'll explore how TTF, NBP, and global LNG markets connect. For now, the key point: markets are physically linked, prices influence each other, but capacity limits define how much convergence is possible.

Quick Check

UK power price: £90/MWh
French power price: €80/MWh (≈ £68/MWh at 1.18 exchange rate)
IFA interconnector capacity: 2 GW
Which direction does power flow? What limits the arbitrage?

7. Summary: Why This Matters for Trading

You've now seen the physical reality beneath energy markets. Before we move forward to spread calculations and P&L management, let's recap the core principles.

Power: The Instant Commodity

Electricity must balance supply and demand within milliseconds. There is no meaningful storage at system scale. Frequency (50 Hz in the UK) is the real-time indicator of balance. NESO dispatches plants via the merit order—cheapest first, marginal plant sets the clearing price. Because there's no buffer, prices are volatile: small changes in demand or renewable output can swing prices by £20–50/MWh hour-to-hour.

Traders must forecast not just demand, but wind output, solar generation, interconnector flows, and which plant will be marginal. Get it right, and you profit. Get it wrong, and you're buying back shortfalls in the Balancing Mechanism at punitive prices.

Gas: The Flexible Commodity

Gas operates on a different timescale. Storage (linepack in pipelines, underground facilities like Rough) provides a time buffer measured in hours or days. Commercial shippers self-balance their portfolios, updating nominations throughout the gas day. NBP is a virtual hub—an accounting fiction that creates a liquid trading point. Because storage dampens volatility, gas prices move more slowly than power.

But physical constraints still matter. When pipelines are full or regional demand spikes, local prices diverge from NBP. Basis risk—the gap between hub prices and physical delivery costs—is always present.

The Connection: Gas-Fired Generation

CCGT plants are the bridge between gas and power markets. When NBP gas rises by £5/MWh (in energy terms), the variable cost of a CCGT rises by roughly £10/MWh (due to efficiency losses and carbon costs). If power prices don't rise by at least that much, the Spark Spread—the plant's profit margin—goes negative, and the plant shuts down.

This linkage is why traders obsess over spread trading. The Spark Spread determines whether a gas plant runs. The Dark Spread does the same for coal (when it was still around). The merit order is dynamic, driven by fuel prices, carbon costs, and renewable output. Understanding the physical dispatch logic is the first step to predicting prices.

Looking Ahead

We've focused on the UK domestic market—gas flowing through NBP, power balancing at 50 Hz, plants dispatched by merit order. But the UK doesn't exist in isolation. Gas arrives as LNG from the US and Qatar (with the US now the dominant supplier). Power prices are influenced by French nuclear output and Norwegian hydro. Carbon costs flow from UK and European policy.

In the next guide, *The Global Energy Stack*, we'll zoom out to see how energy markets connect globally:

- How LNG tankers link NBP to Henry Hub, TTF, and Asian spot markets
- Storage as a tradable asset: seasonal spreads, injection/withdrawal economics
- Carbon markets: how UKA and EUA prices influence dispatch and plant profitability
- Cross-border hub arbitrage and the mechanics of global gas flows

New physical assets reshaping dispatch: Battery Energy Storage Systems (BESS) are rapidly changing the intraday power market. Unlike CCGT plants (which take minutes to ramp), batteries respond in sub-second timescales, capturing frequency response and intraday arbitrage.

BESS trading strategies revolve around cycling economics—how many charge/discharge cycles before degradation—and co-optimizing multiple revenue streams (frequency response, capacity markets, energy arbitrage). We'll explore BESS economics and trading strategies in a dedicated guide. Similarly, algorithmic trading now dominates intraday and Balancing Mechanism markets—systems that optimize dispatch decisions every settlement period based on rolling forecasts and market signals, treating assets as portfolios of real options.

Beyond wholesale prices: This guide has focused on commodity prices—£/MWh for power and gas. But for industrial consumers and suppliers, wholesale energy typically represents only 40% of the final bill. The other 60% includes network charges (TNUoS for transmission, DUoS for distribution), balancing charges (BSUoS), and policy levies (Renewables Obligation, Contracts for Difference, Nuclear RAB). For demand-side traders, optimizing *when* energy is consumed to avoid peak network charges (Triads in the UK) can be as valuable as optimizing the commodity price itself. This has spawned a generation of peak-shaving software and demand response aggregators, turning consumption timing into a tradable asset. These non-commodity costs increasingly drive trading strategies for suppliers and large industrials, and we'll explore them in future guides on market structure and demand response.

Then we'll return to calculations: the Spark Spread guide will show you how to predict whether your gas plant should run, and the Position & P&L guide will show you how to aggregate trades, mark positions to market, and calculate your exposure.

But all of that rests on what you've learned here: the pipes, the wires, the merit order, and the physical constraints that define energy trading.

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