Nuclear Power: Myths, Safety and the Price of Every Choice

22 November 2026 25 min read

> *I wrote this after a good discussion with Claude. Consequentially it has some assumptions in it that I don't fully agree with*

Why this guide exists

Nuclear power is safer than most people assume, but it is not risk-free, and every safety choice has a price somewhere else. This guide is for anyone who wants to judge it on the evidence rather than on films, headlines or slogans.

We will do three things. First, look at how radiation actually behaves. Second, take the most common assumptions about nuclear power and test them. Third, compare a new kind of reactor, the Small Modular Reactor (SMR), with wind and solar on the thing that matters most for climate: carbon over the whole life of the plant.

None of this is about telling you what to think. Consider each point for a moment, let it settle, and then decide for yourself. The best engineers are the ones who can argue both sides.

How radiation works

Radiation is energy given off by unstable atoms as they settle into a more stable form. You are exposed to it every day from rocks, the air, food, space and medical scans. The question is never "is there radiation?" but "how much, what kind, and for how long?"

The three main types

Type	What it is	What stops it	Main danger
Alpha (α)	Two protons and two neutrons	A sheet of paper, or your outer layer of dead skin	Harmful only if breathed in or swallowed
Beta (β)	A fast electron	A few millimetres of aluminium or plastic	Skin burns; harmful inside the body
Gamma (γ)	High-energy light	Thick lead, concrete or water	Passes through the body from outside

Neutrons are a fourth type, produced inside a working reactor. They are stopped by water and concrete, which is one reason reactors sit inside thick concrete and water.

Why neutrons limit a reactor’s life

Neutrons do more than keep the chain reaction going. Fast, high-energy neutrons slam into the metal around the core and knock atoms out of their places in the crystal structure. Over decades, this makes steel harder and more brittle, a process called **neutron embrittlement**. Some atoms in the structure also become radioactive themselves.

The reactor pressure vessel, the thick steel container holding the core, cannot practically be replaced. So how much neutron damage it can safely absorb is one of the main factors deciding how long a reactor can run. Operators place small metal test samples inside the reactor and remove them over the years to measure the damage, and this evidence is what regulators use when deciding whether to extend a plant’s life beyond its original 40 or 60 years.

This problem is even harder for **fusion** reactors. Fusion of hydrogen isotopes releases neutrons carrying roughly seven times more energy than typical fission neutrons, and they also create helium bubbles inside metals. Finding materials that can survive years of this bombardment is one of the biggest unsolved engineering challenges for fusion power, and dedicated test facilities are being built just to study it. Fusion deserves a whole guide of its own, so we will leave it there for now.

Measuring dose

Dose is measured in sieverts (Sv), usually millisieverts (mSv). One mSv is a thousandth of a sievert. Some rough everyday comparisons:

Source	Approximate dose
Eating a banana	0.0001 mSv
Chest X-ray	0.1 mSv
Return flight London to New York	0.08 mSv
Average UK background per year	2.7 mSv
CT scan of the abdomen	8–10 mSv
UK annual limit for radiation workers	20 mSv
Clear increase in long-term cancer risk becomes measurable	around 100 mSv

The figures are approximate; natural background varies a lot by place. Cornwall, with its granite and radon gas, runs several times higher than the UK average.

The inverse square law

For a small source radiating in all directions, the dose rate falls with the square of the distance.

$$I \propto \frac{1}{d^2}$$

Double your distance and you receive one quarter of the dose. Triple it and you receive one ninth. Ten times further away gives one hundredth.

Think of spray paint. Hold the can close and you get a thick patch. Step back and the same paint spreads over a much bigger area, so each part gets far less. Radiation spreads out the same way, over the surface of an ever-growing sphere.

This is why the three rules of radiation protection are **time, distance and shielding**: spend less time near a source, keep further away, and put something dense in between. Distance is free and very powerful.

Two caveats matter. The law applies to a point source; a large contaminated area behaves differently. And it describes radiation travelling through space, not radioactive dust that you breathe in or carry on your clothes, which is why contamination control matters as much as distance.

How a reactor works, in plain terms

A nuclear reactor is a very elaborate kettle. It makes heat, the heat boils water, the steam spins a turbine, and the turbine drives a generator.

The heat comes from fission. A uranium atom absorbs a neutron, splits into two smaller atoms, and releases energy plus two or three new neutrons. Those neutrons go on to split more atoms. This is the chain reaction.

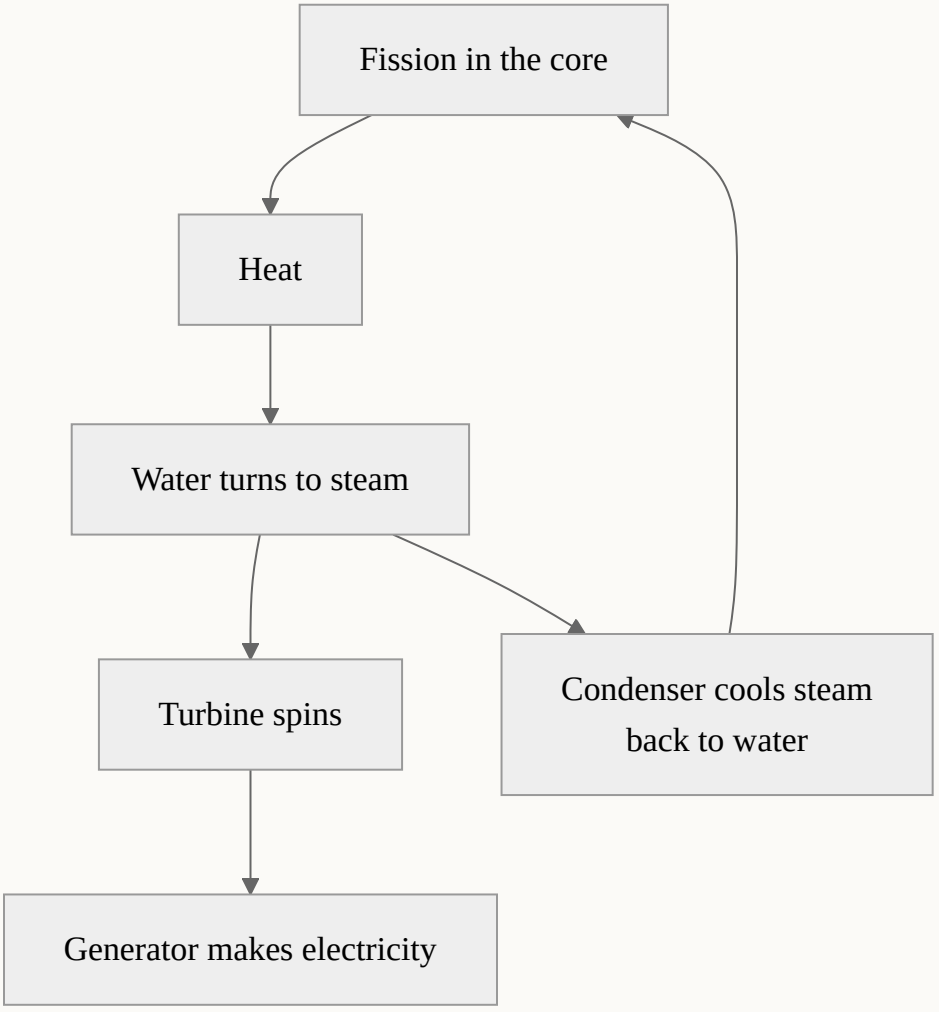

Operators control the reaction with **control rods**, made of materials that soak up neutrons. Push them in and the chain reaction slows or stops. This is called a SCRAM, and in most designs it happens automatically within seconds.

The part most people miss: decay heat

Stopping the chain reaction does not stop the heat. The broken fragments of uranium are themselves radioactive and keep giving off energy as they decay.

At the moment of shutdown, decay heat is about 6–7% of the reactor's full power. After an hour it is around 1–2%, and after a day about 0.5%. That sounds small, but for a large reactor 1% is still tens of megawatts, enough to boil away its cooling water if nothing replaces it.

This single fact explains almost every serious reactor accident. The danger is not an explosion like a bomb; it is losing the ability to carry away heat after shutdown. Consider this for a moment, because every safety system in the next sections is really an answer to one question: *how do we keep the fuel cool when things go wrong?*

Assumptions worth questioning

Some assumptions were built into early reactor designs by engineers; others are held by the public. Both kinds deserve testing.

Assumption 1: "The reactor will always have a grid connection"

Many older reactors were designed on the assumption that, after shutdown, they could draw power from the national grid to run their cooling pumps. If the grid failed, diesel generators on site would take over.

Fukushima Daiichi in 2011 showed what happens when both fail together. The earthquake knocked out the grid connection. The reactors shut down safely and the diesel generators started as designed. Then the tsunami flooded the generators and electrical switchgear, many of which were in basements. With no power for pumps, decay heat melted the fuel in three reactors.

The lesson was not that nuclear power is uncontrollable. It was that designs had quietly assumed two independent systems could not fail for the same reason at the same time. Engineers call this a **common-cause failure**, and hunting for them is now a central part of safety reviews.

Assumption 2: "A meltdown could destroy the world"

This is the most widespread fear, and it is not true. A power reactor cannot explode like a nuclear bomb.

Bomb material is enriched to about 90% uranium-235 and must be crushed together with extreme precision in microseconds. Most power reactor fuel is only 3–5% uranium-235. Physically, it cannot sustain that kind of explosion, just as damp firewood cannot behave like gunpowder.

What a meltdown *can* do is release radioactive material into the local environment. That is serious and can force people from their homes for years. But it is regional, not global.

The two worst accidents give a sense of scale:

- **Chernobyl (1986)**: an unsafe design with no proper containment building, run during a badly managed test. Around 30 plant workers and firefighters died within months. Several thousand thyroid cancers followed in

people exposed as children, most successfully treated. A large exclusion zone remains.
- **Fukushima (2011)**: the earthquake and tsunami killed about 18,000 people. Radiation itself caused very few deaths, but the rushed, long evacuation caused over 2,000 deaths among elderly and vulnerable people through stress and disrupted care.

Think about that and let it settle. At Fukushima, fear of radiation, and the way the evacuation was handled, harmed far more people than the radiation did. That is a lesson about planning and communication, not only engineering.

Assumption 3: “Nuclear waste is an unsolvable problem”

Nuclear waste is real, but the picture is different from green barrels leaking in a cartoon.

- **It is small in volume.** A person’s lifetime share of high-level waste from nuclear electricity would fit roughly in a drinks can. Most of the radioactivity is concentrated in a tiny fraction of the total waste volume.
- **It is fully tracked.** Unlike carbon dioxide from fossil fuels, which goes straight into the air, every gram of spent fuel is accounted for and stored.
- **It gets less dangerous over time.** Radioactivity falls steeply in the first few hundred years, though some elements stay hazardous for many thousands.
- **There is now a permanent solution in practice.** Finland’s Onkalo repository buries spent fuel in copper canisters more than 400 metres down in 1.9-billion-year-old bedrock. Posiva began drilling its first real disposal holes in August 2026 and aims to start final disposal by the end of 2026. The UK is still searching for a site for its own Geological Disposal Facility.

The honest position is that waste is a *political and long-term management* challenge more than a technical mystery. Critics point out that promising to keep something safe for 100,000 years is still a very big promise.

How engineers answer those problems

Modern reactor safety rests on one idea, **defence in depth**: several independent layers, so that no single failure leads to harm. Each mitigation below targets one of the weaknesses from the previous section.

Backup power, done properly

Every reactor has on-site emergency generators to run cooling pumps if the grid fails. After Fukushima, the rules tightened. Generators must be protected from flooding, placed at different heights or locations, and backed by batteries and portable equipment that can be brought in from off site. The aim is to remove common-cause failures, not just to add more of the same kit.

Passive safety: designs that fail safe

The bigger step is to stop relying on power at all. **Passive** systems use gravity, natural circulation and stored pressure rather than pumps and operators.

- **Gravity-fed water**: large tanks sit above the reactor. If power is lost, valves open and water flows down on its own.
- **Natural circulation**: hot water rises and cool water sinks, so coolant can keep moving through the core without pumps.
- **Air cooling of the containment**: the outer steel shell is cooled by air drawn up around it like a chimney.
- **Fail-safe control rods**: in many designs the rods are held up by electromagnets. Lose power and they drop into the core under gravity.
- **Negative feedback**: most modern reactors are built so that as they get hotter, the chain reaction naturally slows. Chernobyl's design did the opposite in some conditions, which is a key reason it ran away.

The Westinghouse AP1000, for example, is designed to cool itself for 72 hours with no power and no operator action. Many SMR designs aim for indefinite passive cooling.

Core catchers

If all else fails and the fuel melts, it can form a very hot molten mass called corium that could eat through the reactor vessel and the concrete below. A **core catcher** is a structure under the reactor designed to receive that molten mass, spread it out so it can cool, and stop it reaching the ground or groundwater.

Newer large reactors such as the EPR design at Hinkley Point C and Russia's VVER-1200 include core catchers. Some other designs choose a different route, flooding the outside of the reactor vessel to keep the melt inside it.

Safer fuels

Conventional fuel is uranium oxide pellets in thin metal tubes (cladding). At very high temperatures the zirconium cladding reacts with steam and releases hydrogen, which caused the explosions at Fukushima.

Safer options include:

- **Accident-tolerant fuels**: coated or alternative cladding that reacts far less with steam and buys operators extra hours.
- **TRISO fuel**: tiny grains of uranium, each wrapped in layers of carbon and ceramic, about the size of a poppy seed. Each grain is its own miniature containment and can survive temperatures above 1,600 °C, far beyond what the reactor normally reaches. It is used in high-temperature gas reactors, including several SMR and advanced designs.

- **Molten salt fuels**: the fuel is dissolved in a liquid salt that operates at normal air pressure, so there is no high-pressure steam to drive material out. If it overheats, a frozen plug melts and the fuel drains into tanks shaped to stop the chain reaction.

Every safety gain has a cost somewhere

There is no free lunch in engineering. Improving one thing almost always makes something else harder, more expensive or less efficient. Learning to spot these trade-offs is more useful than memorising any single fact.

Safety improvement	What it gains	What it costs
TRISO fuel	Survives extreme heat; each grain contains its own radioactivity	Uranium is spread through lots of graphite and ceramic, so spent fuel volume per unit of electricity is much larger
Small reactor cores (SMRs)	Easier to cool passively; less decay heat per unit	More neutrons escape the core, so more material becomes radioactive; one Stanford-led study estimated 2 to 30 times more waste volume per unit of energy
Higher-enrichment fuel (HALEU, up to about 20%)	Longer time between refuelling; smaller cores	Closer to weapons-usable levels, so stricter security; limited global supply
Molten salt fuel	Low pressure; can drain away passively	Corrosive salts; complex, chemically reactive waste; little operating experience
Core catchers	Contain a worst-case melt	Cost, space and extra concrete in a design that may never use it
Passive safety systems	No need for power or operators	Harder to test at full scale in advance; relies on physics behaving exactly as modelled
More backup generators	Fewer ways to lose power	More equipment to maintain, test and secure

A closer look: safer fuel, more waste

Take TRISO as a worked example. The ceramic and graphite layers that make each grain so robust are exactly what makes the waste bulky. You are, in effect, burying the packaging along with the product.

But volume is not the whole story. Repositories are mainly limited by heat and by how easily certain elements could dissolve and move through rock over thousands of years. Some argue TRISO's tough coating is itself a good long-term waste form; others argue more volume means bigger, costlier repositories.

The SMR waste study was also challenged by reactor developers, who said it used out-of-date versions of designs still in development. That is normal science: a claim is made, tested and argued over. Just consider this: when you read a single striking number, ask who produced it, what they assumed, and who disagrees.

Small Modular Reactors: why people are excited

A Small Modular Reactor is a reactor of up to about 300 MW of electricity per unit, built largely from standard parts made in a factory. A typical large reactor is around 1,000–1,600 MW. The idea is to build nuclear plants the way we build aircraft, not the way we build cathedrals.

The case for SMRs

- **Smaller and simpler.** Less fuel and less decay heat per unit makes passive cooling much easier. Many designs can cool themselves indefinitely with no power or operator action.
- **Faster to build.** Factory-made modules shipped to site should cut construction from over a decade to a few years, once production lines are running.
- **Long working life.** Most designs target about 60 years of operation, similar to modern large reactors.
- **Centralised manufacture and maintenance.** Building many identical units lets one expert team learn, improve and repair them, rather than every plant being a one-off. Some very small designs are intended to be swapped out whole and returned to the factory for refuelling.
- **Harder to misuse.** Sealed cores that run for years between refuelling, installation below ground, small fuel inventories and passive safety all reduce what an attacker could achieve. There is simply less to release and fewer ways to force a release.
- **Distributed baseload.** Spreading many smaller reactors across the country, instead of a few giant ones, means one fault takes less power off the grid. They can also sit near factories, data centres or towns that need steady power or heat, giving the network more resilience.

The honest caveats

- **“Faster and cheaper” is still a promise.** The first units of any new design are expensive. NuScale’s first US project was cancelled in 2023 after estimated power prices rose sharply. Ontario’s first BWRX-300 at Darlington is budgeted at about CAD 7.7 billion including shared site works, aiming to run by 2030. Costs are expected to fall with each unit, but that has to be proven.
- **“Small” varies.** The UK’s chosen Rolls-Royce SMR is about 470 MW, larger than the usual definition. In April 2026 a contract was signed to start work on three units at Wylfa on Anglesey.
- **Security risk is lower, not zero.** Some designs use HALEU fuel enriched to nearly 20%, which needs tighter protection. And more sites means more places to guard, more transport of fuel, and more cyber-security to manage.
- **Waste.** As covered in the trade-offs section, small cores may produce more waste volume per unit of energy.
- **Regulation takes time.** Each new design needs years of independent safety review before it can be built.

The realistic view: SMRs are one of the most interesting ideas in energy today, and the physics behind their safety case is sound. Whether they are cheap and quick enough to matter will be decided by the first few dozen actually built, in the next ten years.

Carbon: SMRs, wind and solar over their whole lives

Over their full life, nuclear, wind and solar are all very low-carbon, roughly 50 to 100 times cleaner than coal. Nuclear and onshore wind usually come out lowest; solar varies most, mainly because of where panels are made.

No power station is truly zero-carbon. Mining, making steel and concrete, manufacturing, shipping, building, running, maintaining and finally dismantling it all produce emissions. A **life cycle assessment** (LCA) adds all of these up and divides by every unit of electricity the plant will ever make. The result is grams of CO_2-equivalent per kilowatt-hour (gCO_2e/kWh).

Source	Lifecycle emissions (gCO_2e/kWh)	Typical design life	Where most of the carbon comes from
Large nuclear	5.1–6.4	60–80 years	Mining and enriching uranium; concrete and steel
SMR (studies of designs, not yet operating)	4.6–8.4	about 60 years	Same as above; factory build may lower concrete use
Onshore wind	7.8–16	20–30 years	Steel tower, concrete base, blade materials
Offshore wind	12–23	25–30 years	More steel, specialist ships, sea foundations
Solar PV	8–83	25–35 years	Refining silicon, often using coal-heavy electricity
Gas (for comparison)	403–513	25–40 years	Burning the fuel and methane leaks
Coal (for comparison)	751–1,095	40–50 years	Burning the fuel

Figures are ranges across regions and technologies, mostly from the UN Economic Commission for Europe's 2021 study; the SMR range comes from a 2023 review of published SMR studies.

Why the numbers differ: four factors

1. Production lifespan. The carbon cost of building a plant is spread over every unit it ever produces. A reactor running 60 or more years divides its concrete and steel over far more electricity than a wind turbine replaced after 25. Over one reactor's 60-year life, a wind or solar farm on the same site would typically be built two or three times over.

2. Performance fall-off. Solar panels lose about 0.5% of their output each year, so a panel makes roughly 12% less by year 25. Wind turbines also decline with wear; one UK study found output falling around 1.6% a year. Nuclear plants tend to hold their output and are often upgraded to produce more over time. These are approximate figures and vary by model and maintenance.

3. Manufacturing. This is where solar's wide range comes from. A panel made with coal-heavy electricity can carry several times the carbon of the same panel made with hydro or nuclear power. Nuclear's footprint is

dominated by the fuel chain, mainly mining and enrichment, and modern centrifuge enrichment uses far less energy than older methods.

4. Shipping. Transport is usually a small share, typically a few percent of the total. It matters more for heavy, bulky items like turbine towers and blades, and for panels shipped halfway round the world, but it rarely changes the ranking.

What the table leaves out

These numbers compare one kilowatt-hour to another. They do not include the cost of making wind and solar available when the wind drops or the sun sets. UK onshore wind typically produces around a quarter to a third of its maximum over a year, and UK solar around a tenth. Nuclear runs at 80–90% of its maximum.

Filling those gaps needs batteries, pumped hydro, interconnectors to other countries, backup gas plants, or much more capacity than you strictly need. Each of those has its own carbon footprint. Nuclear, in turn, is harder to ramp up and down quickly to match demand.

So the fairest comparison is not “which source is lowest” but “which *mix* gives a reliable, low-carbon grid at a sensible cost”. Most energy planners, including the UK’s, expect the answer to involve all three.

Subsidies: who really pays for energy?

Almost every kind of electricity generation in the UK receives some form of public support, and every pound of it is paid by the public in the end. The only question is whether you pay through your **energy bill** or through **taxes and government borrowing**.

Why energy gets subsidised at all

Power stations cost billions to build and run for decades, and electricity prices swing wildly. Investors will not risk that money without some guarantee they will get it back. Governments also want things markets do not reward on their own, such as low carbon emissions, energy security, and help for households who struggle to pay.

Support comes in several forms. A few UK examples:

Scheme	Who it supports	How it works	Who pays
Renewables Obligation (closed to new projects)	Older wind, solar and biomass	Extra payment for every unit generated, on top of the market price	Mostly bills; since 2026, 75% moved to general taxation
Feed-in Tariff (closed to new projects)	Small-scale solar and wind, such as home rooftops	Fixed payment per unit generated	Bills
Contracts for Difference	New wind, solar, nuclear and others	Guaranteed “strike price”: topped up when market prices are low, pays money *back* when prices are high	Bills
Regulated Asset Base (Sizewell C)	New nuclear	Consumers pay towards the plant while it is being built, lowering borrowing costs	Bills
Capacity Market	Mainly gas plants and batteries	Payment for being available when demand peaks	Bills
Tax reliefs and decommissioning support	North Sea oil and gas; legacy nuclear clean-up	Tax breaks and public funding	Taxpayers
Warm Home Discount and energy efficiency schemes	Low-income and vulnerable households	Bill discounts and home insulation	Bills

What is actually on a household bill

Under the April–June 2026 price cap, a typical household electricity bill was about £875 a year. It broke down like this:

Part of the bill	Annual cost	Share
Wholesale electricity (includes Contracts for Difference and Capacity Market)	£333	38%
Network: wires, pylons and substations	£214	23%
Supplier operating costs, bad debt and industry charges	£141	16%
Policy costs (Renewables Obligation, Feed-in Tariff, Warm Home Discount, efficiency schemes)	£78	9%
Balancing the grid second by second	£42	5%
VAT, supplier profit and other	remainder	about 9%

The biggest single factor is the wholesale price, and that is mostly set by gas. In Britain's electricity market, the most expensive generator needed at any moment sets the price for *all* electricity sold in that moment. That is usually a gas plant, so cheap wind, solar and nuclear power is often sold at a gas-driven price.

Why does the most expensive generator set the price?

This surprises most people, so it is worth asking why the market works this way.

Each half hour, generators offer to sell power at a price. The grid accepts the cheapest offers first and works up the list until demand is met. This ranking is called the **merit order**. Everyone accepted is then paid the price of the last, most expensive generator needed. This is called **marginal pricing** or "pay-as-clear". It happens in each national market, such as Great Britain's, not on one international market; gas matters so much because gas itself is traded internationally.

The reasons given for this design are:

- **Honest bidding.** If everyone gets the clearing price, each generator's best strategy is to bid what it truly costs to produce one more unit. Under "pay-as-bid", where each is paid its own offer, generators instead try to guess the clearing price and bid just under it. Prices often end up much the same, but less transparent.
- **Efficient use of plants.** The cheapest plants always run first.
- **Investment signals.** High prices at busy times reward building new generation and storage.

The criticism is that this worked well when most power came from fuel-burning plants with similar costs. Wind, solar and nuclear cost a lot to build but almost nothing to run one more hour, so paying them a gas-set price can give them large windfall profits when gas is expensive, as happened in 2022. Contracts for Difference partly fix

this by forcing newer generators to pay back anything above their agreed price, but older plants do not have them.

There is a second fairness argument, running the other way. Nuclear and gas plants can promise to deliver power in any given half hour, day or night, calm or windy. Wind and solar cannot: on a still winter night they may produce almost nothing, yet the grid must still keep the lights on. The market mainly pays for electricity actually delivered, not for being ready to deliver it. So the plants that guarantee the grid works during wind droughts may run less often, earn less, and struggle to cover their costs, while the value of their reliability goes largely unpaid. Economists call this the **missing money** problem.

The UK's Capacity Market exists partly to fix this, by paying plants simply for being available at peak times. Supporters of wind and solar reply that they carry their own hidden costs, such as being paid to switch off when the grid is congested, and that the costs of backup should be shared across the whole system. How to value reliability fairly is one of the central unsolved questions in electricity market design.

Alternatives debated in the UK include splitting the market so low-carbon power is sold separately from gas, and **zonal pricing**, where prices differ by region depending on local supply. The UK government reviewed the market design and in 2025 chose to keep a single national price with reforms, rather than move to zonal pricing.

The impact on industry

The UK has had the highest industrial electricity prices of any country reporting to the International Energy Agency in both 2023 and 2024. Firms paid around 50% more than competitors in France and Germany, and about four times as much as those in the United States.

Analysts disagree about how much of this is caused by subsidies. The consultancy Deloitte estimated that decarbonisation levies made up about 10% of industrial electricity bills, and said their effect was dwarfed by high gas prices. It also pointed out two UK-specific problems: fewer exemptions for heavy industry than in Germany, France or the Netherlands, and the fact that almost all costs go on bills here, while other countries pay some of them from general taxes.

The most energy-hungry industries now get rebates through the British Industry Supercharger. But that relief is itself funded by a levy on everyone else's bills, adding roughly £5–£7 a year to a typical household.

Carbon leakage: exporting the problem

High energy costs have contributed to closures in UK steelmaking and other energy-heavy industries. The blast furnaces at Port Talbot closed in 2024, and the government had to step in to keep British Steel's Scunthorpe works running in 2025.

But Britain still needs steel. When it is made abroad instead, often in countries with coal-heavy grids and no carbon price, the UK's own emissions figures improve while global emissions may actually rise. Jobs and skills leave too. This is called **carbon leakage**.

From January 2027, the UK is introducing a **Carbon Border Adjustment Mechanism** (CBAM). Importers of steel, aluminium, cement, fertiliser and hydrogen will pay a carbon charge based on the emissions made in producing them, similar to what a UK producer would pay. Supporters say this levels the playing field; critics worry it raises costs for UK manufacturers who buy these materials.

The arguments on each side

Arguments that subsidies have gone wrong:

- Schemes can be gamed or poorly designed. Early Feed-in Tariff rates were set so high that they had to be cut sharply, and some biomass subsidies have been dogged by questions over whether the wood burned met sustainability rules.
- Some contracts lock in high prices for decades. Hinkley Point C's price was agreed in 2012 for 35 years, and critics say it was expensive from day one.
- Wind farms in the wrong places are paid to switch off when the grid cannot carry their power, while gas plants elsewhere are paid to switch on.
- Paying through bills is regressive: a family on a low income pays the same levy per unit as a wealthy one, so it takes a bigger share of their income.
- Costly energy drives industry abroad, costing jobs and often increasing global emissions.

Arguments that subsidies are necessary and working:

- Without guaranteed returns, few private investors would build long-lived, expensive plants at all, whether nuclear, wind or anything else.
- Competitive auctions have pushed prices down dramatically, especially for offshore wind.
- Contracts for Difference work both ways. During the 2022 energy crisis, some generators paid money back into the system.
- The deepest cause of high prices is dependence on imported gas. Building more low-carbon capacity, even with support, reduces exposure to future gas price shocks.
- Moving costs from bills to taxes does not make them disappear. It changes who pays, and adds to government borrowing.

Consider this for a moment: nearly every electricity source is subsidised in some way, including fossil fuels and nuclear. So the useful question is not "is it subsidised?" but "is this subsidy good value, who pays for it, and what happens if we don't pay?"

Questions to discuss

1. Should energy subsidies be paid through bills or through taxes? Who wins and who loses under each option?

2. If a UK steelworks closes and we import steel made with coal power instead, has Britain helped or harmed the climate?
3. How could a government design a subsidy that is hard to game?
4. Is it fair for today's bill payers to fund power stations that will serve people for the next 60 years?

Pulling it together

- Radiation is everywhere; what matters is the type, the dose and the time. Distance is your best friend, thanks to the inverse square law.
- The real danger in a reactor is not a bomb-like explosion but losing cooling after shutdown, because of decay heat.
- Early designs made hidden assumptions, such as a reliable grid. Modern designs remove them with passive cooling, protected backup power, core catchers and tougher fuels.
- Every safety gain costs something: tougher fuels and smaller cores can mean more waste volume.
- SMRs are promising for safety, speed, resilience and security, but their costs and build times still have to be proven.
- Nuclear, wind and solar are all very low-carbon. The real question is how to combine them into a reliable grid.

Questions to discuss

1. If you had to choose between a fuel that is safer in an accident and one that produces less waste, which would you pick, and why?
2. Why might people fear nuclear power more than coal, even though coal pollution causes far more deaths each year?
3. Who should decide where a waste repository goes: scientists, the government, or the local community?
4. Is it fair to compare electricity sources per kilowatt-hour if some of them need backup when the weather changes?
5. What would convince you that SMRs really are cheaper and faster to build?

Glossary

Term	Meaning
Core	The part of a reactor where the fuel sits and fission happens
Core catcher	A structure under the reactor designed to hold and cool molten fuel
Decay heat	Heat produced by radioactive fragments after the reactor has shut down
Defence in depth	Several independent layers of protection, so no single failure causes harm
Fission	Splitting a heavy atom, releasing energy and neutrons
HALEU	High-assay low-enriched uranium, enriched to between 5% and 20%
Inverse square law	Dose falls with the square of distance from a point source
Life cycle assessment	Adding up all emissions from cradle to grave
Passive safety	Safety systems that work using physics alone, without power or people
SCRAM	An emergency shutdown of the chain reaction
Sievert (Sv)	The unit of radiation dose to the body
SMR	Small Modular Reactor, up to about 300 MW, largely factory-built
TRISO	Fuel made of tiny coated grains that hold in their own radioactivity

Sources

- UNECE lifecycle figures, summarised by Nucleation Capital
- Review of lifecycle studies including SMRs (Energies, 2023)
- Nuclear waste from small modular reactors (PNAS, 2022)
- Developer responses to the PNAS study (Atomic Insights)
- Onkalo passes safety assessment (ANS, August 2026)

- Posiva begins drilling disposal holes (ANS, August 2026)
- Rolls-Royce SMR contract for Wylfa (NucNet, April 2026)
- Darlington BWRX-300 budget and timeline (OilPrice)
- What costs make up an electricity bill? (House of Commons Library, 2026)
- Electricity prices in Great Britain (House of Lords Library, 2026)
- The impact of higher energy costs on UK businesses (ONS)
- How could the UK lower its industrial electricity prices? (Nesta)
- UK to implement a CBAM in 2027 (SteelOrbis)

Radiation doses, accident figures, capacity factors and degradation rates are approximate values.