SUSTAINABLE AVIATION FUEL — FROM NETWORK TO SCENARIO

READING THE REFUELEU AND UK SAF MANDATES AGAINST THE RESOURCE, UNDER TWO NGFS SCENARIOS

Peter Keutgens, FIA

Independent Research Note · Working Paper 04 / 2026 · Carbon Risk & Capital Allocation Series
Draft for public comment — 2026-07 · July 2026 · model data v0.5.1 (2026-07-02)

Companion to Working Paper 03 / 2026 on the SAF process network, and to Working Papers 01 and 02 / 2026 on carbon pricing as a risk factor and the componentised EU ETS cap. Figures and exhibits are generated from the model tables and regenerate when the data changes.

ABSTRACT

Working Paper 03 (WP03) mapped a producer-agnostic process network for sustainable aviation fuel — feedstock to fuel, static, and through a financial and economic lens. This paper puts the same network in motion. Two NGFS Phase 5 scenarios — Net Zero 2050 and Below 2°C — drive the WP03 network to produce dynamic mandate-gap trajectories, binding-input migrations and scenario-conditional cost decompositions for the European Union (EU27+EFTA) and the United Kingdom side by side. NGFS scenarios are read here as normative not predictive: what an orderly, coordinated transition to a 1.5 °C or 2 °C carbon budget requires from the resource pool, not what current policy will deliver. Read that way, the results quantify the shape of the compliance risk that CFOs, banks, leasing companies and insurers actually price.

Six substantive findings emerge. First, the total-SAF mandate — ReFuelEU's 2 % → 70 % ramp and the UK's 2 % → 23.7 % ramp — can be over-fulfilled under both scenarios once the aviation-allocated resource pool is set: the scenarios' feedstock envelopes are large enough for the mandate volumes to be met physically. Whether that supply is actually built depends on whether SAF's marginal cost clears against fossil-jet-plus-carbon-levy, a rational-refinery check the engine does not yet apply (§2, §6). Either way, the total-mandate volume is not where the compliance risk sits. Second, the sub-mandate for synthetic aviation fuel is where the pool binds: PtL (e-SAF) is CO_2-bound under Net Zero 2050 in the EU through to 2050 and mandate-floor-bound between 2050 and 2080. Third, the binding input migrates across the century — biomass-bound in the middle decades on the bio-pathways whose feedstocks the aviation sector must claim from a rising share of a shrinking pool, CO_2-bound at both ends of the century as the RFNBO-eligible biogenic CO_2 pool ramps slowly pre-2050 and contracts again post-2080. A brief H_2-bound phase for Fischer-Tropsch in the earliest 2030s is visible in the model but does not survive as a headline reading of the chart. Fourth, the UK Mandate and the UK ETS operate together as a two-sided bound on the UK PtL market: statutory buy-out on the price side, CRCF absence through to 2029 on the input side. Fifth, EU absorbs supply first; the UK is residual. Sixth, aviation carbon-intensity thresholds land at the eligibility margin, so small methodological choices swing pathway eligibility.

Scope of this draft. As with WP03, the production boundary stops at the neat certified SAF product — the synthetic paraffinic kerosene prior to any blending with fossil jet — so that we are pricing the fuel the mandate counts, not the finished uplift. CAPEX and fixed OPEX per pathway are excluded — the cost-per-tonne figures are a floor, not a full cost. Deployment ramps, trade flows and per-project construction timelines are excluded — these are Phase 2 extensions that will refine but not invert the results here. The scenario set is NGFS V5.1 (Phase 5); Shell's alternative demand trajectory is discussed as a sensitivity bracket but not fully implemented. The geography is EU27+EFTA plus the United Kingdom — "Greater Europe" as the unit of analysis. Global figures are used only where triangulation adds value.

1. FROM PROCESS NETWORK TO SCENARIO ENGINE

Working Paper 03 (WP03) in this series set out a producer-agnostic process network for SAF, mapped from feedstock to fuel and read through a financial and economic lens. That paper stopped at a static picture: variable operating costs per pathway on the model's canonical tables, a break-even comparison against the SAF market price, and a supply-side ceiling summary that named the binding resources without tracking them over time. WP03 also stated, in its scope note, that "carbon-intensity values and price dynamics are developed in companion threads and only referenced here." This paper is that companion. It takes the WP03 network and sets it in motion.

The instrument that moves it is a scenario. Under the Network for Greening the Financial System (NGFS) Phase 5 scenario set, an orderly path — Net Zero 2050 in its more ambitious form; Below 2°C in its softer variant — describes what a coordinated global carbon-pricing transition to net zero looks like if the policy were in place. The orderly paths are anchored to climate targets from the Intergovernmental Panel on Climate Change (IPCC) — Net Zero 2050 targets the 1.5 °C-aligned carbon budget, Below 2°C targets the 2 °C-aligned one — which is why their normativity is science-based rather than politically chosen: they describe what a coordinated policy response to a physically-defined constraint would require. The scenarios are not forecasts of current policy, and the distinction is more than pedantic: taking NGFS as a forecast would over-state the carbon prices the market is actually about to see and would under-state the compliance gap the market is actually about to face. Read correctly, an orderly scenario is a normative engineering brief for what the transition requires — precisely the thing that the CFO, the underwriter, the leasing residual and the project financier need to price policy risk against. It is also the language in which those readers already work. WP03 mapped the engineering-and-chemistry side of SAF — mass balances, yields, energy intensities, feedstock flows; NGFS scenarios are the operating language of climate risk in finance, the frame in which allocators stress portfolios, banks parametrise transition-risk models and insurers set forward-looking underwriting. The two worlds normally speak past each other. Bringing them through the same engine — driving the engineering network with the financial scenarios — is what allows the SAF question to be priced in the same units as every other transition-risk exposure on the balance sheet, and is the substantive contribution this paper aims to make.

WP03's opening framed EU aviation as caught in a vice: the EU ETS makes the right to emit dearer while ReFuelEU Aviation mandates a rising share of SAF that does not yet exist at scale.

This paper does not soften that picture; it sharpens it. The shape of the closing depends on which scenario is read. Under NGFS Net Zero 2050, EU carbon prices rise to roughly $1,580 per tonne of CO_2 by 2050, aviation demand contracts by about 40 % under the scenario's NiGEM macro-feedback on income and modal choice, and the required SAF absolute tonnage in 2050 comes in below what a naive extrapolation of ReFuelEU × today's demand would suggest. That 40 % is scenario-internal: it does not incorporate several channels of demand or fuel-burn reduction that sit outside the model's frame, including the concentration of the mandate-plus-ETS cost signal on intra-EEA flights only (extra-EEA international departures fall under the weaker CORSIA regime), efficiency gains from air-traffic-control procedure optimisation, and per-flight fuel-burn reductions from continued airframe and engine improvements and the eventual electrification of short-haul routes. Each of those channels would compound with the NGFS contraction where realised; each is discussed as a limitation in §6. But the sub-mandate for e-SAF, meanwhile, runs straight into a structural CO_2 supply constraint that neither the network nor the mandate reasoning alone reveals. Under Below 2°C, the shape shifts: carbon prices peak later and lower, demand contracts less sharply, and the compliance gap widens along different edges. The compliance question is therefore not "how much SAF" but "how much SAF, in which scenario, and where in the century".

WP03 also stopped at the European Union. The 2024 addition of the United Kingdom SAF Mandate to the European regulatory picture — a distinct statutory regime with its own trajectory, its own buy-out mechanism and its own eligibility framing, running in parallel to ReFuelEU Aviation — changes the unit of analysis. This paper treats Greater Europe — the EU27 plus EFTA aggregate that ReFuelEU covers, plus the United Kingdom under SI 2024/1187 — as one market for SAF and reads the two regimes side by side. The two mandates front-load and profile differently; they set buy-out prices at different levels and against different eligibility structures; and they interact through cross-border arbitrage on the small population of e-SAF molecules that will exist in the late 2020s. The UK is neither a footnote to the EU story nor a self-contained parallel: it is a residual market whose behaviour is materially shaped by what its larger neighbour is paying.

This paper covers aviation only. The same NGFS-driven framing applies naturally to any sector competing for the same shared resource pools — most immediately maritime (in the EU under FuelEU Maritime), then heavy industry (steel, cement, chemicals, refining, aluminium) whose EU ETS exposure and hydrogen or electricity intensity make them adjacent claimants on biomass, low-carbon power and carbon dioxide. Extending the reading to each is a natural sequel, and would sensibly proceed in order of ETS-price impact and decarbonisation cost stack. The aviation-only reading here is not compromised by that omission: aviation's competing share of the contested resources is already implicit in the NGFS-derived aviation share, which is scenario-internal to a world in which every sector is bidding.

Three modelling choices carry through the remainder of the paper. First, the scenario engine consumes NGFS variables at their published cadence — carbon prices, primary and secondary energy quantities, biomass and hydrogen prices, aviation demand, and the CCS-to-biomass-supply variable that stands in for the RFNBO-eligible CO_2 pool — and derives daily trends by cubic-spline interpolation between the NGFS anchor points, so any date the paper cites has a defined value. Second, aviation's share of NGFS-published economy-wide resource totals is computed from the scenario itself rather than assumed as a fixed percentage: for each shared input, the aviation share is derived from NGFS's own ratio of aviation's liquid-fuel demand to the total transport-sector liquid-fuel pool, and it rises through the century because road and rail electrify faster than aviation, so aviation's demand for liquid fuel becomes a larger share of the shrinking liquid pool — and by extension the model's assumed share of contested biomass, hydrogen and electricity rises with it. Third, the CO_2 input to the power-to-liquid pathway is priced not at engineering marginal cost but at the opportunity cost of not storing that CO_2 permanently under the EU Carbon Removal Certification Framework — a mechanism whose importance the WP03 draft acknowledged but did not price. That opportunity cost rises with the carbon price, and its magnitude changes the story that a carbon-price-alone reading of e-SAF economics would tell.

The remainder of the paper is structured to be useful to a financial reader who is deciding what to underwrite, finance, allocate to or lease against, without demanding that they follow the modelling in detail. Section 2 sets out the scenario engine and how the NGFS trajectories drive the WP03 network. Section 3 presents the model's substantive results: where the mandate can be over-fulfilled from the scenario resource pool and where it cannot, how the binding constraint evolves from an early hydrogen limitation on Fischer-Tropsch to biomass and ultimately carbon dioxide over the century, why the "high carbon price closes the SAF gap" narrative is incomplete, and how the UK behaves as a residual market for e-SAF next to the higher-clearing EU. Section 4 reads those results four times — for the institutional investor, for the project financier, for the leasing company, and for the insurer — with the mechanism most relevant to each audience placed in the section that leads with it and cross-referenced from the others. Section 5 collects the data exhibits from which the argument is built. Section 6 states the limitations, names the sensitivities, and invites correction from the network of practitioners whose work this paper draws on.

Sources: Working Papers 01, 02 and 03 in the Carbon Risk & Capital Allocation Series. NGFS Phase 5 Scenarios V5.1 (Network for Greening the Financial System, 2025). ReFuelEU Aviation Regulation (EU) 2023/2405; UK SAF Mandate Order (SI 2024/1187); DfT SAF Mandate Compliance Guidance 2026.

2. THE ENGINE: NGFS TRAJECTORIES DRIVING THE WP03 NETWORK

The engine that produces the results in the remainder of this paper has three layers stacked on top of each other. The lowest layer is the NGFS scenario data — variables published by the Network for Greening the Financial System at their standard cadence, ingested into a set of dimensioned tables that keep track of which model, region, scenario and release each fact belongs to. The middle layer is the WP03 process network — the same producer-agnostic set of pathways, mass balances and yields that WP03 documented, unchanged in shape but now consumed as the physical constraint layer against which volumes have to reconcile. The top layer is the supply algorithm — a merit-order allocation that fills the aviation-demand ceiling from the cheapest pathway upwards, subject to resource caps and pathway eligibility. The statutory mandate line is drawn on the same chart as a reference, but is not a stopping rule: the algorithm keeps allocating SAF beyond the mandate percentage until either a physical resource is exhausted or the aviation-demand ceiling is reached. What that says is "how much SAF the aviation-allocated resource pool could physically support", not "how much SAF the market would build" — the cost-vs-fossil-plus-carbon-levy check that a economically rational producer would apply is deferred to §6.

Data ingest. The scenarios in this draft are NGFS Phase 5, version 5.1, using the GCAM 6.0 NGFS integrated assessment model as the source of variables at the native GCAM regional resolution and its Downscaling[GCAM 6.0 NGFS] product for the sub-EU regional splits. From that set the engine consumes ten variables, grouped as follows:

Prices —

- `Price|Carbon` — the shadow carbon price that drives the fossil-jet levy and the CO_2 opportunity-cost mechanism developed in §3.3.
- `Price|Secondary Energy|Hydrogen` — clean-hydrogen price for the FT and PtL pathways.
- `Price|Secondary Energy|Liquids|Biomass` — bio-liquids price, used as the proxy for ethanol and bio-oil.

Resource quantities —

- `Primary Energy|Biomass` — total biomass supply (feeds the FT budget).
- `Secondary Energy|Liquids|Biomass` — bio-liquid fuel production (feeds the alcohol-to-jet and HEFA budgets).
- `Secondary Energy|Hydrogen` — clean-hydrogen production.
- `Secondary Energy|Electricity` — total electricity production.

Aviation demand —

- `Final Energy|Transportation|Aviation` — the closure ceiling against which SAF blending is measured.

Carbon sequestration —

- `Carbon Sequestration|CCS|Biomass|Energy|Supply|Liquids` — the RFNBO-eligible biogenic CO_2 pool available for e-SAF.
- `Carbon Sequestration|CCS|Biomass` (total) — the wider BECCS envelope, for context.

NGFS publishes each of these on a five-year cadence to 2060, then decadal to 2100, with one exception: the downscaled GBR series used for the UK is truncated at 2050 in V5.1, so every UK chart in this paper stops at that year (the fuller note is under Two regions below). The engine interpolates between the anchor points with a natural cubic spline so that any date the paper cites has a defined value, while the annually-reported reference points remain visible on charts as dots for verification.

Two scenarios. The engine runs Net Zero 2050 and Below 2°C — the two orderly paths in the NGFS Phase 5 set that assume coordinated global carbon pricing across the century, with the former on the tighter 1.5 °C-aligned carbon budget and the latter on the softer 2 °C-aligned one. NGFS's disorderly variants and its Current Policies benchmark are not run in this paper: the two orderly scenarios span the range of plausible policy-consistent outcomes, and the results they produce bracket the compliance question well enough for the arguments made in §§3-4. A future extension to the disorderly variants would sharpen the transition-risk reading; a Current Policies read would substantially widen the demand-side envelope by removing the NGFS macro-feedback that produces the aviation U-curve, at the cost of stepping outside the orderly framing on which §1 rests.

There is a further reading of the two-scenario pairing worth naming, because Europe does not decarbonise aviation in isolation. NGFS Net Zero 2050 assumes coordinated global carbon pricing on a 1.5 °C-aligned budget; Europe's mandate architecture — ReFuelEU Aviation, EU ETS, UK SAF Mandate — is written for a world in which the rest of the aviation system is on the same trajectory. It is not currently on that trajectory. ICAO's Carbon Offsetting and Reduction Scheme for International Aviation (CORSIA), which governs international aviation outside the European regulatory bubble, sets a carbon-neutral-growth-from-2020 target that is aligned more with a 2 °C ambition than a 1.5 °C one, and its offsetting mechanism prices non-compliance well below the levels that Net Zero 2050 implies for European aviation. The two scenarios can therefore be read as a Europe-ambition-versus-world-ambition bracket in addition to their orderly-versus-softer-orderly reading: Net Zero 2050 is what Europe is committing to and needs the world to match; Below 2°C is closer to what the world's aviation system is currently being asked to deliver. Where the two levels of ambition remain misaligned, Europe competes against the rest of the world for the same physical resource pool at a carbon-price signal it alone is imposing — a tailwind on the availability of the shared inputs, but not on the demand-side pull that would use them.

Two regions. The geographic unit of analysis is Greater Europe, resolved into two mandate-distinct regions: EU-27 plus EFTA — the ReFuelEU Aviation scope, aggregated inside the engine from the native GCAM 6.0 NGFS regions for EU15, EU12 and European Free Trade Association — and the United Kingdom under SI 2024/1187, drawn from GCAM's Downscaling product with canonical label GBR. Two implementation notes matter for reading the exhibits in §5. First, NGFS V5.1 truncates its downscaled GBR series at 2050, so every UK chart in this paper stops at that year rather than continuing to 2100 as the EU27+EFTA charts do. Second, GCAM's downscaled product does not publish an aviation-specific demand variable for GBR, so UK aviation demand in the engine is derived rather than read: a 2019 anchor from ICAO's Air Transport Reporting is scaled forward by the native GCAM EU15 aviation index, on the reasoning that UK aviation demand behaves like the wider EU15 aviation index but is anchored to its actual pre-COVID scale.

Aviation's share of the shared inputs is derived from the scenario rather than assumed as a fixed fraction, as noted in §1. Concretely, for each shared input — biomass, hydrogen, ethanol, electricity — the aviation share is taken as NGFS's own ratio of aviation liquid-fuel demand to the total transport-sector liquid-fuel pool, on the reasoning that aviation's competing draw on those inputs is well-proxied by its competing draw on the liquid-fuel pool itself. A hard-coded share was rejected because it bakes in the very policy allocation the paper is trying to test. The NGFS-liquids ratio rises from roughly 20 % in 2025 to roughly 40 % by 2050 under Net Zero 2050 EU27+EFTA, and its rise is one of the mechanisms by which the model produces the finding that hydrogen binds first, biomass in the middle years, and carbon dioxide toward the far end of the century. The sensitivity of the results to this choice is examined in §6 (Sensitivity dimensions).

The supply algorithm is a merit-order allocation: for each year and scenario the engine fills the aviation-demand ceiling from the cheapest pathway upwards, subject to each pathway's binding resource cap and to its RED III / SI 2024/1187 eligibility. When one pathway exhausts its input allocation, the next-cheapest steps in until either the total mandate-eligible SAF need is met or every remaining pathway has hit its own cap. Jet fuel is treated as the premium liquid output of each pathway (with co-products credited at their prevailing market prices), so a producer's rational preference is to route joint output toward jet rather than diesel or naphtha. The co-processing pathway sits early in the merit order but is supply-constrained by host-refinery co-processing capacity rather than by feedstock: this bounds it materially below its physical maximum and prevents it from displacing the stand-alone pathways it would otherwise crowd out at low cost.

The WP03 network is consumed as a physical constraint layer, unchanged in shape. Five stand-alone SAF pathways — HEFA, HEFA-oilseed, alcohol-to-jet, Fischer-Tropsch, power-to-liquid — plus co-processing as a sixth (the biogenic fraction of the co-processed jet stream), each with the mass balances, yields and per-tonne input intensities documented in WP03's Appendix A. The supply algorithm iterates over the pathways in merit order — cheapest per tonne of SAF first — and for each pathway allocates its inputs from the pool of that pathway's binding resource until the input runs out or the pathway's own eligibility cap under RED III and SI 2024/1187 binds instead. What each pathway can produce is bounded by three separate things: the physical intensity from WP03, the scenario-derived aviation share of the shared input, and the regulatory eligibility that either counts the fuel toward the mandate or does not. When a pathway's binding constraint switches from one of those to another as the model advances through the century, that switch is what we will call the binding-input migration, and §3.2 traces it.

The cost stack. All monetary values in this paper are quoted in US dollars at a 2010 real base, the currency and base year in which NGFS publishes its scenario prices. The reader converting to nominal 2026 dollars or to euros should apply the appropriate deflator and FX; the paper's substantive conclusions are unchanged by the choice of unit but the absolute magnitudes shift materially. For each pathway the engine computes a $/t SAF number that is the sum of feedstock and energy intensities from the WP03 network multiplied by the NGFS-published prices for the relevant year and scenario. This is the variable-OPEX cost that WP03 introduced, extended in one substantive way. For the power-to-liquid pathway the CO_2 input is priced not at engineering marginal cost — an indicative figure of around $200 per tonne of delivered CO_2 at scale, sitting in the middle of WP03 Exhibit 5 (p. 34)'s range between biogenic point-source at the low end and direct water capture at $100-250 — but at the higher of that engineering floor and the opportunity cost of not storing the CO_2 for a Carbon Removal Certification Framework credit. That opportunity cost rises with `Price|Carbon`: a supplier who could earn `Price|Carbon` per tonne by storing captured CO_2 permanently will not release it to a PtL refinery for less, so PtL must bid up to the same level to keep the input flowing. As the carbon price rises through the century under Net Zero 2050, the effective CO_2 input cost to PtL is therefore pushed materially above its engineering floor and grows in step with the carbon-price trajectory. §3.3 develops this opportunity-cost mechanism and its consequences. CAPEX and fixed operating cost remain excluded; the cost decomposition in §5 is a floor, not a full cost.

A note on the AtJ input pricing. Ethanol has no NGFS-published price of its own; the engine takes NGFS `Price|Secondary Energy|Liquids|Biomass` (a bio-liquid-fuel pool price series that blends crop-derived and cellulosic bio-liquids) and applies ethanol's lower heating value of 26.7 GJ/t to convert US$2010/GJ to US$2010/t ethanol. That proxy sits at the low end of the specific-technology cost spectrum: real-world Annex IX-A cellulosic ethanol (LanzaTech and equivalents) is materially more expensive per litre than the modelled price implies. Feedstock volume is drawn from NGFS `Agricultural Demand|Crops|Energy` via a corn-to-ethanol mass yield of 0.40, so 1 Mt of energy crops maps to 400 kt of ethanol. Two policy caps then constrain what aviation can actually claim: the RED III Article 25 cap on food-and-feed-crop-derived fuels (7 % of aviation transport energy) — which the engine enforces on the crop-ethanol volume available to AtJ — and the RED III Annex IX-B waste-oils cap applied elsewhere in the merit order. The uncapped Annex IX-A cellulosic route is legally available but its scenario-consistent 2030 supply is materially below what NGFS's crop-ethanol pool would suggest. AtJ's paper-cheap number is therefore best read as the ethanol-price bottom under a blended bio-liquid proxy — not as a quotable production cost for any specific commercial project — and the deployment picture in §3.5 is what turns that number into an investable question.

Alongside the per-pathway variable cost, the engine also tracks the carbon levy that falls on any fossil jet residual. Fossil jet not displaced by SAF pays 3.16 tCO_2 per tonne of jet at `Price|Carbon` for the year, which under Net Zero 2050 EU27+EFTA rises from a modest starter today to a peak of roughly $5,000 per tonne of fossil jet in 2050. The aggregate aviation fuel spend the reader sees in §4.3 is the SAF pathway cost stack plus the residual fossil cost plus that levy.

Reproducibility. Every result in the paper carries an engine version tag and a data version tag, stamped on the underlying database rows the results come from. The engine version identifies the merit-order rules, the pathway intensities and the mechanism configuration in force; the data version identifies the NGFS release, the trend-derivation configuration and the loading date. The results in this draft are stamped `v0.5.1 (2026-07-02)` for the data version. As with WP03, any figure quoted from the calculator therefore traces to a specific model state, and later data versions will be visible in the version stamp on each result without changing the calculator's interface. The reproducibility of any figure cited in the paper does not depend on the paper itself being re-issued.

Sources: NGFS Phase 5 Scenarios V5.1 (Network for Greening the Financial System, 2025); GCAM 6.0 NGFS and Downscaling[GCAM 6.0 NGFS] model outputs at native regional resolution. Working Paper 03 (2026) Appendix A pathway network; SI 2024/1187 Article 6(4) and DfT SAF Mandate Compliance Guidance 2026 §5.31 (CI factor formula, CIb = 26.7 gCO_2e/MJ); Regulation (EU) 2023/2405 Annex I(b)(i) (ReFuelEU mandate percentages); ICAO Air Transport Reporting (2019 anchor for UK aviation demand derivation).

3. WHAT THE MODEL SAYS

Six substantive results emerge from running the WP03 network under the two NGFS scenarios across Greater Europe. The first four are about supply — where the mandate can be met from the scenario resource pool, where the binding constraint migrates over the century, why e-SAF viability turns on a mechanism that a carbon-price reading alone misses, and how the United Kingdom behaves as a residual market next to a larger neighbour. The last two are about deployability — the gap between the model's paper economics and what is actually being built, and the engineering choice that the current generation of PtL projects is optimised around. Each subsection anchors on the findings that support it, so a reader tracking the underlying data can trace to the empirical basis; each closes with the pointer to §4's audience-lens reading of the same result.

3.1 The total SAF mandate can be over-fulfilled — but the sub-mandate is where the risk sits

Under both NGFS scenarios and both regions, the aggregate SAF mandate — 2 % in 2025 rising to 70 % by 2050 under ReFuelEU Aviation, 2.0 % rising to 23.7 % under the UK SAF Mandate —

could be comfortably over-fulfilled from 2030 onwards (Exhibit 2 (p. 27)), if the mandate-eligible aviation-allocated resource pool were actually converted into SAF at scale. Fischer-Tropsch alone would carry most of the load in the middle of the century, because lignocellulosic residue — the fibrous non-food plant matter left over from forestry, agriculture and paper production — is the least contested of the shared inputs and its NGFS-published price does not rise as steeply as the price of ethanol, hydrogen or captured CO_2. Under NGFS Net Zero 2050 EU27+EFTA the mandate-eligible pathway output at 2050 would come to roughly 26 Mt of SAF against a 21 Mt required volume (70 % of 30 Mt aviation demand) — an 88.7 % blend of mandate-eligible SAF against total demand, well above the 70 % floor and about 19 percentage points of headroom (Exhibit 2 (p. 27), top panel). Under Below 2°C the same picture is softer in shape but structurally similar: the mandate is still met with room to spare from 2030, again carried mainly by Fischer-Tropsch, with a modestly larger absolute SAF tonnage because aviation demand contracts less sharply under the softer scenario (Exhibit 2 (p. 27), middle panel). The engine's algorithm fills SAF supply from the cheapest pathway upward until the aviation-demand ceiling is reached, subject only to the physical resource caps — it does not stop when the marginal SAF pathway becomes more expensive than fossil jet plus its carbon levy. In practice a economically rational producer would stop at that break-even; the "potential 100 % SAF by 2050" reading of Exhibit 2 (p. 27) should therefore be read as what the scenario resource pool can physically support, not as what an economically rational supply-side response would deliver. Adding that cost-vs-fossil check is Phase 2 work (§6, Merit-order assumptions).

The e-SAF sub-mandate is a different picture. ReFuelEU sets a synthetic-fuels sub-mandate that rises from 1.2 % in 2030 to 35 % by 2050. ReFuelEU stops legislating percentages at 2050; the engine assumes a linear extension of the mandate from 70 % to 100 % over 2050-2080 and of the sub-mandate from 35 % to 50 % across the same window (see §6, Post-2050 mandate ramp assumption), so the results downstream of 2050 are conditional on that extension rule and the sensitivity to a hold-at-2050 alternative is discussed there. Under NGFS Net Zero 2050, PtL is CO_2-bound below the sub-mandate floor pre-2050: NGFS's `Carbon Sequestration|CCS|Biomass|Energy|Supply|Liquids` ramps only slowly through the 2030s, and the RFNBO-eligible biogenic CO_2 pool available in the model is insufficient to support the required e-SAF tonnage. From 2050 to 2080 the sub-mandate is met — barely, mandate-floor-bound rather than resource-bound. After 2080 the story turns again. NGFS treats BECCS as a transitional negative-emissions tool whose role shrinks once the rest of the energy system reaches net-zero: `Carbon Sequestration|CCS|Biomass|Energy|Supply|Liquids` peaks at 126 Mt CO_2/yr at 2050 and contracts to about 80 Mt by 2100, while the extended sub-mandate floor rises through the same period. The two curves cross around 2080 and the sub-mandate gap re-opens.

This is the substantive result that reframes WP03's e-SAF viability gap. WP03 stated the gap in cost terms: EASA 2025 reference prices for synthetic e-SAF land at around €7,520 per tonne — roughly twelve times conventional jet, and above the UK buy-out — with the market-clearing price for a mandated fuel likely being the buy-out itself rather than production cost. Currency note: the paper mixes US$2010 (NGFS scenario prices), €2025 (EASA references) and £2024 (UK statutory buy-outs) figures in their source units — the comparisons are used directionally; converting to a common currency and base year would require an FX + deflator adjustment that we have not attempted. Under the scenario-driven reading here, the gap is not purely a cost gap. It is a resource gap first, structurally binding pre-2050 because the eligibility-safe biogenic CO_2 pool has not yet been built at the scale the sub-mandate requires, and a mandate-floor gap in the middle decades where scenario-consistent supply is available but only just. The audience-lens read of this result — in particular what it means for a project financier's view of when e-SAF capacity is actually needed, and for an insurer's view of eligibility risk on the delivered fuel — is developed in §4.

3.2 The dominant binding constraint migrates from biomass to carbon dioxide, with an early hydrogen constraint on Fischer-Tropsch

WP03 introduced the SAF pathway routes as a ladder from use what exists — co-processing, HEFA — to build an entire new system — Fischer-Tropsch, alcohol-to-jet, power-to-liquid. WP04's engine quantifies which rung of that ladder is binding, on which input, at which point in the century.

In the early 2030s, hydrogen binds first for Fischer-Tropsch (Exhibit 3 (p. 30)). FT's hydrogen intensity of 0.07 t H_2 per t SAF is small per tonne, but at scale — and at NGFS's slow clean-hydrogen ramp through the 2020s — the aviation share of the available hydrogen pool exhausts FT's allocation before its biomass allocation is touched. Under NGFS Net Zero 2050 EU27+EFTA, FT output at 2030 is 1.3 Mt SAF, hydrogen-bound. By 2035 the hydrogen constraint has released — clean-hydrogen production has ramped enough that biomass becomes the binding resource instead — and FT output rises to 7.9 Mt. This is a non-obvious finding because SAF supply narratives typically focus on feedstock scarcity as the near-term bottleneck; in the early 2030s it is not feedstock but hydrogen that limits the most abundant advanced-biofuel pathway. Under Below 2°C the same transition happens, but shifted later by roughly two to three years because the clean-hydrogen ramp is slower.

From the mid-2030s through to 2060, lignocellulosic biomass is the binding constraint for FT and — via the empirical aviation-claimable lipid curve derived from SkyNRG — for HEFA and co-processing. RED III shared caps on Annex IX Part B feedstocks and food-and-feed-crop pathways add a second layer of binding: HEFA-oilseed becomes cap-bound in the 2050s once the rising aviation share of the food/feed pool passes the 7 % of aviation transport energy that Article 25 of RED III allows, and alcohol-to-jet finds the residual RED III headroom compressed accordingly. These interactions are why the model's output at any one date depends not only on physical resource availability but on merit-order priority — a sensitivity we return to in §6.

From 2050 to 2080, PtL is mandate-floor-bound: the CO_2 pool has grown enough that it would in principle support more e-SAF than the sub-mandate requires, but the algorithm produces exactly the sub-mandate floor rather than the pool ceiling because that is the volume the statute demands. The operative constraint has shifted from the resource (CO_2 availability) to the policy floor (the sub-mandate percentage). Then, from 2080 to 2100, the constraint switches once more. NGFS's BECCS|Liquids trajectory contracts materially — from 126 Mt CO_2/yr at 2050 to 80 Mt by 2100 — while the extended sub-mandate rises to 50 %. PtL is again CO_2-bound, and the sub-mandate re-opens.

The result is a three-phase pattern. Hydrogen binds first, in the early 2030s, on the most abundant advanced pathway. Biomass binds in the middle decades, across the bio-pathways whose feedstocks the aviation sector must claim from a rising share of a shrinking pool. Carbon dioxide binds at both ends of the century — first because it does not yet exist at scale in an RFNBO-eligible form, and later because the scenario-consistent BECCS envelope contracts as the wider energy system reaches net-zero. Which end of the century is a binding-input problem depends on which scenario is read. Under Below 2°C, the shape is softer; the sub-mandate rises less steeply and the BECCS envelope contracts less sharply, so the post-2080 re-emergence is milder. Under Net Zero 2050, it is materially present and worth naming.

A related pathway-substitution observation is worth noting. Fischer-Tropsch is the scenario sweet spot through most of the middle of the century — cheapest of the advanced routes because lignocellulosic residue is the least contested shared input, and abundant enough to carry most of the mandate-eligible stack from the mid-2030s through the mid-2040s. Around 2045-2050, however, the position begins to flip: NGFS's biomass price rises steeply toward the carbon-price peak — biomass and BECCS-for-power compete for the same feedstock as the wider energy system reaches for negative emissions — and the FT cost per tonne of SAF converges on PtL's within about $2,000/t at the 2050 peak (Exhibit 4 (p. 32)). FT loses its cost dominance and pathway substitution toward PtL (subject to the CO_2 constraint just discussed) sharpens through the second half of the century. Under Below 2°C the crossover is softer and later.

3.3 e-SAF viability turns on the CO_2 input, not the carbon price

One particularly consequential result in this paper is a mechanism that a straight-through carbon-price reading of e-SAF economics misses. Under a naive reading, a rising `Price|Carbon` is unambiguously good for e-SAF: the fossil-jet levy pushes fossil kerosene's effective cost up, closing the SAF-versus-fossil gap and making the mandated substitution rational. Under a scenario-consistent reading it is more complicated. Rising `Price|Carbon` lifts the fossil-jet cost, yes; but it also lifts the opportunity cost of the CO_2 input that PtL depends on, because CO_2 suppliers can otherwise store their captured molecules and earn a Carbon Removal Certification Framework credit at `Price|Carbon`. Where the removal credit and the fuel-feedstock revenue are both available in principle, the CO_2 supplier will not release CO_2 to a PtL refinery for less than the removal credit could earn. PtL must bid up to the same level to keep the input flowing.

The mechanism relies on one explicit optimisation assumption: a CO_2 supplier able to route captured molecules to either permanent storage or fuel synthesis behaves as a rational profit-maximiser and allocates each tonne to the higher-value use. That assumption is why the fuel-feedstock revenue must clear the removal-credit revenue rather than sit indifferent to it. The engine implements no politically-directed reservation of CO_2 to fuel synthesis, no subsidy tilting the choice toward one destination, and no long-term contract locking supply into fuel at pre-agreed prices. Any of those interventions — a policy floor on the fuel-synthesis price, a mandate to route a share of captured CO_2 to fuels, a strategic-industry contract at a fixed take-off price — would break the mechanism at the point of intervention. Absent such interventions, market-clearing at scale drives the two revenue streams into equality, and the CO_2 opportunity cost is what a PtL producer will actually see.

The engine implements this as a floor: the effective CO_2 input cost to PtL is `max(engineering marginal cost, Price|Carbon)` for each year and scenario. The engineering marginal cost is around $200 per tonne (§2). Under NGFS Net Zero 2050 EU27+EFTA, `Price|Carbon` rises to roughly $850 per tonne of CO_2 at 2030, $1,580 per tonne at 2050, and eases back to $1,200 per tonne by 2100. Multiplied through the PtL pathway's 6 tonnes of CO_2 per tonne of SAF, the effect on the CO_2 input line of the pathway cost is:

- 2030: pre-mechanism $1,200/t SAF (6 × $200); post-mechanism ~$5,100/t (6 × $850). An increase of about 325 %.

- 2050: pre-mechanism $1,200/t SAF; post-mechanism $9,480/t (6 × $1,580). An increase of about 690 %.
- 2100: pre-mechanism $1,200/t SAF; post-mechanism $7,200/t (6 × $1,200). An increase of about 500 %.

The consequence for the overall PtL cost stack is roughly a doubling at the carbon-price peak (Exhibit 4 (p. 32)). Empirically, running the engine at EU27+EFTA × NZ2050 × 2050 produces a PtL cost per tonne of SAF of $7,700 without the mechanism and $15,826 with it — a rise of 106 %. At 2100 the same slice moves from $6,619 to $12,863, a rise of 94 %. Under Below 2°C the effect is present but smaller: `Price|Carbon` peaks lower and later, so the CO_2-input opportunity cost rises less steeply and the PtL cost stack roughly 1.5× at its own peak, rather than doubling.

Two consequences follow. First, e-SAF's absolute cost trajectory under NGFS Net Zero 2050 is substantially higher than a pure input-cost-plus-carbon-price reading would suggest. The "high carbon price closes the SAF gap" narrative is incomplete: it depends on whether captured CO_2 is allocated to storage or to fuel synthesis, and the market-clearing allocation at high carbon prices skews toward storage. Where that is the case, PtL cost climbs in step with the carbon price rather than in opposition to it, and the cross-over against carbon-taxed fossil jet is not automatic. This is one of the paper's more consequential results for a CFO or project-finance reader, because it changes the sign of a variable that most sector models treat as unambiguously favourable.

Second, the mechanism shifts levels but not the rank ordering of pathways. PtL remains the most expensive pathway in absolute terms under both scenarios and both regions; Fischer-Tropsch remains the mid-cost route; HEFA and co-processing remain the cheapest. An investor building a PtL refinery should care about the mechanism intensely because it changes the numerator of a project return calculation; a merit-order planner allocating supply to pathways should note the mechanism but not change their pathway ranking on the basis of it. The mechanism matters for absolute SAF cost narratives and for cost-of-mandate calculations; it does not change which pathway wins in a merit-order allocation.

The CO_2 constraint on PtL also carries into the Below 2°C scenario, not only Net Zero 2050. Under Below 2°C EU27+EFTA, PtL still sits below the ReFuelEU synthetic sub-mandate in the pre-2050 window because NGFS's biogenic CO_2 pool ramps only slowly under Below 2°C as well — the constraint's shape is the same as under Net Zero 2050, with slightly later timing and slightly lower magnitude. The finding that "the sub-mandate is where compliance risk sits" is therefore robust across the paper's two scenarios, not a property of the more aggressive scenario alone. That robustness matters for readers who take Below 2°C as the more plausible near-term operating envelope: the CO_2 constraint on e-SAF supply is not a Net-Zero-2050-only artefact.

A sensitivity dimension worth naming: the engine treats the removal credit as trading at full parity with `Price|Carbon` (a multiplier of 1.0). Reviewers may want a 0.5 case — removal credits traded at a discount because of certification friction and market liquidity — and a 1.2 case for a durability premium. The mechanism's direction is unchanged in either sensitivity; its magnitude scales. A future refinement would also address whether policy might mandate captured CO_2 to be directed at fuel synthesis rather than storage; where such a policy is in place, opportunity cost does not bite, and PtL's CO_2 input line returns to the engineering floor.

3.4 The United Kingdom — parallel regime, residual market

The United Kingdom is neither a shrunken EU nor a self-contained parallel. Its SAF Mandate under SI 2024/1187 sets a distinctively shaped trajectory — front-loaded and plateaued rather than steep and rising — and the interaction of that shape with a neighbouring market that has no statutory buy-out ceiling produces investment behaviour that neither market considered in isolation would predict.

The UK Mandate trajectory itself is worth naming plainly (Exhibit 6 (p. 36)). In 2025 the UK Main Obligation of 2.041 % effectively matched ReFuelEU's 2 %. By 2030 it is 10.6 %, materially larger than ReFuelEU's 6 %. Through the 2030s the UK Main Obligation rises to 23.7 % by 2040 and plateaus there — the same rate holding for every subsequent obligation period. ReFuelEU by contrast keeps rising, reaching 34 % by 2040, 42 % by 2045 and 70 % by 2050. The UK PtL sub-obligation follows the same profile: 0.6 % by 2030 (ReFuelEU: 1.2 %), rising to 4.5 % by 2040 (ReFuelEU: 10 %) and plateaued there. The UK is a near-term HEFA market that does not become a deep e-SAF market.

The asymmetry is worth naming for what it is: the UK is materially more ambitious than the EU in the near term and materially less ambitious in the long term. At 2030 the UK Main obligation is roughly 76 % higher than ReFuelEU's in percentage terms, and the PtL sub-obligation kicks in earlier. By 2050 ReFuelEU has risen to roughly 3× the UK's held-rate. Near-term compliance pressure — supplier procurement effort, buy-out exposure, the tempo of new-plant FID decisions — therefore falls harder on UK-obligated suppliers than on their EU counterparts through the mid-2030s, before the balance flips decisively toward the EU from ~2035 onwards. The Revenue Certainty window for UK non-HEFA SAF is now (the 2027-2028 FID window in §3.5); the equivalent EU-side pressure lies a decade later. The investment-timing consequences of that difference are developed in §4.2.

That shape has two immediate consequences. First, in the near term, UK PtL is two-sided bounded. On the demand side, SI 2024/1187 Article 21 sets a PtL buy-out at £0.145/MJ ≈ £6,235 per tonne of SAF ≈ roughly $7,900 per tonne — an effective ceiling on what any UK fuel supplier will rationally pay for compliance-eligible PtL. On the input side, the UK ETS Authority is targeting end-2028 legislation and end-2029 operational delivery for greenhouse-gas removals integration into the UK ETS. Until then, UK CO_2 suppliers have no carbon-price-level alternative use for their molecules, so — in a mirror-image of §3.3's mechanism — the opportunity-cost line does not bite in the UK. Counter-intuitively this makes UK PtL CO_2 input paper-cheap through to 2029. But the paper-cheapness is theoretical: with no removal-credit revenue path, no one invests in UK CO_2 capture, so there is no supply. UK PtL delivered through to 2029 is effectively zero on a market basis. The narrow exception is a small number of Power BECCS projects underwritten one-by-one by DESNZ (the Department for Energy Security and Net Zero) under bespoke Contracts for Difference — enough to prove the technology, not enough to move the compliance picture.

Second, in the medium term, the UK Mandate operates as a price ceiling rather than merely a fall-back. ReFuelEU has no statutory buy-out; non-compliance triggers administrative penalties under each member-state implementation, with the combined exposure — financial penalty plus back-billing plus supply-chain stigma — materially higher than the UK's £0.145/MJ for any commercially serious supplier. e-SAF is fungible across the EU-UK regulated market. A producer with an e-SAF cargo can sell into either region and will rationally route to the higher-clearing market first. That is the EU. UK demand absorbs what EU absorption leaves.

The empirical claim is on the record. LanzaTech's SAF Café presentation of April 2026 stated the position directly: the UK buy-out price sits materially below EU non-compliance penalties, and — as their deck put it — "Non-UK eSAF will be supplied to the EU NOT UK." Combined with WP04's model result that the required e-SAF tonnage under the UK PtL sub-obligation is small — of order 0.4 Mt at 2040 against a UK Main Obligation of about 2.3 Mt — the finding is that the UK PtL sub-obligation will predominantly be met by buy-out through the 2030s, not by fuel volumes. The mandate becomes a Treasury revenue line rather than a fuel-supply line.

A third dimension of UK exposure is worth flagging: single-licensor concentration on the Fischer-Tropsch side. The UK FT-MSW first-of-a-kind projects that had reached public disclosure by 2026 — Altalto at Immingham, NEXTGEN SAF's FT block at North Tees, Sustainable Molecules' commercial FOAK, all identified through the SAF Café series — Altalto at Immingham, NEXTGEN SAF's FT block at North Tees, Sustainable Molecules' commercial FOAK — uses the same Fischer-Tropsch reactor and catalyst package: Velocys microFTL™. Velocys is not distressed. The company was taken private in January 2024 by a consortium of climate investors (Carbon Direct Capital, Lightrock, GenZero, Kibo Investments) with over $40 million of growth capital, announced a 30 % cost reduction on microFTL™ deployment in February 2026, and signed a manufacturing scale-up partnership with Morimatsu. But the concentration risk stands. Reactor manufacturing throughput, licensing decisions, and any common-mode technical issue propagate across the UK FT-MSW fleet simultaneously. Diversification across FT licensors (Sasol, Johnson Matthey/BP, ARA) would not eliminate but would distribute the exposure.

Finally, there is a post-2050 data horizon on the UK side that reads as a caveat rather than a finding. NGFS V5.1 truncates its downscaled GBR series at 2050, so every UK chart in this paper stops at that year. The extension methodology under discussion — anchoring UK aviation on the native GCAM EU15 series and treating the plateaued mandate rate as effectively held — is documented in the engine, but the paper does not claim UK trajectories beyond 2050 as reproducible from NGFS itself. This is a genuine implementation limitation, not an analytical choice.

3.5 Paper economics do not equal deployable capacity

WP03's Exhibit 1 (p. 25) gave a break-even cost per tonne of SAF for each pathway under a variable-OPEX floor with co-products credited at market prices. WP04 extends that by feeding NGFS input prices through the same pathway network, so the paper economics per tonne evolve over time with the scenario. Under either scenario at either region, the ordering is stable: HEFA and co-processing at the low end, Fischer-Tropsch mid-range, alcohol-to-jet mid-to-high, PtL well above the rest. The cost trajectories are shown in §5's Exhibit 4 (p. 32) and their audience-lens reading in §4.1.

What the paper cost cannot show is deployability. Every UK non-HEFA first-of-a-kind SAF project we track from 2026 industry disclosures is bunched into a Final Investment Decision window between end-2027 and end-2028, targeting commercial operation between end-2030 and 2034. Their combined nameplate is order-of-magnitude twenty times too small for the UK mandate they exist to serve, and each depends on either a Revenue Certainty strike price or a targeted grant to reach FID at all. That gap between paper economics and deployable capacity is the practical constraint on the SAF supply trajectory, and it is invisible in the merit-order calculation.

The concrete picture — a representative sample rather than a comprehensive list, drawn from developers who have presented at the UK Sustainable Aviation Fuel Café series and for whom we have direct company disclosures. The series was hosted monthly online by Innovate UK Business Connect as part of its Sustainable Aviation Fuel Innovation Programme, supporting the UK Government's target of at least 10 % SAF in the UK jet fuel mix by 2030. It ran from November 2023 with published recordings and slide decks accessible from the series page, and it functions as a genuine primary-source register of UK developer intentions, project timelines and technical approaches. Other UK SAF projects at various stages of development exist and are not included here, either because they are earlier stage than pre-FEED, because they are less publicly disclosed, or because we have not yet reviewed the relevant developer's materials. A comprehensive UK SAF project register — cross-checked against Advanced Fuels Fund awards and DfT disclosures — would sharpen the reading but does not change its shape: the deployment window is short, the aggregate nameplate small, and every entrant needs a fixed-price offtake floor to close FID.

- Altalto, Immingham (FT-MSW via MyRechemical gasification + Velocys microFTL™ + Topsoe upgrading): 22 kt/yr SAF; FID target December 2027; commercial operation date December 2030; over £30m invested pre-FID; £27m plus £3m Advanced Fuels Fund awards committed.
- NEXTGEN SAF, University of Sheffield / E.ON (point-source-carbon-to-liquid (PCtL) demonstrator at Blackburn Meadows + commercial plant at North Tees, using CO_2 captured from a nearby industrial emitter rather than from atmospheric direct air capture): 10 kt/yr demonstrator, operational 2029-2030; 73 kt/yr commercial, operational 2033-2034; £901.6m CAPEX at the commercial step; asking a £15m Series A round for FEED. Lifecycle GHG intensity 26.7 gCO_2e/MJ, which we return to in §3.6.
- Sustainable Molecules (FT-MSW via SuMo advanced gasification + Velocys microFTL™ + COX EPC): 13 kt/yr SAF; FID target Q1-Q3 2028; commercial operation date Q4 2030. Self-quoted as ">1.3 % of the UK 2030 mandate".
- LanzaTech DRAGON II, Saltend Humberside (alcohol-to-jet via LanzaTech gas fermentation + LanzaJet upgrading): 80 kt/yr SAF (67 kt advanced RCF plus 13 kt eSAF); pre-FID activity Q2 2026 to Q4 2027; FID end-2027 contingent on Revenue Certainty; commercial operation date Q2 2031.
- Zero Petroleum, Salt End (PtL Direct-FT® in strategic partnership with Technip Energies): pre-FEED only; Class 4 cost estimate developed; strategic timeline dependent on Revenue Certainty stage gates.

Five plants across three pathways, all needing FID inside 30 months of one another, all needing a fixed-price offtake floor to close, all together delivering an order of magnitude below what the UK 2030 mandate needs. This is the deployment bottleneck WP03 gestured at without naming; it is a first-order limit on how quickly the merit-order model's implied supply can be built.

External triangulation supports the reading (Exhibit 8 (p. 39)). The IATA / Worley global outlook of April 2026 projects SAF capacity growing from 22 Mt/yr in 2030 to 413 Mt/yr in 2050 at a 16 % compound annual growth rate (CAGR), with PtL/e-SAF reaching about 43 % of the 2050 mix and Europe contributing 42 Mt/yr. Whether that 42 Mt/yr is "sufficient" depends on the demand basis: it is roughly double the ReFuelEU-plus-UK-Mandate requirement under NGFS Net Zero 2050 (~21 Mt/yr at 70 % of ~30 Mt EU aviation demand plus 23.7 % of ~2 Mt UK), still comfortably above the NGFS Below 2°C requirement (~29 Mt/yr), and at or below the equivalent number under WP03's Concawe reference demand basis (~40-60 Mt/yr). So the "structurally short" reading applied in WP03 was demand-basis-dependent; under this paper's NGFS basis, aggregate mandate volumes look broadly achievable and the substantive constraint sits on the sub-mandate side (§3.1). Trade and imports — discussed under the audience lenses in §4 — matter for whether the domestic-vs-imported split of that mandate volume plays out politically and for whether the sub-mandate is met by physical fuel or by buy-out.

3.6 The dual-market credit-neutral point

NEXTGEN SAF's commercial 200 t/day PCtL design at North Tees reports a lifecycle GHG intensity of 26.7 gCO_2e/MJ. The number is engineered rather than incidental. It sits at the exact intersection of two policy tests in two jurisdictions.

On the UK side, SI 2024/1187 Article 6(4) scales delivered SAF into SAF Certificates via `(CF - CS) / (CF - CR)` — where CF is the 89 gCO_2e/MJ fossil-jet baseline, CS the fuel's actual lifecycle intensity, and CR a statutory reference of 26.7 gCO_2e/MJ. At CS = CR the CI factor is exactly 1: one litre delivered earns one Certificate. On the EU side, Regulation (EU) 2023/2405 Article 3(12) requires renewable fuels of non-biological origin to meet the RED III Article 29a 70 % lifecycle-savings threshold, which against the same 89 baseline gives an eligibility ceiling of 26.7 gCO_2e/MJ. Above 26.7 the fuel is not counted as a synthetic aviation fuel under ReFuelEU at all.

The number lands at the engineering optimum for both regimes at once. Move above 26.7 and the fuel loses ReFuelEU eligibility outright and takes a UK credit reduction on top; move below and the producer pays for cleanliness that only the UK notional-credit formula rewards, while the EU side gets no marginal benefit. NEXTGEN's design sits exactly at the reference: maximum revenue from both regimes without over- or under-cleaning.

The 26.7 is a calculated design value rather than an operational measurement — plants engineer to a specific renewable-electricity mix, biogenic CO_2 source, catalyst and process yield, and report the calculated intensity that configuration produces under the standard methodology. Real-world operational variance is typically several gCO_2e/MJ around the design point, and this matters for the underwriter (§4.4): excursions above the reference land in a binary ReFuelEU eligibility cliff, while excursions below reward only the UK side. Any EU move to tighten the RFNBO 70 % threshold plays out on the two regimes asymmetrically — a binary cliff on the EU side, a scaled reduction on the UK side.

For a project designed to serve both markets — as most UK developers must, given the routing dynamics of §3.4 — the eligibility-cliff risk is asymmetric across the two regimes. That asymmetry is what §4.4 develops as underwriting-adjacent exposure.

Sources: Regulatory: SI 2024/1187 (UK SAF Mandate Order); Reg (EU) 2023/2405 Article 3, Annex I; RED III (Directive (EU) 2023/2413) Article 29a; DfT SAF Mandate Compliance Guidance 2026 §5.31. Industry (developer briefings from the UK Sustainable Aviation Fuel Café series, hosted by Innovate UK Business Connect at https://iuk-business-connect.org.uk/events/sustainable-aviation-fuel-cafes/): LanzaTech (DRAGON II, April 2026); Altalto / Velocys; NEXTGEN SAF / University of Sheffield / E.ON; Sustainable Molecules / SUEZ / COX; Zero Petroleum / Technip Energies. External triangulation: IATA / Worley SAF at Scale outlook (April 2026).

4. READING THE SCENARIOS THROUGH A FINANCIAL AND UNDERWRITING LENS

The results in §3 are the same set of model outputs read from four different seats around the table. The institutional investor asks a portfolio question — how much aviation-sector exposure to carry and at what discount rate. The project financier asks a bankability question — will the plant close at FID and pay back its capital stack. The leasing company asks a residual-value question — how do fuel-cost trajectories move the residual of an aircraft over a ten-to-fifteen-year lease term. The insurer asks an underwriting question — what are the loss triggers to price, and how correlated are they across the projects on the book. The four readings are followed by a closing note (§4.5) that gives their shared underlying risk factor a single name — policy risk — and shows how the same government-delivery question surfaces differently on each of the four balance sheets. This section walks through each audience in turn, placing the mechanism most relevant to each in the subsection that leads with it and cross-referencing from the others.

4.1 Institutional investors — scenario dispersion as a portfolio-level risk factor

By institutional investors we mean the asset managers, pension funds and insurance-company general accounts whose books hold aviation-adjacent exposures — direct positions in SAF producers, positions in airlines and aviation-service companies, positions in oil majors with SAF pipelines, or a broader transition-risk mandate that touches the sector. The §3 results feed each of those readings, and this subsection is written for all of them.

For a long-hold portfolio, the question posed by §3 is not so much which scenario is right — NGFS scenarios are normative not predictive, and §1 argued that the practical way to use them is to price the range they bracket. It is how far apart the scenarios sit at each date, and what the size of that gap implies for the discount rate applied to aviation-sector exposures.

The scenario gap is not uniform across the paper's headline results. On aggregate SAF supply, the two scenarios are close: the resource pool would support mandate volumes under both, and the two pathway mixes at 2050 look similar in shape (Exhibit 2 (p. 27)). On e-SAF sub-mandate compliance, the two scenarios also broadly converge — both hit the sub-mandate floor from 2050 to roughly 2080. Where the scenarios open up is on the cost trajectory of e-SAF, and this is where the portfolio question sits (Exhibit 4 (p. 32)).

Under NGFS Net Zero 2050, `Price|Carbon` peaks at roughly $1,580 per tonne of CO_2 at 2050 in the EU27+EFTA aggregate. Under NGFS Below 2°C, the same variable peaks materially lower and later — of the order of half at its own peak. Multiplied through the six tonnes of CO_2 per tonne of SAF that PtL requires, and applied via the opportunity-cost mechanism developed in §3.3, that carbon-price gap produces a PtL cost gap between the two scenarios of the order of five thousand dollars per tonne of SAF at the mid-century peak, on top of a base variable cost that is itself high. For a portfolio-level exposure to SAF pathway economics, that gap is the scenario risk. The underlying NGFS trajectory shapes for both scenarios are reproduced in the Appendix as the ten-variable trend slice the engine consumes.

Three consequences follow for portfolio construction. First, aviation-sector exposures that pay off under NGFS Below 2°C are not simply a "less-aggressive" version of the Net Zero 2050 story — they are structurally cheaper by amounts that dominate the aggregate SAF spend calculation. A book that is long fossil-jet resilience does better under Below 2°C; a book that is long PtL producer margins does worse. Second, the sensitivity dimension the paper calls out around the CO_2 opportunity-cost mechanism — whether the removal credit trades at full parity with `Price|Carbon` (multiplier 1.0), at a certification-friction discount (say 0.5), or at a durability premium (say 1.2) — is a portfolio dial for the manager willing to take a view on how the removal-credit market clears. Third, the divergence between scenarios means that an allocator's choice of scenario reference case is itself a material portfolio decision — one that is often made implicitly in transition-risk stress-tests and that this paper's results make explicit. The underwriter's version of the same question, framed around loss triggers rather than portfolio construction, is developed in §4.4.

4.2 Banks and project finance — Revenue Certainty as the FID condition

For a project finance desk, the §3 results reframe SAF underwriting from a cost question to a bankability question. The variable-OPEX cost curves in §3.5 are informative but not sufficient. What matters at the desk is whether the plant reaches Final Investment Decision, and every developer we track in §3.5 has said the same thing about what has to happen for that: Revenue Certainty.

The mechanism is direct. Under either NGFS scenario, SAF production costs at every non-HEFA pathway sit above the UK buy-out price of roughly $7,900 per tonne through the 2030s, and above every plausible fossil-jet-plus-carbon-levy break-even for at least the same window. Selling the fuel at market price alone does not clear the project's cost stack. Selling the fuel at the mandate-compliance-forced price, capped in the UK by the buy-out price and in the EU by the implicit administrative-penalty ceiling, still does not clear it. What is required is a third revenue line — a fixed-price offtake floor, whether structured as a Contract for Difference with a strike price above the buy-out or as an equivalent Advanced Fuels Fund grant absorbing part of the CAPEX. Without it, the mandated fuel is not produced and the mandate is met by buy-out payments to the Treasury (§3.4).

For the banker, this changes what "bankable" means. In a traditional project-finance underwrite, "bankable" asks whether offtake revenue at the going market price covers debt service — i.e. whether the market is willing to pay enough for the fuel to service the loan. For non-HEFA SAF, that test fails from the start. The buyer of last resort is the obligated aviation fuel supplier under SI 2024/1187, and that buyer will never pay more than the statutory buy-out price (which is roughly $7,900 per tonne of PtL SAF) — because paying the buy-out itself discharges the same compliance obligation for the same amount. In other words, the market price for compliance-eligible SAF is bounded above, not below, by law. And for every non-HEFA pathway the engineered production cost sits above that ceiling through the 2030s. There is therefore no market-clearing price at which offtake revenue services the debt. The project is only bankable if a third party — the government, via a Contract for Difference or an Advanced Fuels Fund grant — commits to pay the gap between the buy-out ceiling and the production cost. That commitment turns the underwrite from a market-risk question into a policy-delivery question: will the CfD strike price be set high enough, and will the instrument actually be in force at COD?

Compliance is a stack, not a binary. A supplier meets the Main Obligation through a mix of four instruments:

- fossil jet for the non-mandated portion of their volume;

- HEFA SAF up to the HEFA cap in SI 2024/1187 Article 22, which stands at 100 % of the Main Obligation pre-2027 and drops progressively to 42.16 % by 2040 and thereafter;
- non-HEFA compliance-eligible SAF (Fischer-Tropsch, alcohol-to-jet and — for the PtL sub-obligation — power-to-liquid) for whatever residual the HEFA cap leaves, at any price up to the buy-out;
- buy-out payments to the Treasury for anything not procured through the three routes above.

The ETS cost on the fossil residual is a separate compliance regime that does not fall on the fuel supplier — in either jurisdiction. In the UK, the SAF mandate under SI 2024/1187 lands on the fuel supplier, while the UK ETS Aviation allowance surrender obligation under the Greenhouse Gas Emissions Trading Scheme Order 2020 lands on the aircraft operator for flights within its scope. In the EU, ReFuelEU Aviation under Regulation (EU) 2023/2405 lands on the fuel supplier, while EU ETS Aviation under Directive 2003/87/EC lands on the aircraft operator for intra-EEA flights. Both jurisdictions have the same two-party structure: mandate on the supplier, allowance-surrender on the operator. Both regimes attach to the same physical fuel but bear on different balance sheets. For SAF offtake structures the practical consequence is that the supplier's willingness to pay for SAF is anchored on the mandate compliance value (the £0.145/MJ buy-out in the UK; administrative-penalty avoidance in the EU), while an airline's willingness to pay for SAF reflects the avoided ETS allowance cost on the emissions the SAF displaces. A commercially structured SAF offtake needs to see both willingness-to-pay lines to know what price the fuel actually clears at.

The buy-out ceiling bites specifically on the non-HEFA residual. HEFA is already commercial at prices below the buy-out and clears through a normal market. Non-HEFA SAF is where the market-price ceiling problem sits: for Fischer-Tropsch, alcohol-to-jet and power-to-liquid, a rational supplier will not pay more than the buy-out price of roughly $7,900 per tonne, because paying the buy-out itself discharges the same obligation for the same amount. That upper bound sits below the fuel's engineered production cost for every non-HEFA pathway through the 2030s.

There is therefore no market-clearing price at which offtake revenue from a non-HEFA SAF plant covers the plant's cost stack. The applicable test becomes not whether the market price covers debt service but whether the strike price under the Revenue Certainty contract, once agreed, sits above the fuel's engineered production cost including CAPEX amortisation and O&M. That is a policy question at least as much as it is a project question. The project passes if the government's willingness to pay under the CfD mechanism exceeds the project's cost stack, and fails if it does not. Whether the CfD mechanism itself is legislated in time and funded adequately is the shared policy-risk question developed in §4.5.

A related question a banker will ask is whether a UK PtL producer, structurally short in a UK residual market and paid off a UK buy-out ceiling, could sell into the EU instead — where there is no statutory buy-out cap and the effective willingness-to-pay for compliance-eligible fuel is set by the higher EU administrative-penalty ceiling. On the physics and on the mandates themselves the answer is yes: an RFNBO-eligible UK PtL molecule counts as ReFuelEU-eligible synthetic aviation fuel and is fully tradable into the EU market. The commercial consequence is that a rational UK PtL producer routes marginal supply to the higher-clearing market first; the UK receives the residual. That is precisely the mechanism §3.4 develops for the UK residual-market picture, and it is why Exhibit 7 (p. 38) draws the thick supply arrow to the EU market and the thin one to the UK. The banker's follow-through is that any project-finance model of a UK PtL plant should not assume the UK buy-out as the effective clearing price for the plant's own output — the achievable price is the EU-clearing one, with a UK residual off-take at the buy-out only for whatever fraction the EU market has not absorbed.

Three consequences follow. First, the timing of Revenue Certainty legislation and instrument design becomes the direct gating item on the FID window that §3.5 flags for UK non-HEFA projects: the end-2027-to-end-2028 window is only reachable if the CfD mechanism is legislated in time and its strike prices are commercially adequate. Second, the pre-FID financing gap — Altalto's £30m spend before FID, NEXTGEN SAF's £15m Series A ask for FEED, and the developer-borne equity of the other UK projects — is a financing product distinct from the debt-and-mezzanine stack that project finance desks are set up to underwrite. It is developer-borne equity spent on FEED, environmental permitting, land acquisition and pre-FID engineering before the project has any secured contract cash flows to underwrite senior debt against, and it is entirely lost if the project does not reach FID. That risk profile — full downside if the project stops, no in-life cash flow, all-or-nothing dependence on a policy instrument that may or may not arrive — sits closer to venture capital than to traditional project finance, and the equity investors bearing it therefore price it materially wider than a project-finance (PF) senior desk would price the same aggregate exposure once the CfD is in hand. Bridging the gap between what venture capital charges for pre-FID risk and what project finance charges for post-FID debt is where a distinct financing structure — pre-FID facility, staged draws, milestone-triggered convertible instruments — needs to sit, and where a specialised transition-finance product could earn a return that neither pure equity nor senior debt captures cleanly today. Third, the co-timing of FIDs across the UK pipeline (§3.5) means that if one flagship project falls out of the window, the political case for the CfD mechanism itself weakens for the others — an implicit correlation across nominally independent projects that traditional project-finance covenant structures do not naturally capture. The leasing perspective on the aggregate cost trajectory that this bankability question drives is in §4.3; the underwriter's version of the co-timing exposure is in §4.4.

4.3 Leasing companies — fuel-cost pass-through and residual-value risk

For an aircraft leasing company, the §3 results reach the balance sheet through fuel-cost trajectories and their effect on airline lessees' willingness to pay for older versus newer airframes. The most direct anchor is the aggregate aviation fuel spend under NGFS Net Zero 2050 EU27+EFTA, which rises from roughly $30 billion per year today to something on the order of $430 billion per year at 2050 and stays near that level to 2100 — a roughly fifteen-fold increase in absolute annual spend on an aviation demand base that itself contracts by about 40 % over the same period (Exhibit 5 (p. 34)). The composition shifts through the same period from a fossil-jet-plus-carbon-levy stack to a nearly-all-SAF stack, with the carbon levy peaking around 2050 and declining as fossil residuals shrink.

That trajectory has three lease-relevant consequences. First, the fuel-cost pass-through position matters more than it does today. In a low-fuel-cost environment airlines absorb variance internally; in a high-fuel-cost environment the fuel line becomes a first-order determinant of route economics, and lessees push harder to renegotiate lease terms or shorten commitment windows. The lessor's ability to hold long-dated fixed-rate leases through a rising-fuel-cost regime depends on whether the underlying lease structures pass through the cost or absorb it, and on whether the counterparty's credit position is robust to a doubling and re-doubling of its per-hour operating cost.

Second, fleet renewal accelerates when fuel-cost trajectories are steeper than airframe efficiency improvement rates. If the SAF-plus-carbon-levy stack pushes per-tonne jet costs up faster than new-generation aircraft reduce fuel burn per available seat kilometre, the efficiency premium on the newer airframe grows and the residual value of an older airframe compresses faster than a straight-line depreciation curve would suggest. That is a residual-value risk sitting on the lessor's balance sheet, and it is a function of the scenario at least as much as of the specific fleet composition — the same aircraft carries a materially different residual under Net Zero 2050 versus Below 2°C by the mid-2040s.

Third, and connecting back to §4.2, the mandate-step timing under ReFuelEU and the UK SAF Mandate does not align with typical fleet-renewal cycles. Mandate percentages step up at 2030, 2035, 2040 and (in the UK) plateau at 2040, while lease terms typically run seven to fifteen years. A ten-year lease originated in 2027 crosses the 2030 mandate step and the 2035 step, and re-lease conditions at the end of that term depend on whether the compliance cost of operating an older-airframe lessee has moved materially against them by then. The underwriter's version of the same overlap, cast as a loss-trigger correlation across projects and lessees, is in §4.4.

4.4 Insurers — eligibility cliffs, single-vendor concentration, deployment-timeline slippage

By insurer in this subsection we mean the specialty markets that underwrite project-side exposures on SAF plants — construction-phase (Contractors' All Risks / Erection All Risks), operational-phase business interruption and property damage, offtake-default under Revenue Certainty structures, technology-obsolescence, and correlated common-mode failure. These are industrial-energy and infrastructure insurance lines, distinct from the traditional aviation-insurance market that underwrites hull, liability, war and terrorism on operating aircraft. Whether the same carrier writes both is a matter of firm-level line-of-business scope; the products are different underwriting exercises.

Traditional aviation insurance is materially insulated from the paper's mechanisms, because SAF's fuel-cost trajectory does not directly affect the risk of flying. Where the fuel-cost trajectory does touch the traditional aviation market is indirectly, via airline creditworthiness — a doubling and re-doubling of per-hour operating cost compresses margins at weaker carriers, which affects premium-collection risk on aviation cover and, more directly, the primary insured exposure for trade-credit and financial-guarantee insurance products written on airline receivables. That corridor between fuel-cost trajectory and airline credit is a real but distinct loss channel; it is not developed further here.

For the project-side insurer on SAF plants, then, the §3 results are read as loss triggers and correlation exposures rather than as portfolio positions. Three trigger classes stand out.

The first is eligibility-cliff exposure on RFNBO-eligible synthetic aviation fuel. The mechanism developed in §3.6 gives producers a strong incentive to engineer to exactly 26.7 gCO_2e/MJ — the dual-market credit-neutral point — but the paper argued there that 26.7 is a calculated design value under a specified methodology, not a physically-locked operating condition. Real-world operational variance around the design point is asymmetric across the two regimes: an excursion below 26.7 rewards the producer under the UK regime and is neutral under the EU regime, while an excursion above 26.7 crosses the ReFuelEU eligibility threshold and voids the fuel's counting as a synthetic aviation fuel for the duration of the excursion. The financial loss on a given excursion depends on volumes and prices of the affected consignment. For an operational-phase cover on an e-SAF producer, this is a distinct loss trigger — one whose expected loss is asymmetric around the reported design point — and it depends materially on plant-operating discipline plus methodology and reporting choices that no insurer traditionally has visibility on.

The second is single-vendor concentration on the Fischer-Tropsch side. §3.4 documented that the UK FT-MSW first-of-a-kind projects surfaced in the public record by 2026 all use the same reactor and catalyst package — Velocys microFTL™. The vendor itself is not distressed and is actively scaling manufacturing, but the concentration exposes the UK FT-MSW fleet to correlated technical, licensing and manufacturing-throughput risks that would not exist in a diversified vendor mix. For an insurer underwriting technology-obsolescence or common-mode-failure cover across the UK SAF book, this is a correlation exposure of the sort that traditional aviation reinsurance treaties do not naturally scope. The magnitude — plants representing an aggregate meaningful fraction of the UK 2030-2035 non-HEFA supply — is enough that a Velocys-specific event during commissioning or early operation would produce a correlated loss across multiple insured projects. The reasoned response is a treaty design that names the technology dependence explicitly and prices the correlation, rather than assuming it away in the individual-project underwrite.

The third is deployment-timeline bunching. Every UK non-HEFA FOAK we track has an FID target inside a 30-month window between end-2027 and end-2028, and a commercial operation target inside a 40-month window between end-2030 and 2034 (§3.5). Slippage on one project increases the political and commercial risk of slippage on the others — because the Revenue Certainty mechanism itself becomes less credible if the flagship projects it was designed to enable do not close on schedule (§4.2). For construction-phase insurance and offtake-default cover, this bunching implies a correlation structure that a traditional independent-project view of aviation risk does not capture. The insurer who ignores it is under-pricing the tail; the insurer who prices it correctly needs a view on the CfD mechanism's political viability as an underwriting input.

The four audience readings above are not mutually exclusive. Every allocator eventually asks a bankability question; every banker eventually asks a residual-value question; every leasing company eventually asks a compliance-cliff question; every insurer eventually asks a scenario-dispersion question. The subsection breakdown here is a matter of where each mechanism has its strongest first-order effect, and the cross-references make the second-order links explicit. The final subsection names the underlying risk factor that runs through all four.

4.5 Policy risk — the shared exposure across all four audiences

The four subsections above have treated four audiences as if they held four different exposures. There is a fifth reading that ties them together: policy risk — the risk that the governments whose ETS and mandate together create the SAF-versus-jet cost differential fail to close it through the instruments they alone control.

The mechanism that creates the SAF-versus-jet cost differential is the ETS itself — a policy instrument that generates substantial government revenue from allowance auctions. Under the EU ETS, that revenue accrues to Member States as national auction income, to the Innovation Fund and the Modernisation Fund via specific carve-outs, and — indirectly via CBAM — to the Social Climate Fund. Under the UK ETS, allowance auction revenue accrues to the UK Consolidated Fund. Closing the SAF cost gap that the same ETS creates is, in principle, potentially within reach of the revenue the ETS generates — either through direct subsidy (Advanced Fuels Fund awards in the UK, Innovation Fund grants in the EU) or through fixed-price offtake mechanisms (Contracts for Difference for SAF in the UK; equivalent instruments contemplated for the EU). We do not attempt the revenue-vs-gap arithmetic here — that turns on national-treasury share of auction income, cross-sectoral competition inside the Innovation Fund, and the political discretion around Fund allocations, none of which are quantified in this paper. In practice, alignment is imperfect and politically negotiated. The EU Innovation Fund's mandate is cross-sectoral: SAF projects compete for the same pot as steel, cement, chemicals and hydrogen. The Modernisation Fund's brief is regional rather than aviation-specific. National governments hold discretion over their share of auction revenue and are not required — or, in some cases, not fully briefed on a need — to earmark it for closing the SAF supply-side gap that their own mandate has opened. The industry's investment timeline, however, is not fungible: SAF plants need Revenue Certainty now to reach FID inside the 2027-2028 window (§3.5), even though the mandate they will serve accumulates ETS revenue over the twenty years thereafter. This mismatch — between when the revenue arrives and when the plants must be committed — is the political-economy edge that runs through every one of §§4.1-4.4.

That risk lands differently on each of the audiences already considered:

- For the institutional investor (§4.1), policy risk sits in the scenario probability weighting. NGFS Net Zero 2050 as a portfolio reference case presupposes that the orderly transition actually delivers — i.e. that ETS revenue is deployed in ways that support the mandated supply build-out. If it is not, aviation exposures priced to Net Zero 2050 mis-price the actual delivered trajectory.
- For the project financier (§4.2), policy risk is the CfD instrument's political viability. The banker's applicable test — will the CfD strike price cover engineered production cost — reduces to a question about whether the CfD gets legislated in time and is funded adequately.
- For the leasing company (§4.3), policy risk lands on the realisation of the modelled jet-fuel-plus-carbon-levy trajectory itself. The residual-value story in §4.3 depends on that trajectory rising steeply enough to make newer-airframe efficiency premia compress older-airframe residuals. Three policy variables drive that rise: the actual mandate percentages delivered (versus those written into the statute), the ETS carbon-price path, and the buy-out level (which sets the ceiling on what suppliers pass through to airlines). If any of the three falls short of the modelled Net Zero 2050 path — the mandate is softened, ETS ambition is reduced, the buy-out is raised — the actual fuel-cost trajectory is lower than the model implies, older airframes compress less than the residual-value forecast built off that path expects, and the lessor carries a mark-to-model loss directly attributable to policy delivery.
- For the insurer (§4.4), policy risk is the correlation node that ties the individual project underwrites together. A policy failure — CfD delay, mandate softening, buy-out revision — cascades across every UK non-HEFA SAF project simultaneously, converting what was priced as independent-project exposure into a single correlated loss trigger.

The four exposures are not four separate risks. They are four faces of the same underlying policy-delivery question, seen from four different balance sheets. An allocator, banker, lessor and insurer holding UK non-HEFA SAF exposure are all long the same variable: the willingness and technical capacity of the UK and EU governments to close, with allowance-auction revenue and CfD instruments they alone control, the cost gap their own mandate has opened. That is worth naming as a single risk factor rather than four separate ones. The chart-pack in §5 does not attempt to quantify it — policy delivery is not a modellable parameter in the NGFS scenario set — but the reader who has walked through §§4.1-4.5 is now equipped to price it as a cross-cutting overlay on the specific findings.

Sources: Institutional context: EU Carbon Removal Certification Framework (Regulation (EU) 2024/3012 as adopted); EU Innovation Fund and Modernisation Fund (Directive 2003/87/EC as amended, Articles 10a and 10d); Social Climate Fund (Regulation (EU) 2023/955); Advanced Fuels Fund (Department for Transport, UK); UK Contracts for Difference for Power BECCS (DESNZ).

---

5. DATA EXHIBITS

The eight exhibits below each read against a specific finding in §3 and a specific audience in §4. Charts are generated from the model tables and regenerate deterministically when the data version changes.

Exhibit 1. NGFS aviation demand under Net Zero 2050 and Below 2°C, EU27+EFTA and UK.

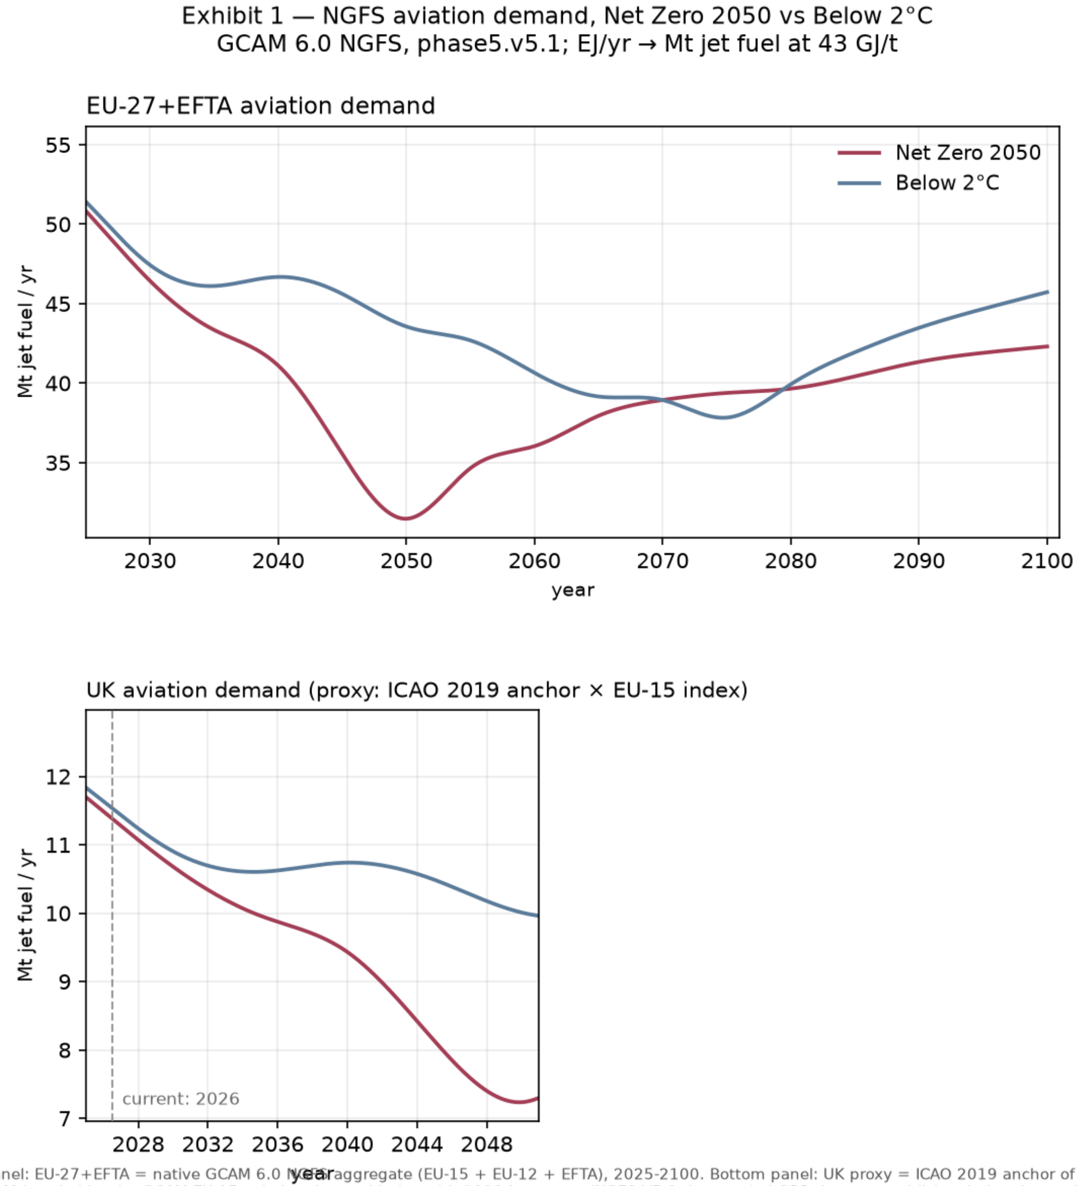

Top panel: EU-27+EFTA = native GCAM 6.0 NGFS aggregate (EU-15 + EU-12 + EFTA), 2025-2100. Bottom panel: UK proxy = ICAO 2019 anchor of 0.546 EJ/yr (≈12.5 Mt) scaled by the GCAM EU-15 aviation-demand index with 2020 base year (NGFS V5.1 downscaled GBR does not publish aviation demand); clipped at 2050 for visual consistency with other UK exhibits. The U-curve under Net Zero 2050 is a scenario-internal outcome of NGFS's NiGEM macro-feedback on income and modal choice — a normative not predictive reading (§1). Channels outside the model's frame (intra-EEA vs extra-EEA cost signal, ATC efficiency, airframe/engine improvements, short-haul electrification) would compound with the contraction where realised (§6).

Vertical axis: NGFS `Final Energy|Transportation|Aviation`, converted from exajoules per year to equivalent million tonnes of jet fuel at 43 GJ/t. Horizontal axis: 2025 to 2100 for the EU27+EFTA panel (top); 2025 to 2050 for the UK panel (bottom). Each panel carries both scenarios as two separate lines.

The reader is looking at a scenario-internal outcome of NGFS's macro-economic layer (NiGEM), not an aviation demand forecast. Under Net Zero 2050, EU27+EFTA aviation demand contracts by about 40 % from 2025 to 2050 and partially recovers toward 2100 — a U-curve produced by NiGEM's feedback on income and modal choice under the scenario's assumed globally coordinated carbon-price trajectory. Under Below 2°C the contraction is softer and later; the U is shallower. The two shapes are normative reflections of what the orderly transition requires, not predictions of what current policy is delivering (§1). They also exclude several channels through which demand or per-flight fuel-burn could evolve independently: the concentration of the mandate-plus-ETS cost signal on intra-EEA flights only (extra-EEA international departures fall under the weaker CORSIA regime, §4); efficiency gains from air-traffic-control procedure optimisation; per-flight fuel-burn reductions from continued airframe and engine improvements; and the eventual electrification of short-haul routes. Each of those channels would compound with the NGFS contraction to the extent it materialises. Industry-body demand outlooks — IATA's passenger-and-cargo twenty-year forecast, ICAO Long-Term Aspirational Goals, Airbus Global Market Forecast, Boeing Commercial Market Outlook — routinely project continued net growth in aviation activity to 2050. Their forecasts and the NGFS scenarios sit inside different modelling frames and answer different questions; neither is "correct" independently of the framing question the reader is asking.

A note on the UK derivation. The UK line shown in the bottom panel is a proxy, not a native NGFS output. NGFS V5.1's downscaled GBR product does not publish an aviation-specific demand variable, so the engine takes the ICAO Air Transport Reporting 2019 UK aviation activity figure as an absolute anchor and scales it forward using the native GCAM EU15 aviation index — which does extend to 2100. The reasoning is that UK aviation demand plausibly follows the wider EU15 aviation index in shape but is grounded in its own pre-COVID scale; the full derivation is set out in §2. Because this is a proxy, the UK panel is read with a light interpretive touch — it shows the shape of the demand response NGFS would imply for the UK given no distinct downscaled series, not a stand-alone UK aviation forecast. The derived UK series could in principle be extended to 2100 via the EU15 index, but this exhibit stops it at 2050 for visual consistency with the other UK panels in Exhibits 2-5, which are bounded by NGFS's downscaled-GBR truncation at that year and cannot be extended.

Reads directly against §1 and §2; the leasing-company implications of the fuel-cost trajectory these demand curves scale are developed in §4.3.

Exhibit 2. Stacked SAF supply with total and e-SAF sub-mandate lines.

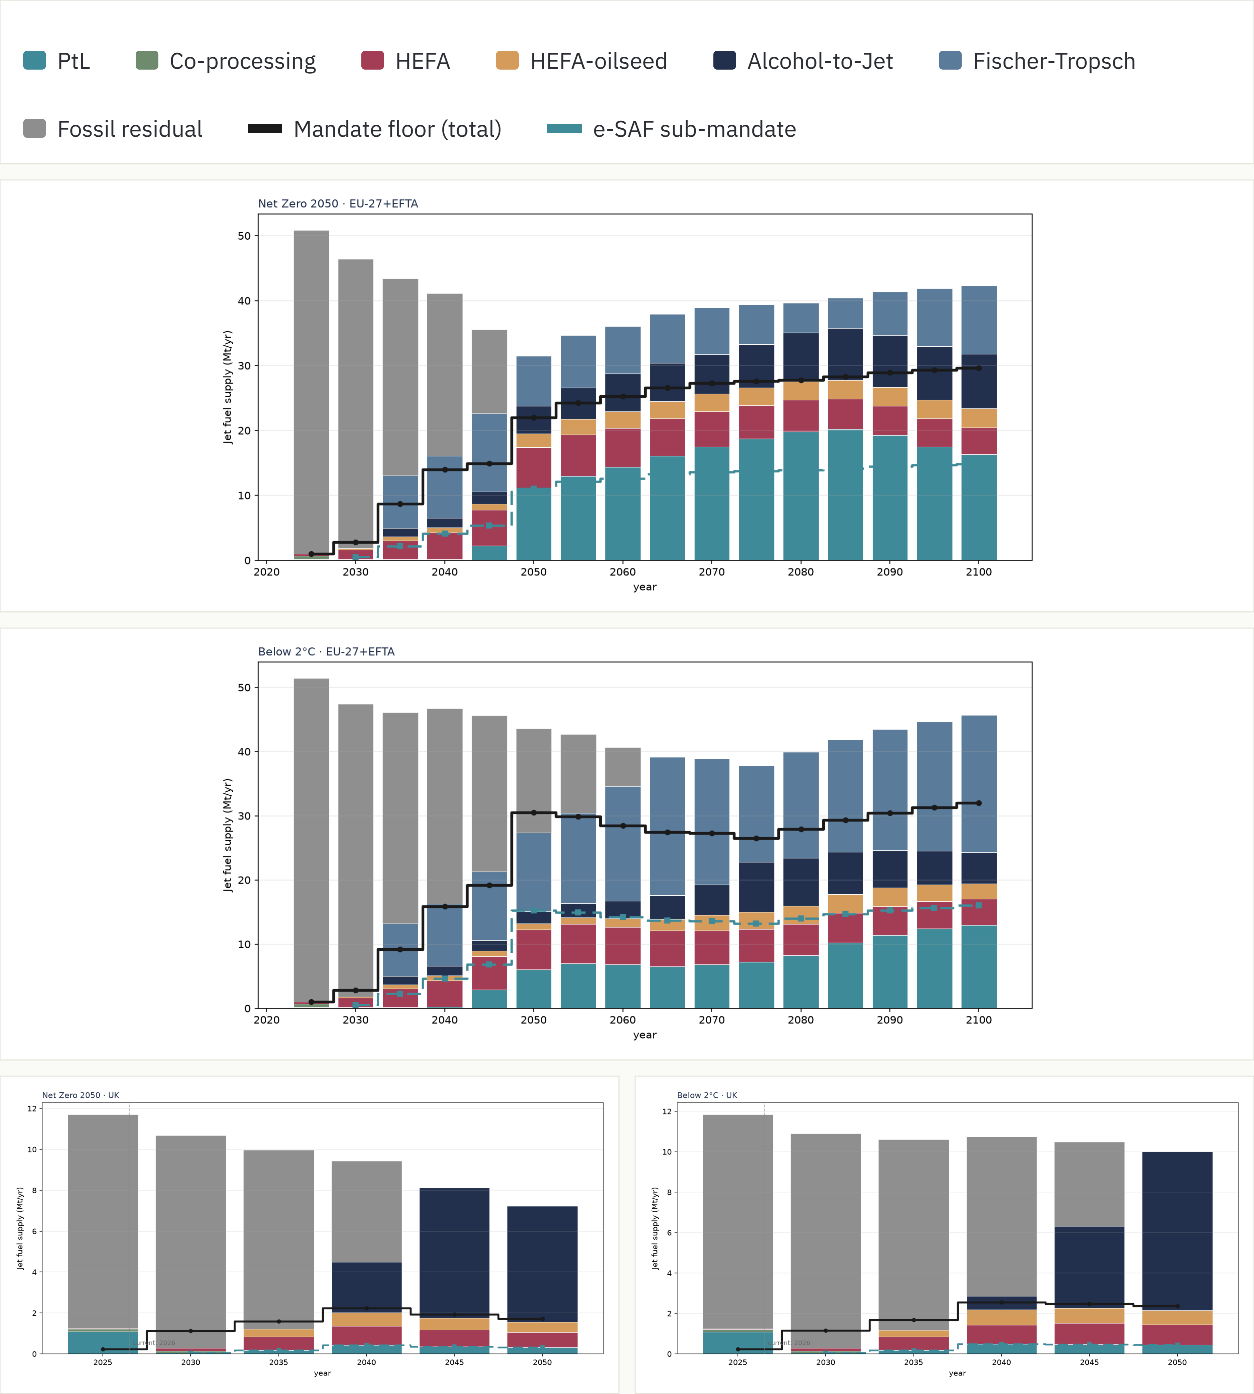

Vertical axis: SAF supply in million tonnes per year, stacked by pathway (Power-to-Liquid at the bottom of each column, then Co-processing, HEFA, HEFA-oilseed, Alcohol-to-Jet, Fischer-Tropsch) with the fossil jet residual — aviation demand minus total SAF supply — shown as a grey band at the top of each stack. Horizontal axis: 2025 to 2100 for the EU27+EFTA rows, 2025 to 2050 for the UK bottom row. Pathway colours are the paper's standing palette (PtL teal, Co-processing olive, HEFA crimson, HEFA-oilseed amber, Alcohol-to-Jet navy, Fischer-Tropsch slate) held constant across Exhibits 2, 3, 4 and 5.

Two overlays run across the pathway stack. The heavy solid line with markers is the total-SAF mandate floor (ReFuelEU: 2 % → 70 % over 2025-2050, extended in the engine to 100 % by 2080; UK: 2.0 % → 23.7 % over 2025-2040 and plateaued at 23.7 % thereafter). The dashed line is the sub-mandate floor for synthetic aviation fuel — e-SAF for the EU, PtL for the UK — drawn in teal to match the PtL bar underneath. The mandate floor doubles as the mandate-limited SAF supply level — the volume the statute would require if a supplier procured exactly the mandate percentage and no more — so no third overlay is needed to make that reading visible.

Reading the pathway stack correctly. The stack is our engine's aviation-allocation of the NGFS resource pool, not the NGFS scenario itself. NGFS scenarios are normative, globally-balanced descriptions of what an orderly transition to 1.5 °C (Net Zero 2050) or 2 °C (Below 2°C) would look like at the whole-system level; they give the resource pool (biomass, hydrogen, CO_2, electricity) and demand-side signals (carbon price, macro-feedback), but they do not specify aviation's SAF share within that pool. Multiple aviation-allocation trajectories are consistent with the same scenario. The stack shown here reflects one such allocation — the one the engine produces, in which aviation receives its scenario-derived share of each shared input (§2) and, within that share, the merit-order fills SAF up to the resource cap. So the stack answers "how much SAF could aviation produce under our specific allocation rules given the scenario's resource pool", not "how much SAF the scenario says aviation will produce".

The gap between the pathway stack and the solid mandate line is therefore discretionary aviation supply — SAF the scenario's resource pool could support under our allocation rules if built, but that current policy does not require. Whether that discretionary supply is built depends on whether corporate SAF procurement, airline-driven demand, or Revenue Certainty mechanisms pull volumes above the statutory floor. It is not the "orderly-transition-vs-current-policy" gap. Whether the actual trajectory of SAF supply is on-track to the scenario's climate objective, on a different but still-viable path, or behind schedule is a separate question that reality-vs-scenario comparison would need to answer — a task deferred to a companion working paper (WP05, in scoping) that will use accumulating compliance data from 2025 onwards.

Given that framing, three substantive readings emerge from the chart. First, the total SAF mandate line can be over-fulfilled from 2030 onwards under both scenarios under our allocation rules — the aviation-allocated resource pool would support SAF above the 70 % ReFuelEU floor by 2050, whether that supply is actually built depends on a cost-vs-fossil-plus-carbon-levy check the engine does not yet apply (§6). Fischer-Tropsch carries most of the load in the middle of the century because lignocellulosic residue is the least contested of the shared inputs and its NGFS-published price does not rise as steeply as ethanol, hydrogen or captured CO_2. Second, the sub-mandate is where the resource pool actually binds inside our allocation: the teal PtL bar at the bottom of the stack is below the dashed sub-mandate line through to 2050 under Net Zero 2050 EU27+EFTA — that is the RFNBO-eligible CO_2 supply constraint developed in §3.1 as the primary substantive finding, and it is present even at maximum aviation-allocation. Between 2050 and 2080 PtL rises to sit exactly at the sub-mandate line (mandate-floor-bound rather than resource-bound), and after 2080 the gap re-opens as NGFS's `Carbon Sequestration|CCS|Biomass|Energy|Supply|Liquids` contracts. Third, the gap between the solid mandate line and the pathway stack is where the paper's institutional-investor and banker readings sit (§4.1 and §4.2) — under our allocation rules, the pool says what aviation could produce, the mandate says what statute requires, and the delta is discretionary supply that only gets built if demand pulls above the mandate.

HEFA-oilseed is drawn as a visually distinct amber layer to make clear that it is produced but not mandate-eligible. A rational supplier produces oilseed-based renewable jet (plus its co-products — protein meal, renewable diesel, naphtha) whenever the joint economics clear the market, but under RED III and SI 2024/1187 the food/feed-crop feedstock excludes it from counting toward the mandate. Separating it visually keeps the mandate-gap read honest: what the reader wants to compare against the mandate lines is the mandate-eligible stack (everything except the amber layer), not gross physical SAF production. Aggregate gross SAF plus fossil residual sums to the aviation demand ceiling by construction.

The UK panels stop at 2050 because the downscaled GBR variables feeding the UK supply algorithm are truncated by NGFS V5.1 at that year (§2). The plateau in the UK mandate lines from 2040 is not a modelling shortcut: SI 2024/1187 sets 23.7 % Main Obligation and 4.5 % PtL sub-obligation as held-rate figures for every obligation period from 2040 onwards. The UK panels also do not carry the buy-out cap line — that overlay is reserved for Exhibit 4 (p. 32) (cost per pathway), where it belongs on a price axis rather than a volume axis.

Reads directly against §3.1 (total mandate can be over-fulfilled from the resource pool, sub-mandate is where the risk sits) and §3.2 (binding-input migration). The banker's view of the mandate-gap-vs-buy-out question is in §4.2; the leasing-company reading of the aggregate fuel-cost dynamics that scale off this supply-and-residual picture is in §4.3.

Exhibit 3. Binding-resource map — which input binds each pathway each year, per scenario × region.

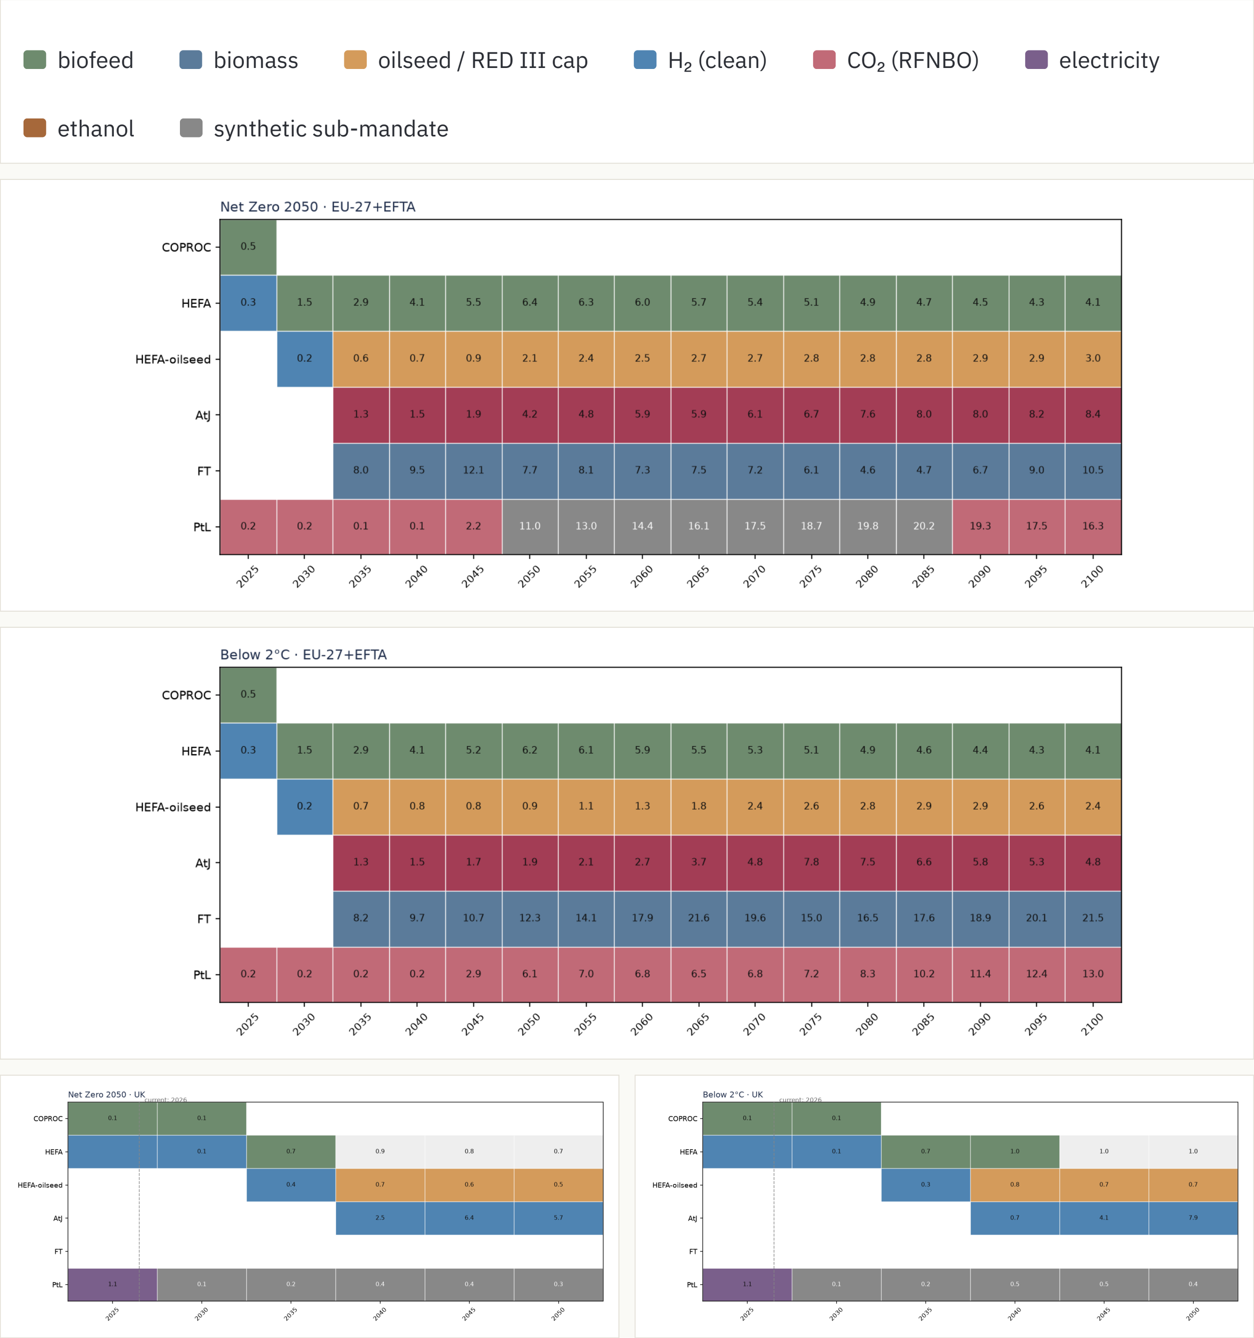

The chart is a matrix heatmap. Rows are the six SAF pathways (Power-to-Liquid, Co-processing, HEFA, HEFA-oilseed, Alcohol-to-Jet, Fischer-Tropsch). Columns are years (2025 to 2100 for EU27+EFTA; 2025 to 2050 for UK). Each cell is coloured by which input is binding for that pathway in that year — the constraint the model identifies as the operative limit on how much SAF that pathway could produce. Inside each non-empty cell, the number in white or black is the SAF tonnage that pathway did produce that year, in million tonnes per year, so the reader sees the binding input and the volume together. Empty cells mean the pathway produced no SAF that year — either merit-order exhaustion of the pathway's shared input, or a hard cap already binding on a pathway earlier in the order.

The binding-input palette uses seven colours that are deliberately partly consistent with the pathway palette in Exhibit 2 (p. 27), so the reader who has learned the pathway colours can carry three of them across without checking a legend:

- biofeed (green) — waste oils and animal fats under RED III Annex IX-B. Same green as COPROC in E2 (COPROC's binding input is biofeed).
- biomass (slate-blue) — lignocellulosic residues. Same slate-blue as FT in E2 (FT's binding input is biomass).
- oilseed (amber) — food/feed crops. Same amber as HEFA-oilseed in E2 (its binding input at scale is the RED III food/feed cap on oilseed).
- h2 (mid-blue) — clean hydrogen; distinct from FT's slate-blue.
- co2 (rose) — RFNBO-eligible biogenic CO_2 for PtL.
- synthetic_sub_mandate (grey) — the operative constraint is not a physical resource but the ReFuelEU e-SAF floor or the UK PtL sub-obligation. The pathway is producing exactly what the mandate requires and no more, because the resource pool has headroom above the floor.
- red_iii_food_feed_cap (amber, matches oilseed) — the RED III Article 25 7 %-of-transport-energy cap on food/feed feedstocks; caught in the same amber to signal that the physical feedstock is present but the cap rather than the pool is doing the binding.

Three phases of the binding-input migration developed in §3.2 are visible directly on the map. Phase 1, early 2030s: the FT row shows hydrogen-mid-blue for a short period — that is the "H_2-bound FT" finding (F4), a non-obvious early-decade result driven by NGFS's slow clean-hydrogen ramp. Phase 2, mid-2030s to 2060: the FT row shifts to slate-blue biomass-bound; the HEFA-oilseed row moves to amber cap-bound as the rising aviation share of the food/feed pool crosses the RED III 7 % cap in the 2050s. Phase 3, 2050-2080 and then again post-2080: the PtL row alternates between rose (co2-bound) pre-2050, grey (synthetic_sub_mandate-bound) 2050-2080, and rose again post-2080 as NGFS's `Carbon Sequestration|CCS|Biomass|Energy|Supply|Liquids` contracts. That three-cell colour sequence on the PtL row is the paper's sub-mandate story compressed into a single visual glance.

Reads directly against §3.2 (binding-input migration H_2 → biomass → CO_2). Where the binding switch matters for the banker or the insurer is developed in §4.2 and §4.4 respectively.

Exhibit 4. Cost per pathway under the NGFS carbon-price trajectory, with the UK SAF Mandate buy-out overlay on the UK panels.

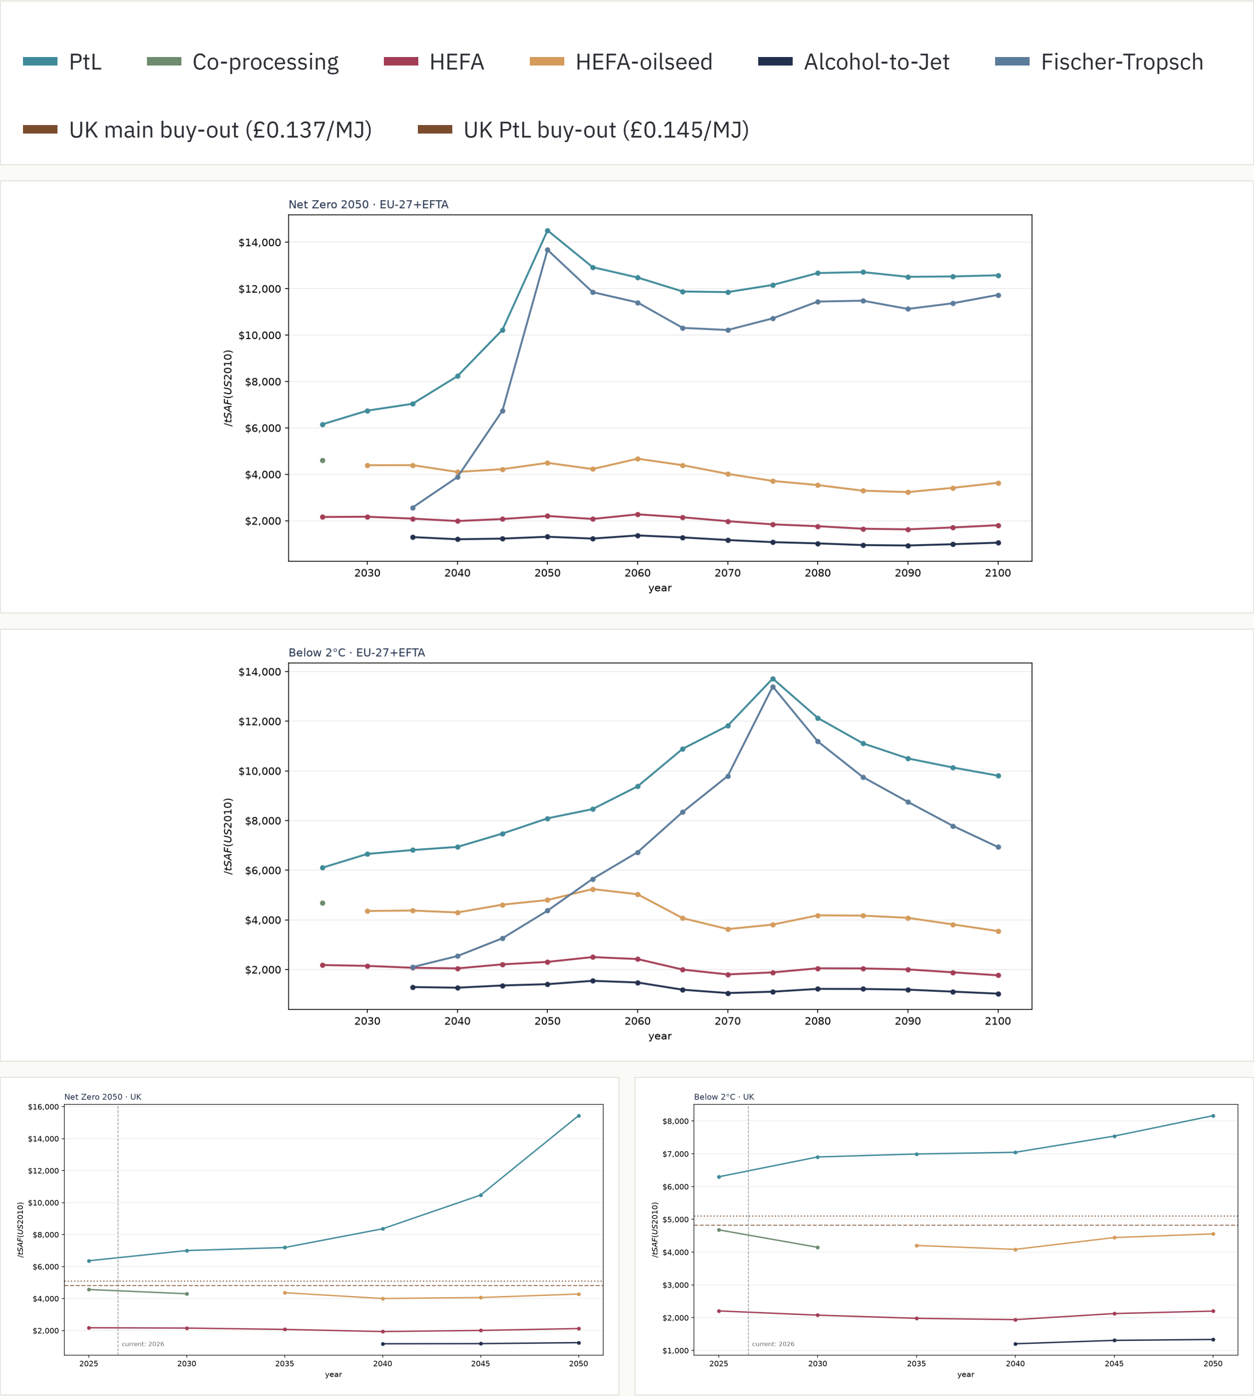

Vertical axis: scenario-internal SAF production cost per pathway, in US$ (2010 base year) per tonne of SAF. Horizontal axis: 2025 to 2100 for EU27+EFTA rows; 2025 to 2050 for the UK bottom row. Six lines per panel, one per pathway, using the same palette introduced in Exhibit 2 (p. 27) (PtL teal, Co-processing olive-green, HEFA crimson, HEFA-oilseed amber, Alcohol-to-Jet navy, Fischer-Tropsch slate-blue). The cost is the sum, for each year, of the pathway's input intensities multiplied by the NGFS-published price of each input at the relevant scenario × region. It is a variable-OPEX floor: CAPEX, fixed OPEX, and downstream blending are excluded, so the numbers represent the minimum a plant would need to earn on the neat SAF product before capital recovery is considered (§2).

Three substantive readings follow.

First, the rank ordering is stable. Under either scenario, at either region, and across every year the model runs, the ranking on paper variable-OPEX runs Alcohol-to-Jet and HEFA at the low end, Co-processing and HEFA-oilseed above them, Fischer-Tropsch materially higher and rising toward mid-century, and PtL well above everything else. That AtJ sits at the bottom on paper economics but is not the merit-order winner in practice is precisely the paper-vs-deployable split developed in §3.5 — AtJ's low modelled cost reflects ethanol input pricing and a favourable stoichiometry, but its deployment is bottlenecked by ethanol availability at aviation-scale volumes and by process yield rather than by feedstock cost.

Second, the PtL curve dominates the picture — and it dominates for a specific reason. The teal PtL line rises sharply through the middle of the century under Net Zero 2050, reaching roughly a doubling of its own base variable cost around the 2050 carbon-price peak. That rise is not driven by a rising cost of clean hydrogen or renewable electricity alone. It is driven by the opportunity-cost mechanism developed in §3.3: as `Price|Carbon` rises, biogenic CO_2 suppliers can earn a Carbon Removal Certification Framework credit at that price by storing the CO_2 instead of selling it to a PtL refinery, so the effective CO_2 input cost to PtL rises to match `Price|Carbon`. Under Below 2°C the same mechanism operates but is materially smaller — `Price|Carbon` peaks lower and later, so the PtL line rises to roughly 1.5× its base cost at the softer scenario's own peak, rather than doubling. The scenario dispersion on this pathway is the largest single-line movement in the exhibit.

Third, the UK panels carry two dashed horizontal buy-out lines — SI 2024/1187 Article 21 sets two statutory buy-out prices, one for the Main obligation at £0.137/MJ (≈ £4,700 / tonne SAF ≈ ~\$7,450/t US\$2010 using the CI-formula conversion set out in §3.4) and one for the PtL sub-obligation at £0.145/MJ (≈ £6,235 / tonne SAF ≈ ~\$7,900/t). The Main buy-out is the operative ceiling for all mandate-eligible SAF (HEFA, AtJ, FT, COPROC and their derivatives); the PtL buy-out is the higher ceiling that applies only when the fuel is being counted toward the PtL sub-obligation. Any pathway whose modelled cost sits above that dashed line is one for which a rational UK fuel supplier would pay the buy-out to the Treasury rather than procure the fuel (the argument developed in §4.2). The PtL line sits above the buy-out for most of the 2030s and often thereafter. The AtJ and FT lines cross the buy-out mid-decade in Net Zero 2050 and stay above it for extended periods. Only HEFA and Co-processing sit reliably below the buy-out. The chart makes the “market-price ceiling” argument in §4.2 visually concrete: for non-HEFA UK SAF, the modelled cost is above the price the market clears at, so the plants only reach FID if Revenue Certainty pays the difference.

A note on the 2025-2029 UK PtL curve. Under §3.4's analysis, UK CO_2 suppliers face no removal-credit alternative until the UK ETS Authority integrates greenhouse-gas removals from end-2029, so the opportunity-cost mechanism does not bite in the UK in that window. The PtL line in the UK panels reverts to its engineering-cost floor for those years — visibly lower than the corresponding EU27+EFTA panels above — but is bounded above by the same £0.145/MJ buy-out cap regardless. The UK is paper-cheap for PtL production cost through to 2029 while being effectively zero for PtL supply, precisely because no CO_2 suppliers invest without a removal-credit revenue path.

Reads directly against §3.3 (opportunity-cost mechanism), §3.4 (UK two-sided bound), §3.5 (paper vs deployable), §4.2 (bankability), and §4.4 (insurer-relevant eligibility-cliff exposure). The absence of a CAPEX/OPEX layer is a Phase 2 extension flagged in §6.

Exhibit 5. Aggregate aviation fuel spend decomposition — SAF, fossil residual, carbon levy.

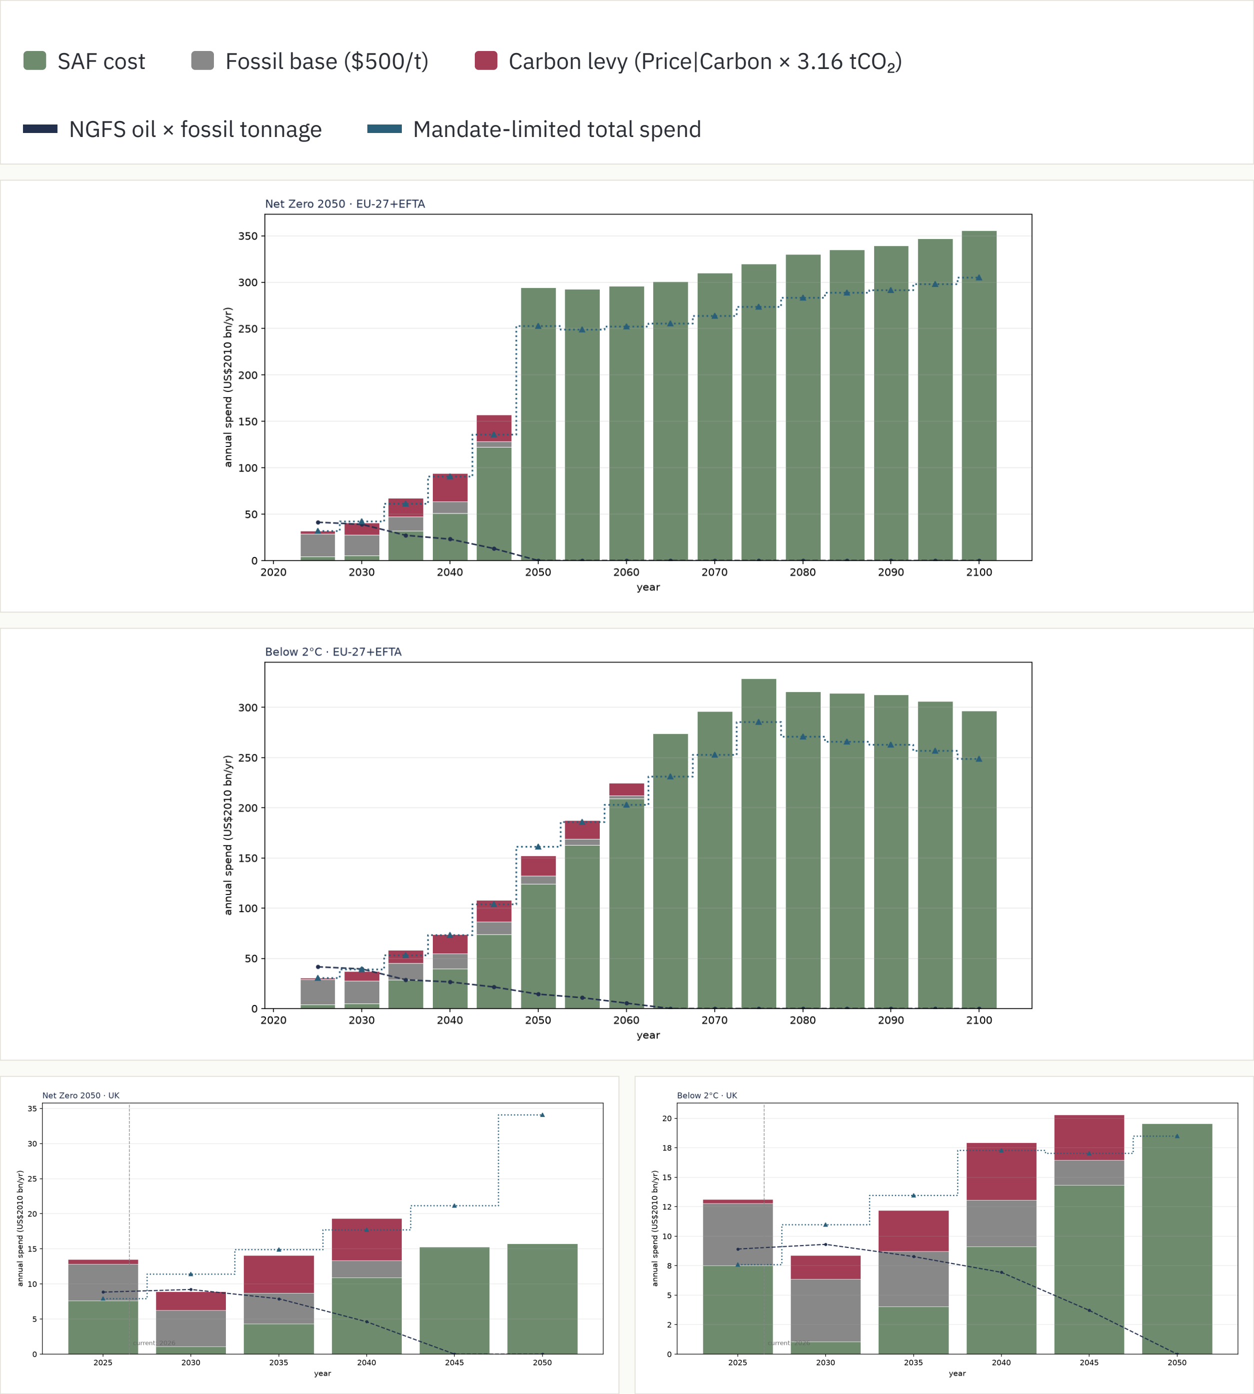

Vertical axis: aggregate annual aviation fuel spend, in US$ (2010 base year) billions per year. Horizontal axis: 2025 to 2100 for EU27+EFTA rows; 2025 to 2050 for the UK bottom row. Each column is a stacked bar decomposing the year's total spend into three components:

- SAF cost stack — for each of the six pathways in the merit-order allocation, the pathway's tonnage multiplied by its variable-OPEX $/t SAF from Exhibit 4 (p. 32). Coloured by pathway using the standing palette (PtL teal, Co-processing olive-green, HEFA crimson, HEFA-oilseed amber, Alcohol-to-Jet navy, Fischer-Tropsch slate-blue). CAPEX and fixed OPEX are excluded, so this component is a floor rather than a full cost (§2).
- Fossil residual base cost — aviation demand minus total SAF, priced at a fixed US\$500 per tonne of fossil jet (Concawe / IEA refined-jet base — a deliberate simplification, held con-

stant across the century so the reader can decompose the levy separately). Coloured light grey to match the fossil residual band in Exhibit 2 (p. 27). A dashed reference line elsewhere in the chart shows what the fossil residual would cost if priced at NGFS `Price|Secondary Energy|Liquids|Oil` × the fossil-jet lower heating value instead — a sanity-check overlay, not the stacked component.

- Carbon levy on fossil residual — the fossil residual tonnage multiplied by 3.16 tCO_2 per tonne of fossil jet and by `Price|Carbon` for that scenario × year. Coloured dark grey to sit visually above the fossil base within each stack.

Three substantive readings follow.

A framing note that carries from Exhibit 2. The SAF cost component in the stack is priced against our engine's diagnostic aviation-allocation of the NGFS resource pool — the same "how much SAF could aviation produce under our specific allocation rules given the scenario's resource pool" that E2 shows in volume terms. It is not a projection of what SAF supply the market will actually deliver, and it is not the SAF cost under mandate-limited compliance. The mandate-limited alternative would show a materially smaller SAF cost component (near-mandate-percentage of aviation demand) and a correspondingly larger fossil residual base cost plus carbon levy through the 2020s and early 2030s — closer to the 2-4 % UK SAF share actually being delivered in 2025-2026. The chart pack that accompanies this paper carries both readings; the caption below and the substantive discussion in §4.3 refer to the diagnostic-allocation reading unless explicitly noted.

First, the aggregate spend rises sharply under the diagnostic allocation. Under NGFS Net Zero 2050 EU27+EFTA, the total climbs from roughly $30 billion per year at the 2025 anchor to roughly $430 billion per year in 2050 and stays near that level through to 2100 — a fifteen-fold increase in absolute annual spend on an aviation demand base that itself contracts by about 40 % over the same period (Exhibit 1 (p. 25)). Per-tonne fuel cost rises roughly twenty-fold. Under Below 2°C the trajectory is materially flatter — the peak is lower and later — but the same directional shift is present.

Second, the composition of the stack shifts through the century. At the 2025 anchor, spend is dominated by the fossil residual base cost; by 2050 the composition is dominated by SAF cost plus the carbon levy on a shrinking fossil residual; by 2100 the spend is nearly all SAF cost, because fossil jet has been almost entirely displaced under our allocation and the carbon levy component collapses to near zero. The transition years — roughly 2035 to 2060 — are where all three components are simultaneously large, and where the peak spend lives.

Third, the carbon levy peaks around 2050 and then declines. The levy is the product of two variables moving in opposite directions: `Price|Carbon` rising through to 2050 and then easing back, and the fossil residual tonnage shrinking as SAF displaces it. The rising side of the levy is what drives the 2035-2060 peak; the falling side is what makes the 2100 stack nearly all-SAF. This mechanism is referenced in §4.3's residual-value discussion — for a leasing company, the peak of the levy is where fuel-cost trajectories are steepest and where fleet-renewal pressure on older airframes is strongest.

The UK panels are smaller in absolute magnitude — roughly 8-10 % of the EU27+EFTA numbers, reflecting UK aviation demand as a share of Greater Europe — and stop at 2050 (§2). The composition shift is broadly similar in shape to the EU panels, but the plateau in the UK Mandate percentages from 2040 (§3.4) means the SAF cost component of the UK stack does not continue to grow at the EU rate through the 2040s. The UK stack's composition in 2050 is therefore relatively heavier on fossil-plus-levy than the EU27+EFTA panel above it.

Reads directly against §3.1 (the mandate translating into a nearly-all-SAF stack post-2080), §4.3 (leasing-company residual-value dynamics), and §6 (the CAPEX/OPEX exclusion is a Phase 2 extension that would raise the SAF component further). The composition shift also underlies §4.1's institutional-investor reading — a book long fossil-jet resilience does better under Below 2°C's softer trajectory, and a book long PtL producer margins does worse under it.

Exhibit 6. ReFuelEU Aviation vs UK SAF Mandate — side-by-side comparison.

Dimension	ReFuelEU Aviation (EU27+EFTA)	UK SAF Mandate
Legal instrument	Regulation (EU) 2023/2405 (Annex I schedule)	SI 2024/1187 (The Renewable Transport Fuel Obligations (Sustainable Aviation Fuel) Order 2024)
In force from	1 January 2025	1 January 2025
Obligated party	Aviation fuel suppliers into Union airports	Aviation fuel suppliers into UK airports, above the 15.9 TJ/yr small-supplier reporting threshold
Geographic scope of the obligation	Fuel uplifted at Union airports (intra- and extra-EEA departures)	UK aviation fuel supply above the threshold
Main mandate 2025	2.0 %	2.0 %
Main mandate 2030	6.0 %	10.6 %
Main mandate 2035	20.0 %	15.9 %
Main mandate 2040	34.0 %	23.7 % (plateaued from this year onwards)
Main mandate 2050	70.0 %	23.7 % (held rate)
Synthetic sub-mandate 2030	1.2 % biennial average, 0.7 % annual minimum	0.6 %
Synthetic sub-mandate 2035	5.0 %	1.8 %
Synthetic sub-mandate 2040	10.0 %	4.5 % (plateaued from this year onwards)
Synthetic sub-mandate 2050	35.0 %	4.5 % (held rate)
Statutory buy-out mechanism	None. Non-compliance triggers administrative penalties under each Member State's implementation	Main obligation £0.137/MJ (~£4,700 / tonne SAF); PtL sub-obligation £0.145/MJ (~£6,235 / tonne SAF ≈ ~$7,900/t)

Dimension	ReFuelEU Aviation (EU27+EFTA)	UK SAF Mandate
Feedstock cap on HEFA / Annex IX-B	Annex IX Part B feedstocks capped at 1.7 % of Member State transport energy (multi-sector; SAF's share is negotiated)	SI 2024/1187 Article 22: HEFA-derived SAF capped at 100 % of the Main obligation pre-2027, tapering to 42.16 % by 2040 and held there
Food-and-feed-crop exclusion	Excluded from the mandate under Reg 2023/2405 Article 3(5), consistent with RED III Annex IX	SI 2024/1187 Article 25: food/feed-crop-derived SAF capped at 7 % of aviation transport energy
GHG intensity structure	70 % lifecycle savings threshold against the 89 gCO$_2$e/MJ fossil comparator = 26.7 gCO$_2$e/MJ eligibility ceiling for synthetic aviation fuels	CI-factor formula in Article 6(4): certificate quantity scaled by (CF − CS) / (CF − CR) with CF = 89 gCO$_2$e/MJ and CR = 26.7 gCO$_2$e/MJ reference intensity. CR is a scaling reference, not an eligibility floor
Compliance certificate	Proof-of-sustainability via voluntary schemes recognised by the Commission; retired against the mandate share	SAF Certificates issued by the DfT Administrator; retired against the obligated volume
Paired ETS obligation on aircraft operator	EU ETS Aviation, Directive 2003/87/EC as amended — allowance surrender on intra-EEA emissions	UK ETS Aviation under the Greenhouse Gas Emissions Trading Scheme Order 2020 — allowance surrender on flights within scope
Sources	Regulation (EU) 2023/2405 Article 3, Annex I; RED III (Directive (EU) 2023/2413) Article 29a(1)	SI 2024/1187 Articles 6, 21, 22, 25 and Schedule 1; DfT SAF Mandate Compliance Guidance 2026 §5.31

Three structural readings follow from the comparison.

First, the mandate trajectories are shaped very differently in time. The UK is materially more ambitious near-term — at 2030 the UK Main obligation is 76 % higher than ReFuelEU's in percentage terms, and the PtL sub-obligation kicks in earlier — but plateaus from 2040 onwards. ReFuelEU keeps rising through to 2050 and reaches roughly three times the UK's held rate by mid-century. Near-term compliance pressure falls harder on UK-obligated suppliers; long-term SAF supply build-out has to happen mostly to serve the EU. This is the asymmetry §3.4 develops and §4.2 turns into a Revenue-Certainty timing question.

Second, the buy-out mechanism differs materially in structure, not just in level. The UK carries a statutory buy-out at £0.145/MJ for PtL (roughly $7,900 / tonne SAF) that a rational supplier can and will pay rather than procure fuel above that price — the mechanism that produces the market-price-as-ceiling reading in §4.2. The EU has no statutory buy-out: non-compliance triggers administrative penalties under each Member State's implementation, priced above the UK's ceiling for any commercially serious supplier. The consequence — cross-border arbitrage on e-SAF molecules routing to the higher-clearing EU market — is developed in §3.4.

Third, both regimes converge on 26.7 gCO$_2$e/MJ but for different reasons. Under ReFuelEU, 26.7 gCO$_2$e/MJ is the eligibility ceiling for synthetic aviation fuel (a 70 % lifecycle-savings threshold against the 89 gCO$_2$e/MJ fossil comparator, from RED III Article 29a(1)). Under the

UK Mandate, 26.7 gCO$_2$e/MJ is the reference intensity CR in the CI-factor formula — a scaling anchor rather than an eligibility floor. The coincidence is what §3.6 calls the dual-market credit-neutral point — a producer engineering to exactly 26.7 clears the EU eligibility threshold and earns exactly 100 % of the UK notional credit, which is why NEXTGEN SAF and its peers land there by design. The two regulatory tests happen to point at the same physical number for different reasons, and the engineering optimisation that emerges is not accidental.

Reads directly against §3.4 (UK trajectory and buy-out ceiling), §3.6 (dual-market credit-neutral point), §4.2 (bankability + Revenue Certainty window), and §4.5 (political-economy timing of ETS revenue vs FID window).

Exhibit 7. UK as a residual e-SAF market — buy-out price ceiling, cross-border arbitrage, and the EU-absorbs-supply illustration.

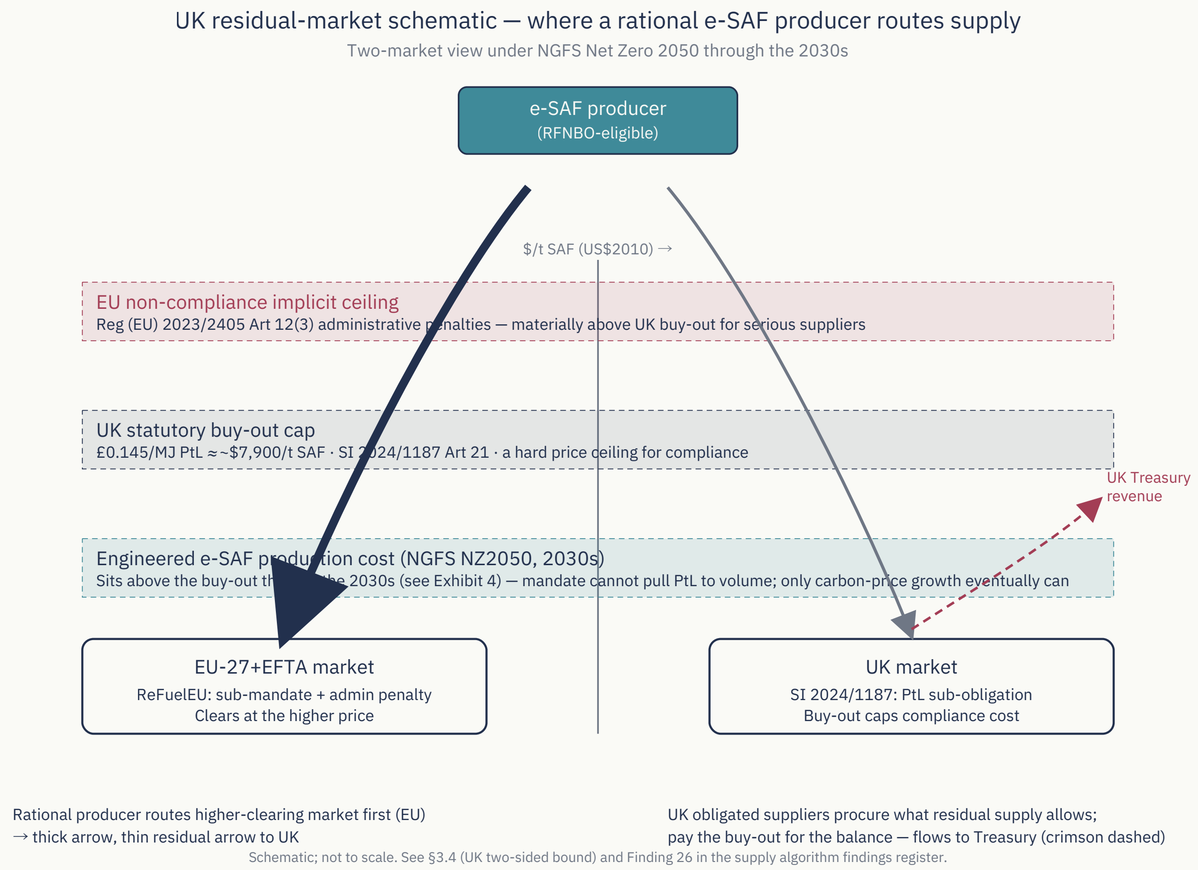

UK residual-market schematic — where a rational e-SAF producer routes supply

Two-market view under NGFS Net Zero 2050 through the 2030s

Rational producer routes higher-clearing market first (EU)
→ thick arrow, thin residual arrow to UK

UK obligated suppliers procure what residual supply allows;
pay the buy-out for the balance — flows to Treasury (crimson dashed)

Schematic; not to scale. See §3.4 (UK two-sided bound) and Finding 26 in the supply algorithm findings register.

The diagram illustrates the cross-border arbitrage mechanism that produces the UK-as-residual-market reading in §3.4. At its centre are two markets — the EU27+EFTA and the UK — connected by a common pool of e-SAF supply. Three price levels are shown on a vertical axis, from top to bottom:

- The EU non-compliance implicit ceiling — administrative penalties under Reg (EU) 2023/2405 Article 12(3), priced by each Member State's implementation. For any commercially serious supplier, the combined exposure (financial penalty plus back-billing plus supply-chain stigma) prices materially above the UK statutory buy-out.
- The UK statutory buy-out ceiling for the PtL sub-obligation: £0.145/MJ ≈ approximately $7,900 per tonne SAF (SI 2024/1187 Article 21).

- A producer's engineered e-SAF production cost, which sits above both under NGFS Net Zero 2050 in the 2030s (Exhibit 4 (p. 32)).

An e-SAF producer with a molecule to sell chooses where to route it. Because the EU market clears above the UK buy-out, a rational producer routes the molecule to the EU first — earning the higher clearing price and closing that market's contractual position. Only what the EU does not absorb — because of contracting counterparties reaching their sub-mandate targets, seasonality of demand, or logistical constraints — is available to the UK market. UK obligated suppliers then face a residual pool of e-SAF supply, procure what they can at prices up to the buy-out, and pay the buy-out to the Treasury for anything not procured.

The empirical statement supporting the mechanism appears verbatim in LanzaTech's SAF Café presentation of April 2026: "Non-UK eSAF will be supplied to the EU NOT UK." Combined with the UK PtL sub-obligation being small in absolute tonnage (of order 0.4 Mt at 2040 against a UK Main obligation of ~2.3 Mt, §3.4), the finding is that the UK PtL sub-obligation is predominantly met by buy-out payments to the Treasury through the 2030s, not by procurement of PtL fuel. The mandate becomes a Treasury revenue line rather than a fuel-supply line.

Two structural consequences of the mechanism deserve mention. First, the arbitrage does not require any commercial malicious behaviour — it emerges from rational routing decisions by independent producers facing two markets with different clearing prices, and would emerge equally under any two-market structure with materially different non-compliance costs. Second, the arbitrage is not a bug in the UK Mandate's design; it is the direct consequence of pairing a statutory buy-out ceiling (which caps what the UK will pay) with a fungible molecule (e-SAF crosses borders freely under standard fuel-logistics arrangements). Correcting it would require either lifting the UK buy-out materially, introducing a UK-of-origin requirement in the mandate, or providing UK-specific procurement subsidies at prices above the buy-out through a Contract-for-Difference mechanism. Options are discussed in §4.2 and §4.5.

Reads directly against §3.4 (UK Mandate as e-SAF price ceiling), §4.2 (Revenue Certainty as the only mechanism that can outbid the EU on UK-origin e-SAF), and §4.5 (whether closing the arbitrage sits in the UK Treasury's discretionary spending priorities).

Exhibit 8. External triangulation against the IATA / Worley 2050 SAF supply outlook.

Segment	IATA / Worley 2050 outlook (SAF at Scale, April 2026)	This paper (WP04) at 2050
Global SAF capacity	22 Mt/yr (2030) → 413 Mt/yr (2050) at 16 % CAGR	Not modelled at global scale (WP04 scope: Greater Europe)
Pathway mix at 2050, share of global capacity	PtL / e-SAF ~43 % (fastest growth); FT / AtJ from residues ~26 %; HEFA ~15 % (dropping from ~75-80 % in 2030 as feedstock caps bind); FT-MSW ~7 %; other ~9 %	Diagnostic aviation-allocation for EU27+EFTA under NZ2050 shows PtL and FT together carrying most of the mandate-eligible stack; HEFA + Co-processing settle around ~15-20 %; HEFA-oilseed appears as an ineligible layer
Europe SAF supply 2050	≈ 42 Mt/yr	—

Segment	IATA / Worley 2050 outlook (SAF at Scale, April 2026)	This paper (WP04) at 2050
Greater Europe SAF required at mandate rate under NGFS Net Zero 2050, 2050	—	≈ 21 Mt/yr (70 % × 30 Mt EU27+EFTA aviation demand + 23.7 % × ~2 Mt UK aviation demand)
Greater Europe SAF required at mandate rate under NGFS Below 2°C, 2050	—	≈ 29 Mt/yr (70 % × 41 Mt EU27+EFTA + 23.7 % × ~4 Mt UK)
Greater Europe SAF required under Concawe reference demand at 2050 (WP03 basis)	—	≈ 40-60 Mt/yr (70 % × Concawe 57-87 Mt EU demand range)

Four structural readings follow.

First, the pathway mix aligns broadly between the two outlooks. PtL and e-SAF as the dominant late-century pathway, HEFA plateauing / declining in relative terms as its feedstock pool is exhausted by aviation and road demand competing for it, Fischer-Tropsch and Alcohol-to-Jet from residues carrying the middle of the century, MSW-derived FT emerging as a smaller but distinct route. That the two independently-constructed outlooks arrive at broadly consistent shape gives some confidence that WP04's pathway-substitution story is not idiosyncratic to the model.

Second, whether Europe's 42 Mt/yr supply figure is "sufficient" depends materially on which demand basis is used. Under NGFS Net Zero 2050 — this paper's core scenario — the Greater Europe mandate-driven requirement at 2050 is roughly 21 Mt/yr, i.e. approximately half the IATA/Worley Europe supply figure. Under NGFS Below 2°C the requirement is roughly 29 Mt/yr, still comfortably below 42 Mt. Under Concawe's reference demand (used in WP03), the requirement is 40-60 Mt/yr, at or above the IATA/Worley figure. WP03's "structurally short" reading applied against Concawe demand; WP04's NGFS-based reading implies Europe supply is broadly adequate to the mandate rate under either scenario — with the sub-mandate CO_2-bound within that (§3.1 and §3.3) as the substantive constraint.

Third, the global outlook of 413 Mt/yr in 2050 is materially larger than the Greater Europe mandate-driven requirement. Whether that global surplus is available to Europe through imports depends on the trade dimension deferred to §6 and Phase 2, and on the certification-and-logistics infrastructure that carries it. The IATA/Worley regional breakdown puts North America at 104 Mt/yr, North Asia at 66 Mt/yr, and Central & South America at 60 Mt/yr — none of which are far below Europe's 42 Mt/yr and each of which is far above its own domestic mandate-driven requirement.

Fourth, the two triangulation exercises together carry a policy implication worth naming. The claim that "Europe cannot meet the SAF mandate on domestic supply" is scenario-conditional and demand-basis-conditional; the claim that "the sub-mandate cannot be met on RFNBO-eligible CO_2 under NGFS scenarios" (§3.1) is robust across demand bases. The paper's structural finding is on the sub-mandate side, not on the total-mandate side.

Reads directly against §3.5 (pathway supply gap) and §3.1 (sub-mandate CO_2 constraint). External source: SAF at Scale — How Developers are moving from Concept to FID, Worley Consulting for IATA, presented 29 April 2026.

---

6. LIMITATIONS, SENSITIVITIES, AND AN INVITATION TO COOPERATE

This is a working draft and should be read as one. The quantities are scenario-conditional and public-source-derived; the production boundary stops at the neat certified SAF product, so blending, distribution and capital cost are excluded; NGFS scenarios drive the demand and resource envelopes rather than being asserted as forecasts; and the geographic focus is Greater Europe rather than the global aviation system some of the markets connect to. Each of these is a known edge of the model, named here so that it can be argued with.

Excluded scope

Several extensions to the model are Phase 2 work — flagged as tasks in the underlying task register — and are not carried in this draft:

- CAPEX and fixed OPEX per pathway (task #81). The cost curves in Exhibit 4 (p. 32) are variable-OPEX floors; full-cost curves adding CAPEX amortisation and fixed OPEX would raise every pathway's line materially and would change the absolute magnitudes of the §4.3 aggregate-spend read. Rank ordering across pathways would likely be preserved, but crossovers against fossil-jet-plus-carbon-levy would shift.
- Per-pathway industrial deployment ramps (task #69). The engine assumes a pathway produces whatever the merit-order allocates to it in a given year, subject to resource caps. Real plants ramp on FID-to-COD timelines and cannot be spun up in a single year. Deployment ramps would smooth the early-decade jumps in the pathway stack.
- CO_2 source split (task #71). The engine treats RFNBO-eligible CO_2 as a single pool from NGFS `Carbon Sequestration|CCS|Biomass|Energy|Supply|Liquids`. In reality DAC and biogenic point-source have different costs, eligibility properties and geographic distributions. A split would refine the §3.3 opportunity-cost mechanism.
- Trade and imports (task #89). The engine treats each region as a closed system. In reality e-SAF is globally tradable, and the EU is likely to import a material share of its e-SAF from lower-input-cost regions. This is the domestic-vs-imported split Exhibit 8 (p. 39) refers to.
- Shell demand alternative (task #88). NGFS demand contracts under Net Zero 2050; Shell's Aviation and energy to 2100 has aviation demand growing through to 2100. A Shell-alternative run would bracket the demand envelope alongside the NGFS reading.
- Intra-EEA / extra-EEA cost split (task #87). The engine applies the full ETS carbon levy to all EU fossil residual. In reality only intra-EEA emissions face ETS surrender; extra-EEA international departures fall under CORSIA, which prices non-compliance materially below ETS. Splitting the two would move §4.3's aggregate-spend numbers.
- Technology learning curves (task #82). Pathway intensities and unit costs are held at present-day values through to 2100. Real-world learning would reduce CAPEX and per-

tonne cost over time — particularly for PtL / e-SAF — and would shift the crossover economics against fossil-plus-levy.

Sectoral scope

This paper covers aviation only. Maritime (under FuelEU Maritime) and heavy industry — steel, cement, chemicals, refining, aluminium — are natural sequels using the same NGFS-driven framing, sequenced by ETS-price impact and decarbonisation cost stack (task #154 tracks the extension).

Sensitivity dimensions

Even within the paper's stated scope, seven dimensions materially move the results.

Aviation share for shared inputs. The engine computes aviation's share of shared NGFS resources from the scenario's own aviation-to-total-transport-liquids ratio (§2). Alternative allocation rules — a hardcoded 10-25 % share, or aviation's own draw of the specific input — produce SAF supply figures 30-50 % different, with the largest sensitivity on biomass and hydrogen budgets.

Post-2050 mandate ramp assumption. Between 2050 and 2080, the engine extends the ReFuelEU mandate from 70 % to 100 % at a constant rate. A hold-at-70 % alternative would sustain a fossil residual through to 2100; a slower ramp would push the aggregate-spend peak later. The rate of the extension materially moves the §4.3 aggregate-spend trajectory.

CO_2 removal-credit multiplier. The opportunity-cost mechanism prices CO_2 at `max(engineering floor, Price|Carbon × multiplier)`. A multiplier of 1.0 (full parity with removal credits) is the central case; 0.5 (certification-friction discount) and 1.2 (durability premium) are plausible sensitivities. The mechanism's direction is unchanged in either sensitivity; the magnitude scales.

Demand-envelope bracket NGFS vs Shell. NGFS demand contracts under Net Zero 2050; Shell's Aviation and energy to 2100 has demand growing. The compliance-gap picture sharpens materially under Shell — the same NGFS-derived supply ceilings leave a much larger fossil residual and a larger mandate non-compliance gap.

Pathway yields and per-tonne input intensities. The pathway network is inherited unchanged from WP03 Appendix A, which uses public-source-derived yields (JEC/JRC, GREET, Concawe conventions) at present-day plant designs. Learning-by-doing at scale would compress input intensities for the pre-commercial pathways — most materially for PtL, where a 20 % reduction in electrolyser electricity consumption per tonne of H_2 would reduce PtL's cost stack by roughly 15 % under NGFS Net Zero 2050 mid-century prices. Yield sensitivity does not, on the paper's slice, invert the rank ordering across pathways, but it compresses the gap between PtL and the biomass pathways in the second half of the century.

Merit-order assumptions. The engine assigns SAF supply to pathways in strict cost order — cheapest per tonne first — subject to resource caps. Task #124 flags a cost-aware refinement in which the marginal pathway is compared to fossil-plus-carbon-levy at each merit-order step; where SAF's marginal cost exceeds fossil-plus-levy, a economically rational producer would not build. That refinement would leave the pathway rankings unchanged but would truncate the pathway stack at the point where SAF supply becomes economically over-served. The direction of every §3 finding is preserved; the absolute magnitudes at high-carbon-price years contract by 10-20 %.

Technological breakthroughs outside the modelled pathway set. The engine treats the six pathways as fixed. A breakthrough elsewhere — direct hydrogen combustion in aircraft (a step out of the SAF category entirely), synthetic-biology-based bio-jet routes at commercially disruptive yields, direct-air-capture-plus-methanol-to-jet at a step-change cost — would displace part of the modelled pathway mix and change which constraint binds. None of these are impossible over the paper's 2025-2100 horizon; none are on public evidence expected on scale before 2035 at the earliest.

Trade, imports, and geopolitical constraints. The engine treats each mandate region as a closed system for supply. In reality e-SAF and biomass-derived SAF are globally tradable, and the EU is likely to import a material share of its e-SAF from lower-input-cost regions (task #89). Geopolitical friction on that trade — carbon-border measures, RFNBO-country certification schemes, protectionist responses to the concentration of certified capacity in specific supplier regions — would raise landed cost and shift the domestic-vs-imported balance discussed alongside Exhibit 8 (p. 39). These channels operate on cost levels and on supply-chain concentration risk (§4.4) rather than on the resource-constraint findings in §3, but they are material for a project-finance or lessor reading.

Resource-allocation framing — market-clearing, not political allocation

One assumption embedded across the paper deserves being named explicitly, because it will bear the weight of much of the reasonable criticism the paper will attract. The engine allocates shared inputs to aviation as if the pool clears in a market: aviation receives its scenario-derived share of biomass, hydrogen, CO_2, and electricity because the NGFS-liquids ratio says that is the size of aviation's competing draw. In practice, national and supranational allocation may operate through very different channels — political priority-setting between decarbonisation sectors, targeted subsidies routing resource pools toward strategic industries, military and dual-use demand carving out reserved shares of hydrogen and low-carbon electricity, long-term contractual commitments locking supply into off-take agreements before the market clears, and regional protectionism restricting cross-border resource flows to preserve domestic industrial capacity.

Each of those channels can move aviation's realised share of a given resource pool materially away from the market-clearing benchmark the model uses. Aviation could receive more than its NGFS-liquids ratio suggests (a strategic-industry allocation of hydrogen to aviation as a national-security priority; a Revenue Certainty instrument that under-writes an aviation-dedicated CO_2 off-take), or less (biomass reserved by policy for co-firing in dispatchable power generation; hydrogen prioritised for steel decarbonisation on carbon-leakage grounds). The finding-level results in §3 rest on the market-clearing benchmark. Wherever the political-allocation channel dominates, the aviation slice will diverge from what the paper models, and the direction of the divergence will depend on which sector the political process is trying to protect at the time.

This is not a limitation to hide behind. The paper's substantive contribution — that binding constraints migrate from hydrogen through biomass to CO_2 across the century — depends on the shape of the resource pool available to aviation, not on aviation's exact percentage share of it. A political re-allocation that gives aviation a larger biomass share still leaves biomass as a constraint of a broadly similar character in the mid-century, just at a different absolute magnitude. What a political-allocation channel does change is who bears the compliance risk in §3.1 and §3.4 — a mandate that has to be met from a resource pool the market did not allocate to aviation is a mandate paid for by whichever party (supplier, airline, taxpayer via buy-out) ends up carrying the price of that reallocation. That is the risk factor §4.5 names as cross-cutting policy risk, and it is the correct reading of the resource-allocation caveat rather than a challenge to the finding-level results themselves.

Demand-side channels outside the model's frame

Four channels through which aviation fuel demand or per-flight fuel-burn could evolve independently of NGFS's macro-feedback — and which would compound with the NGFS contraction to the extent they materialise — are not priced by the engine. They are named in §1 and referenced from Exhibit 1 (p. 25)'s caption:

- The concentration of the mandate-plus-ETS cost signal on intra-EEA flights. Extra-EEA international departures fall under CORSIA, which prices non-compliance well below ETS, so the demand-response signal for those flights is weaker.
- Air-traffic-control procedure optimisation. Efficiency gains from ATC (Single European Sky implementation, continuous-descent operations, sector-based flow management) reduce fuel consumption without changing demand.
- Airframe and engine improvements. Continued technological progress reduces per-flight fuel burn at any given demand level.
- Short-haul electrification. As battery-electric and hybrid-electric aircraft enter service on short-haul routes, the corresponding jet-fuel demand shifts out of the SAF/fossil-jet equation entirely.

Reality-vs-scenario uncertainty

NGFS scenarios describe orderly, globally-balanced paths to 1.5 °C (Net Zero 2050) or 2 °C (Below 2°C) at the whole-system level. They do not assert that the world is on such a path, nor do they specify aviation's SAF share within the resource pool they describe. Three possibilities are compatible with the empirical evidence available in 2026:

1. On-track with policy lag. SAF supply and demand are broadly on-track to what the scenario's aviation slice would require, with current policy delivery lagging modestly and expected to catch up as Revenue Certainty mechanisms mature and the pipeline of announced projects reaches COD.
2. Different-but-still-viable path. SAF supply and demand are on a different but still-viable path to the same climate objective, with other decarbonisation channels (electrification, demand-side measures, hydrogen for aviation, offsets) doing more of the substitution work than the scenario allocates to SAF.
3. Behind. SAF supply and demand are behind the trajectory the scenario would require, and the climate objective drifts as a result.

On the balance of evidence available in mid-2026 — the ~2-4 % UK SAF share actually being delivered in 2025-2026 against a scenario-consistent aviation allocation of much higher share, the concentration of UK non-HEFA FID targets in the 2027-2028 window with no closed FIDs to date, and the absence of large-scale Revenue Certainty commitments — (3) looks the most likely reading. The paper does not attempt to distinguish rigorously between the three; that distinction is the subject of a companion working paper (WP05, in scoping) which will use accumulating compliance data from 2025 onwards to measure where reality sits against both the scenario resource envelope and the mandate schedule.

Version stamps and reproducibility

Every result in the paper carries an engine version tag and a data version tag, stamped on the underlying database rows the results come from. The engine version identifies the merit-order rules, pathway intensities and mechanism configuration in force; the data version identifies the NGFS release, trend-derivation configuration and loading date. The results in this draft are stamped `v0.5.1 (2026-07-02)` for the data version. Any figure quoted from the calculator therefore traces to a specific model state, and later data versions will be visible in the version stamp on each result without changing the calculator's interface — the reproducibility of any figure cited in the paper does not depend on the paper itself being re-issued.

An invitation to cooperate

Above all this paper is an invitation to cooperate. The system it describes is being built now, by companies whose disclosures and presentations underpin much of the data here, and by regulators whose statutory instruments and compliance guidance set the frame the plants have to operate within. The model will only be as good as the scrutiny it survives.

Correction and contribution are welcomed from:

- Chemists and process engineers on pathway math, mass balances, LCA methodology, and the calculated design values that produce §3.6's dual-market credit-neutral reading. Susan Van Dyk (LSU AgCenter) has been a consistent source; the JEC and GREET communities are the natural comparators.
- SAF developers on project data — Altalto (Velocys), NEXTGEN SAF (University of Sheffield / E.ON), Sustainable Molecules (SuMo / SUEZ / COX), LanzaTech (DRAGON I and II), Zero Petroleum (Technip Energies), and the wider SAF Café presenter network.
- Policy analysts and regulators on mandate reading — the DfT Low Carbon Fuels Delivery Unit (saf-compliance@dft.gov.uk), EASA's reference-price working groups, SkyNRG's SAF outlook team, and the Innovate UK Business Connect Sustainable Aviation Fuel Innovation Programme that hosts the SAF Café series.

The point of making the assumptions explicit and the sources named is precisely to make that conversation possible. The LinkedIn post accompanying publication carries the same invitation as a public feedback channel. WP05 (in scoping) will use whatever accumulates in that channel plus the compliance data emerging from DfT and Member State reporting to reality-check the scenarios.

REFERENCES

Working Papers in the Carbon Risk & Capital Allocation Series. Working Paper 01 / 2026 — Carbon as a Capital Risk Factor: A Distribution-Based Framework for Strategic Investment

Decisions under the EU ETS. Working Paper 02 / 2026 — The EU ETS in Numbers: A Componentised Reconstruction of the Cap, 2008-2025. Working Paper 03 / 2026 — SAF — From Feedstock to Fuel: A Producer-Agnostic Process Network for the Financial Modelling of SAF and e-SAF. Working Paper 05 / 2026 (in scoping) — Empirical modelling against the NGFS scenarios (companion paper).

Regulation and statutory instruments. Regulation (EU) 2023/2405 of the European Parliament and of the Council of 18 October 2023 on ensuring a level playing field for sustainable air transport (ReFuelEU Aviation). Directive (EU) 2023/2413 of 18 October 2023 (RED III), Article 29a and Annex IX Parts A and B. Directive 2003/87/EC as amended, establishing the EU Emissions Trading Scheme (EU ETS Aviation). Regulation (EU) 2024/3012 establishing the EU Carbon Removal Certification Framework. Regulation (EU) 2023/955 establishing the Social Climate Fund. SI 2024/1187 — The Renewable Transport Fuel Obligations (Sustainable Aviation Fuel) Order 2024 (UK). SI 2020/1265 — Greenhouse Gas Emissions Trading Scheme Order 2020 (UK ETS). Department for Transport, SAF Mandate Compliance Guidance 2026 (published December 2025); DfT RTFO and SAF Mandate Technical Guidance. ICAO, Carbon Offsetting and Reduction Scheme for International Aviation (CORSIA).

NGFS scenario data. Network for Greening the Financial System, Phase 5 Scenarios V5.1 (2025) — GCAM 6.0 NGFS integrated assessment model outputs and Downscaling[GCAM 6.0 NGFS] regional splits.

Pathways (developer briefings from the UK Sustainable Aviation Fuel Café series, hosted by Innovate UK Business Connect under its Sustainable Aviation Fuel Innovation Programme). Altalto (Immingham FT-MSW) / Velocys — presenter February 2026. NEXTGEN-SAF, University of Sheffield / E.ON — presenter March 2026 (Sheffield demonstrator; North Tees commercial PCtL). Sustainable Molecules (SuMo Ltd) / SUEZ / COX / Velocys — presenter April 2026. LanzaTech DRAGON Projects (Port Talbot; Saltend) — presenter Greg Archer, April 2026. Zero Petroleum / Technip Energies — presenter April 2026 (Salt End, Hull, PtL Direct-FT). Series URL: https://iuk-business-connect.org.uk/events/sustainable-aviation-fuel-cafes/.

External SAF supply outlook (Exhibit 8 (p. 39)). Worley Consulting for IATA, SAF at Scale: How Developers are moving from Concept to FID, 29 April 2026.

Demand outlooks. Concawe Report 5/26 — Impact of Fit for 55 on EU transport energy demand, 2026 (WP03 demand basis, referenced in Exhibit 8 (p. 39) as the higher-demand comparator). Shell, Aviation and energy to 2100 (the growth-oriented demand-envelope alternative referenced in §6). Industry aviation-activity forecasts referenced in Exhibit 1 (p. 25)’s caption: IATA Twenty-Year Passenger and Cargo Forecast; ICAO Long-Term Aspirational Goals (LTAG); Airbus Global Market Forecast; Boeing Commercial Market Outlook.

Yields, energy, economics, and techno-economic assessment. Concawe Report 17/22 — E-Fuels: a techno-economic assessment, 2026. IEA Bioenergy Task 39 — demonstration facilities database and reports, 2024. IRENA, Green Hydrogen Cost Reduction, 2020. IRENA, Renewable Capacity Statistics 2026 and Renewable Energy Statistics 2025. Argonne National Laboratory, GREET model. JEC / JRC Well-to-Tank study for road, rail, marine and aviation fuels. Susan Van Dyk (LSU AgCenter) — SAF pathway data and life-cycle assessment (§6 correspondence).

Feedstock, biomass, LCA methodology. JRC133505 — Biomass supply and uses in the EU, 2023. JRC140117 — Biomass supply and uses in the EU, 2025 update. JRC142012 — Updating GHG default values (RED II Annex V and VI), 2026. Department for Transport, RTFO and SAF Man-

date list of feedstock materials (HTML publication, updated 19 February 2026; URL: https://www.gov.uk/government/publications/rtfo-and-saf-mandate-feedstock-materials-used-for-creating-low-carbon-fuels). Department for Transport, Use of voluntary schemes as evidence of RTFO and SAF Mandate compliance (URL: https://www.gov.uk/government/publications/use-of-voluntary-schemes-as-evidence-of-rtfo-and-saf-mandate-compliance). Department for Transport, RTFO and SAF Mandate Standard Data workbook (December 2024, effective January 2025).

Costs and market. EASA, 2025 Aviation Fuels Reference Prices for ReFuelEU Aviation — briefing note and methodology document, 2026. The eFuels Gap, SAF Congress presentation, 2026. SAS (Scandinavian Airlines), Aviation Insights — the need for e-SAF in Scandinavia, 2026. CE Delft, Sustainable Carbon for aviation and shipping, 2025.

Scientific reference. Intergovernmental Panel on Climate Change (IPCC) — Sixth Assessment Report (Working Groups I and III) for the 1.5 °C and 2 °C carbon-budget references anchoring the NGFS scenarios.

Model source. Every figure quoted from the calculator carries an engine-version and data-version tag traceable to the underlying database rows. Results in this draft are stamped `v0.5.1 (2026-07-02)`.

APPENDIX · NGFS TRENDS SLICE

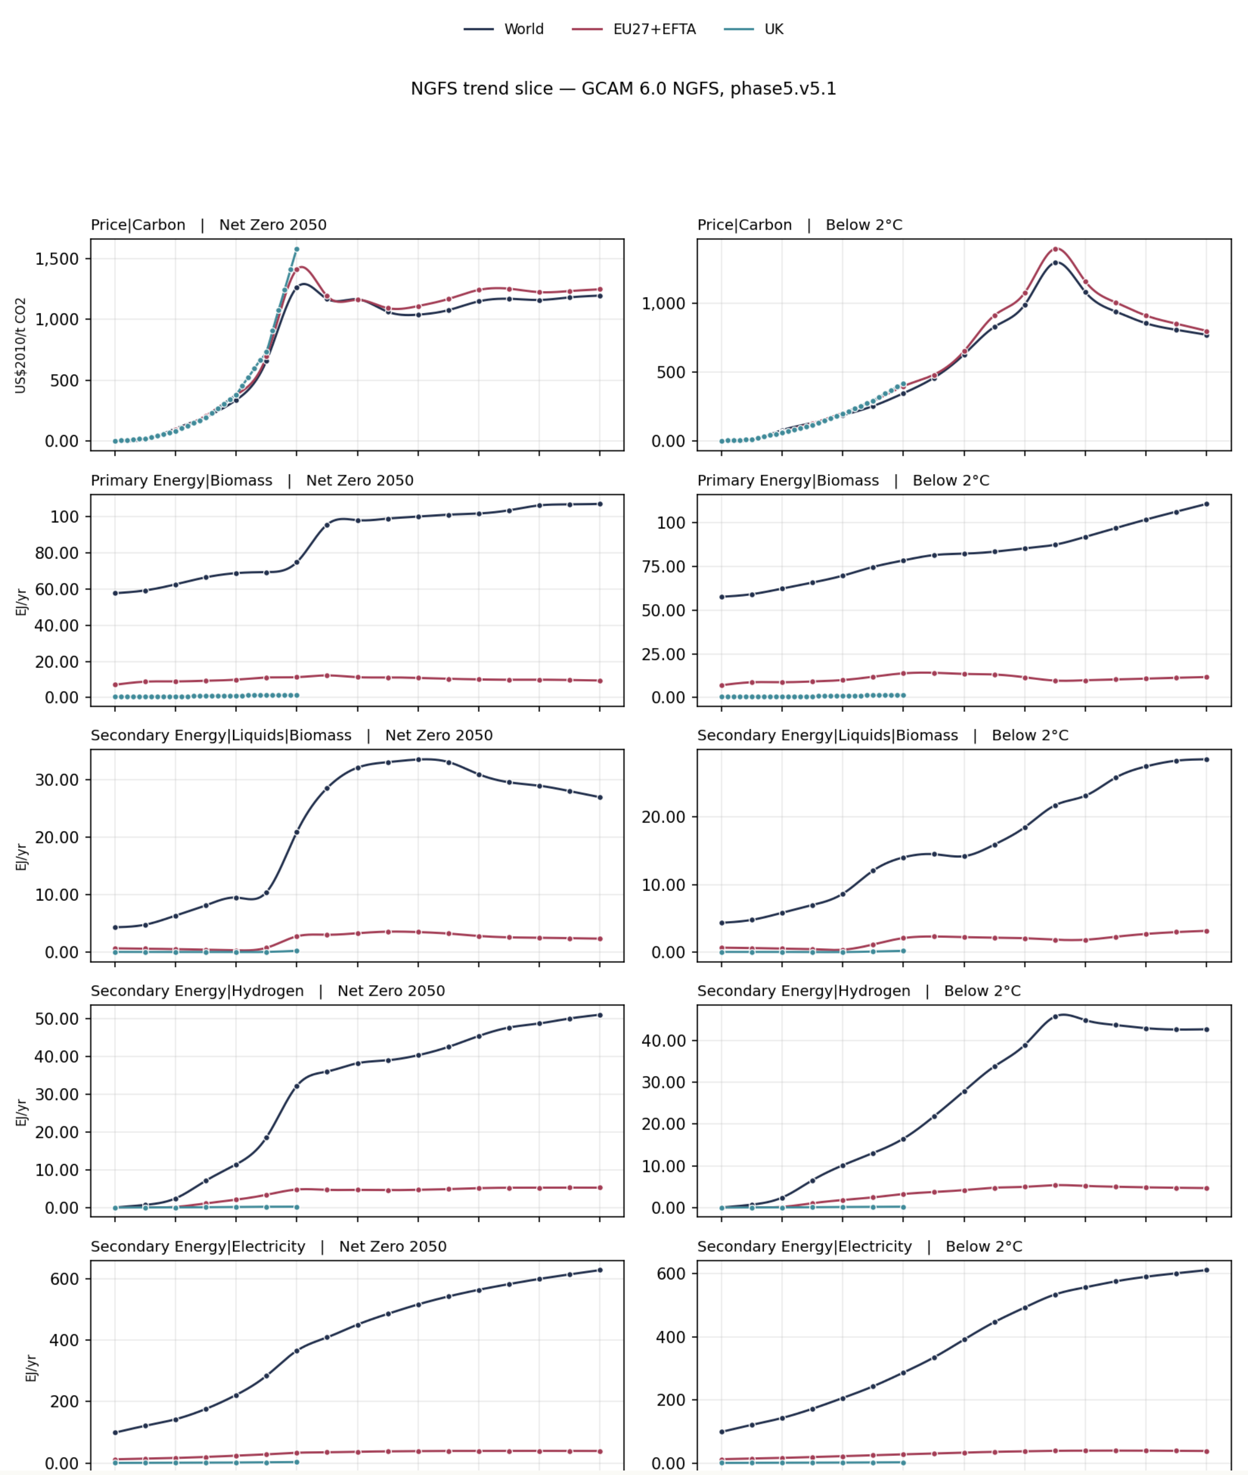

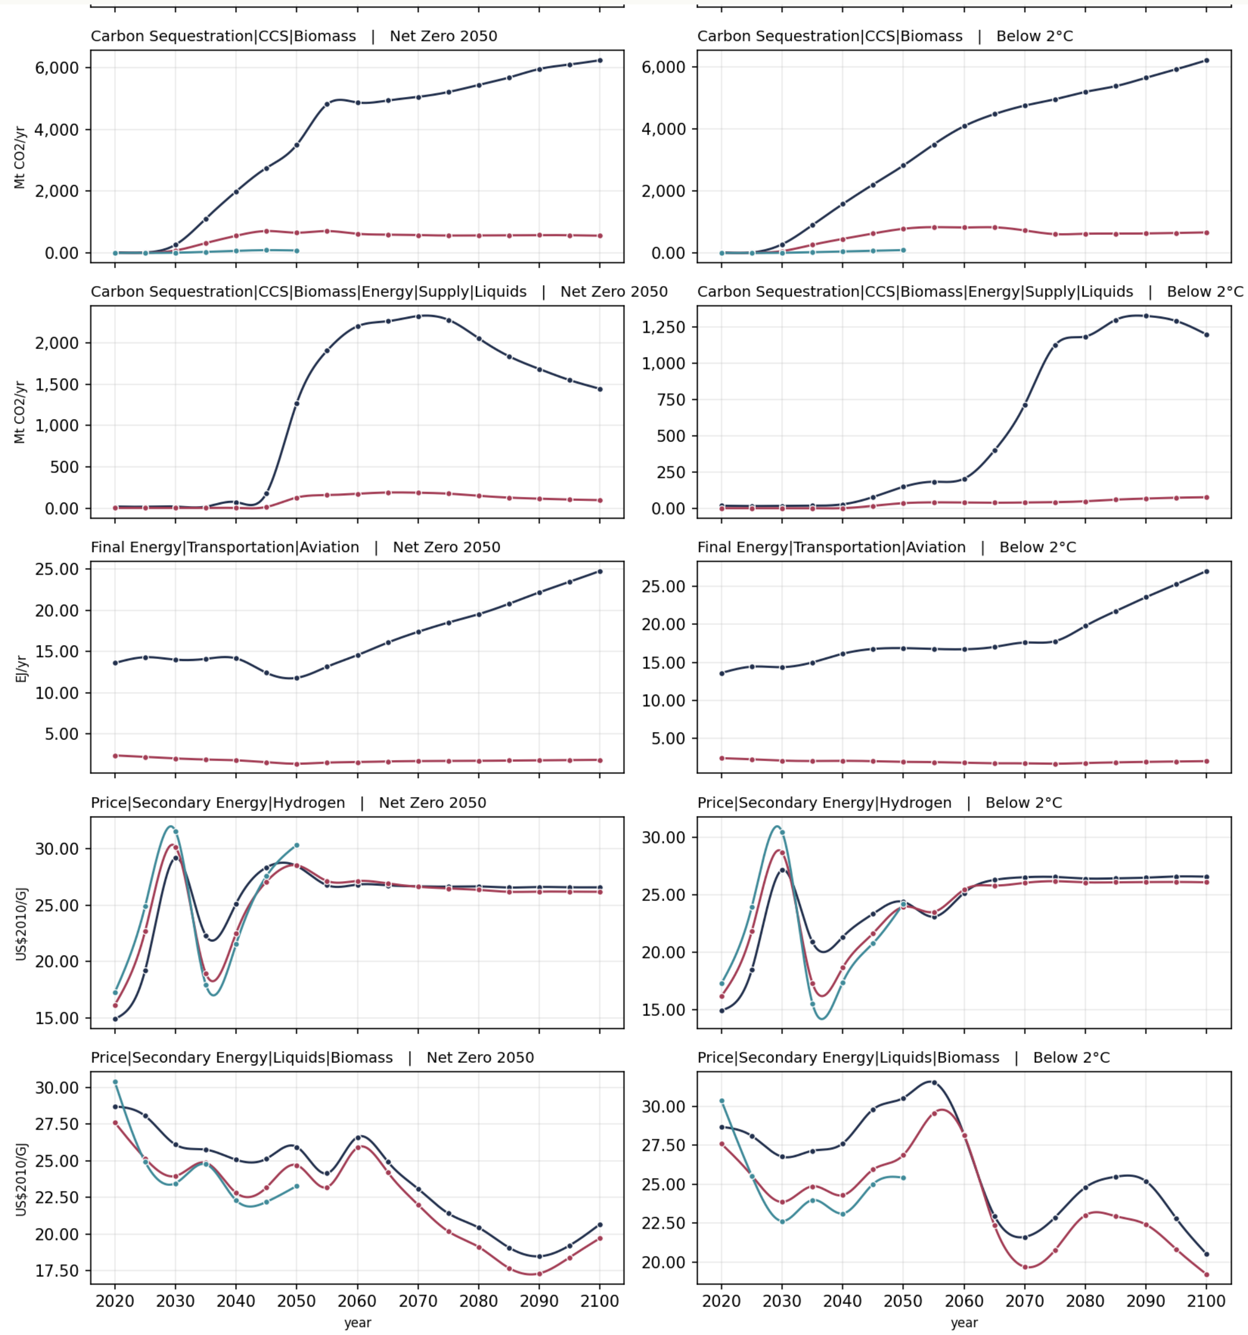

10-yr 2060-2100). World + EU27+EFTA from native GCAM 6.0 NGFS; UK from Downscaling[GCAM 6.0 NGFS] with canonical_label=GBR. NGFS V5.1 truncates downscaled GBR at 2050 — UK lines therefore stop at 2050 even w

The appendix figure shows the ten NGFS variables the engine consumes (§2), rendered as daily trends from cubic-spline interpolation between the NGFS reference points, for the three aggregated regions (World, EU27+EFTA, UK) under both scenarios. Reference points are shown as dots for verification. UK panels stop at 2050 because NGFS V5.1 truncates the downscaled GBR series at that year; the UK panel therefore does not appear for variables not published in the downscaled product (aviation demand, BECCS|Liquids). Reads directly against §2 (data ingest) and §3.1 (BECCS|Liquids trajectory).

The views expressed are solely those of the author and do not represent any affiliated institution. This paper is part of ongoing independent research and should not be interpreted as investment, engineering or regulatory advice.